

Firepower Management Center Configuration Guide, Version 6.2

First Published: 2017-01-23

Last Modified: 2019-04-25

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED “AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2017–2019 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1 **Getting Started With Firepower** 1

Quick Start: Basic Setup 1

Installing and Performing Initial Setup on Physical Appliances 2

Deploying Virtual Appliances 2

Logging In for the First Time 3

Setting Up Basic Policies and Configurations 4

Firepower Devices 5

Firepower Features 6

Appliance and System Management Features 6

High Availability and Scalability Features by Platform 7

Features for Detecting, Preventing, and Processing Potential Threats 9

Integration with External Tools 10

Switching Domains on the Firepower Management Center 10

The Context Menu 11

Firepower Online Help and Documentation 12

Top-Level Documentation Listing Pages for FMC Deployments 13

License Statements in the Documentation 14

Supported Devices Statements in the Documentation 14

Access Statements in the Documentation 14

Firepower System IP Address Conventions 15

Additional Resources 15

PART I **Your User Account** 17

CHAPTER 2 **Logging into the Firepower System** 19

Firepower System User Accounts 19

Firepower System User Interfaces 21
Web Interface Considerations 23
Session Timeout 24
Logging Into the Firepower Management Center Web Interface 24
Logging Into the Web Interface of a 7000 or 8000 Series Device 25
Logging Into the Firepower Management Center with CAC Credentials 26
Logging Into a 7000 or 8000 Series Device with CAC Credentials 26
Logging Into the CLI on 7000/8000 Series, ASA FirePOWER, and NGIPSv Devices 27
Logging Into the Command Line Interface on Firepower Threat Defense Devices 28
Logging Out of a Firepower System Web Interface 29

CHAPTER 3 **Specifying User Preferences** 31
User Preferences Introduction 31
Changing Your Password 31
Changing an Expired Password 32
Specifying Your Home Page 32
Configuring Event View Settings 33
Event View Preferences 33
File Download Preferences 34
Default Time Windows 35
Default Workflows 37
Setting Your Default Time Zone 37
Specifying Your Default Dashboard 37

PART II **Firepower System Management** 39

CHAPTER 4 **Firepower System User Management** 41
User Roles 41
Predefined User Roles 42
Custom User Roles 43
Example: Custom User Roles and Access Control 44
User Account Privileges 44
Overview Menu 44
Analysis Menu 45

Policies Menu 49
Devices Menu 51
Object Manager Menu 52
Cisco AMP 53
Deploy Configuration to Devices 53
System Menu 53
REST VDI Menu 55
Help Menu 55
Managing User Roles 55
Activating and Deactivating User Roles 56
Creating Custom User Roles 57
Copying User Roles 57
Editing Custom User Roles 58
User Role Escalation 59
Setting the Escalation Target Role 59
Configuring a Custom User Role for Escalation 60
Escalating Your User Role 60
User Accounts 61
Managing User Accounts 61
Creating a User Account 62
Editing a User Account 63
Assigning User Roles in Multiple Domains 63
Converting a User from Internal to External Authentication 64
User Account Login Options 64
Command Line Access Levels 66
Creating Local User Accounts for the FTD CLI 67
Firepower System User Authentication 69
Internal Authentication 70
External Authentication 71
LDAP Authentication 72
Required Information for Creating LDAP Authentication Objects 72
CAC Authentication 74
Configuring CAC Authentication 74
Creating Basic LDAP Authentication Objects 76

Creating Advanced LDAP Authentication Objects 78
LDAP Authentication Server Fields 81
Identifying the LDAP Authentication Server 82
LDAP-Specific Fields 83
Configuring LDAP-Specific Parameters 85
LDAP Group Fields 87
Configuring Access Rights by Group 88
LDAP Shell Access Fields 89
Configuring LDAP Shell Access 90
Testing LDAP Authentication Connections 91
Troubleshooting LDAP Authentication Connections 92
RADIUS Authentication 93
Creating RADIUS Authentication Objects 94
Configuring RADIUS Connection Settings 96
Configuring RADIUS User Roles 98
Configuring RADIUS Shell Access 99
Defining Custom RADIUS Attributes 100
Testing RADIUS Authentication Connections 101
Single Sign-on (SSO) 102
Configuring SSO 102

CHAPTER 5 **Licensing the Firepower System** 105
About Firepower Licenses 105
Requirements and Prerequisites for Licensing 106
License Requirements for Firepower Management Center 106
Firepower Management Center Virtual Licenses 106
Evaluation License Caveats 106
Smart vs. Classic Licenses 107
License Firepower Threat Defense Devices (FTD) 107
Smart Software Manager (CSSM) 108
Periodic Communication with the License Authority 108
Service Subscriptions for FTD Features 108
FTD License Types and Restrictions 109
Base Licenses 110

Malware Licenses for Firepower Threat Defense Devices 111
Threat Licenses 111
URL Filtering Licenses for Firepower Threat Defense Devices 112
Licensing for Export-Controlled Functionality 112
Licensing for FTD Clusters 113
Create a Smart Account to Hold Your Licenses 113
How to Configure Smart Licensing with Direct Internet Access 114
Obtain a Product License Registration Token for Smart Licensing 115
Register Smart Licenses 116
Smart Software Manager On-Prem Overview 117
How to Deploy Smart Software Manager On-Prem 117
Assign Licenses to Multiple Managed Devices 119
View FTD Licenses and License Status 120
FTD License Status 120
Move or Remove Licenses from FTD Devices 121
Transfer FTD Licenses to a Different Firepower Management Center 122
If FTD License Status is Out of Compliance 122
Deregister a Firepower Management Center from the Cisco Smart Software Manager 122
Synchronize a Firepower Management Center with the Cisco Smart Software Manager 123
Troubleshoot FTD Licensing 123
License Classic Devices (Firepower 7000/8000 Series, ASA FirePOWER, and NGIPSv) 124
Product License Registration Portal 124
Service Subscriptions for Firepower Features (Classic Licensing) 124
Classic License Types and Restrictions 125
Protection Licenses 126
Control Licenses 127
URL Filtering Licenses for Classic Devices 128
Malware Licenses for Classic Devices 128
VPN Licenses for 7000 and 8000 Series Devices 129
Classic Licenses in Device Stacks and High-Availability Pairs 129
View Your Classic Licenses 130
Identify the License Key 130
Generate a Classic License and Add It to the Firepower Management Center 131
How to Convert a Classic License for Use on an FTD Device 132

Assign Licenses to Managed Devices from the Device Management Page 134
Additional Information about Firepower Licensing 135

CHAPTER 6 System Updates 137
About System Updates 137
Requirements and Prerequisites for System Updates 138
Guidelines and Limitations for System Updates 139
Upgrade System Software 139
Update the Vulnerability Database (VDB) 139
Manually Update the VDB 140
Schedule VDB Updates 141
Update the Geolocation Database 141
Manually Update the GeoDB (Internet Connection) 142
Manually Update the GeoDB (No Internet Connection) 142
Schedule GeoDB Updates 143
Update Intrusion Rules 143
Update Intrusion Rules One-Time Manually 145
Update Intrusion Rules One-Time Automatically 145
Schedule Intrusion Rule Updates 146
Best Practices for Importing Local Intrusion Rules 147
Import Local Intrusion Rules 148
Rule Update Log 149
Intrusion Rule Update Log Table 149
Viewing the Intrusion Rule Update Log 150
Fields in an Intrusion Rule Update Log 150
Viewing Details of the Intrusion Rule Update Import Log 152
Maintain Your Air-Gapped Deployment 153

CHAPTER 7 Backup and Restore 155
About Backup and Restore 155
Requirements for Backup and Restore 156
Guidelines and Limitations for Backup and Restore 157
Best Practices for Backup and Restore 158
Backing Up FMCs or Managed Devices 161

Back up the FMC 161
Back up a Device from the FMC 163
Back up a 7000/8000 Series Device Locally 163
Create a Backup Profile 165
Restoring FMCs and Managed Devices 166
Restore an FMC from Backup 166
Restore a 7000/8000 Series Device from Backup 167
Manage Backups and Remote Storage 168
Backup Storage Locations 169

CHAPTER 8 **Configuration Import and Export** 173
About Configuration Import/Export 173
Configurations that Support Import/Export 173
Special Considerations for Configuration Import/Export 174
Requirements and Prerequisites for Configuration Import/Export 175
Exporting Configurations 175
Importing Configurations 176
Import Conflict Resolution 177

CHAPTER 9 **Task Scheduling** 179
About Task Scheduling 179
Requirements and Prerequisites for Task Scheduling 179
Configuring a Recurring Task 180
Scheduled Backups 181
Schedule FMC Backups 181
Schedule Local 7000 & 8000 Series Device Backups 182
Configuring Certificate Revocation List Downloads 182
Automating Policy Deployment 183
Nmap Scan Automation 184
Scheduling an Nmap Scan 184
Automating Report Generation 185
Specify Report Generation Settings for a Scheduled Report 186
Automating Firepower Recommendations 187
Software Update Automation 188

Automating Software Downloads 189
Automating Software Pushes 190
Automating Software Installs 190
Vulnerability Database Update Automation 191
Automating VDB Update Downloads 192
Automating VDB Update Installs 192
Automating URL Filtering Updates Using a Scheduled Task 193
Scheduled Task Review 194
Task List Details 195
Viewing Scheduled Tasks on the Calendar 195
Editing Scheduled Tasks 196
Deleting Scheduled Tasks 196

CHAPTER 10 **Data Storage** 197
Data Stored on the FMC 197
Purging Data from the FMC Database 198
External Data Storage 199

CHAPTER 11 **Firepower Management Center High Availability** 201
About Firepower Management Center High Availability 201
Roles v. Status in Firepower Management Center High Availability 202
Event Processing on Firepower Management Center High Availability Pairs 202
AMP Cloud Connections and Malware Information 202
URL Filtering and Security Intelligence 202
User Data Processing During Firepower Management Center Failover 203
Configuration Management on Firepower Management Center High Availability Pairs 203
Firepower Management Center High Availability Behavior During a Backup 203
Firepower Management Center High Availability Split-Brain 203
Upgrading Firepower Management Centers in a High Availability Pair 204
Troubleshooting Firepower Management Center High Availability 204
Requirements for Firepower Management Center High Availability 205
Hardware Requirements 205
Software Requirements 206
License Requirements for FMC High Availability Configurations 206

Prerequisites for Firepower Management Center High Availability 207
Establishing Firepower Management Center High Availability 208
Viewing Firepower Management Center High Availability Status 209
Configuration Data Synced between Firepower Management Centers during High Availability 210
Configuring External Access to the FMC Database in a High Availability Pair 211
Using CLI to Resolve Device Registration in Firepower Management Center High Availability 211
Switching Peers in a Firepower Management Center High Availability Pair 212
Pausing Communication Between Paired Firepower Management Centers 212
Restarting Communication Between Paired Firepower Management Centers 212
Changing the IP address of a Firepower Management Center in a High Availability Pair 213
Disabling Firepower Management Center High Availability 213
Replacing FMCs in a High Availability Pair 214
Replace a Failed Primary FMC (Successful Backup) 215
Replace a Failed Primary FMC (Unsuccessful Backup) 216
Replace a Failed Secondary FMC (Successful Backup) 217
Replace a Failed Secondary FMC (Unsuccessful Backup) 218

CHAPTER 12 **Device Management Basics** 219
About Device Management 219
About the Firepower Management Center and Device Management 219
What Can Be Managed by a Firepower Management Center? 220
Beyond Policies and Events 220
About Device Management Interfaces 221
Management Interfaces on 221
Management Interface Support Per Device Model 221
Network Routes on Device Management Interfaces 222
NAT Environments 223
Management and Event Traffic Channel Examples 224
Requirements and Prerequisites for Device Management 226
Complete the FTD Initial Configuration Using the CLI 226
Add a Device to the FMC 229
Delete a Device from the FMC 231
Add a Device Group 232
Configure Device Settings 233

Managing System Shut Down 233
Edit Management Settings 234
Update the Hostname or IP Address in FMC 234
Modify Management Interfaces at the CLI 235
Edit General Settings 240
Edit License Settings 240
Edit Advanced Settings 241
Configure Automatic Application Bypass 241
Inspect Local Router Traffic 242
Configure Fastpath Rules (8000 Series) 243
Change the Manager for the Device 244
Reestablish the Management Connection if You Change the FMC IP Address 244
Identify a New FMC 245
Switch from Firepower Device Manager to FMC 246
Switch from FMC to Firepower Device Manager 247
Viewing Device Information 248
Device Management Page Information 249
General Information 249
License Information 250
System Information 250
Health Information 250
Management Information 251
Advanced Settings 251

PART III **System Monitoring and Troubleshooting** 253

CHAPTER 13 **Dashboards** 255
About Dashboards 255
Firepower System Dashboard Widgets 256
Widget Availability 256
Dashboard Widget Availability by User Role 257
Predefined Dashboard Widgets 258
The Appliance Information Widget 258
The Appliance Status Widget 258

The Correlation Events Widget 259
The Current Interface Status Widget 259
The Current Sessions Widget 260
The Custom Analysis Widget 260
The Disk Usage Widget 264
The Interface Traffic Widget 264
The Intrusion Events Widget 265
The Network Compliance Widget 266
The Product Licensing Widget 266
The Product Updates Widget 266
The RSS Feed Widget 267
The System Load Widget 267
The System Time Widget 267
The White List Events Widget 267
Managing Dashboards 268
Adding a Dashboard 269
Adding Widgets to a Dashboard 269
Configuring Widget Preferences 270
Creating Custom Dashboards 270
Custom Dashboard Options 270
Customizing the Widget Display 271
Editing Dashboards Options 272
Modifying Dashboard Time Settings 272
Renaming a Dashboard 273
Viewing Dashboards 274

CHAPTER 14 **Health Monitoring** 275
Requirements and Prerequisites for Health Monitoring 275
About Health Monitoring 275
Health Modules 277
Configuring Health Monitoring 282
Health Policies 283
Default Health Policy 283
Creating Health Policies 283

Applying Health Policies 284
Editing Health Policies 285
Deleting Health Policies 286
The Health Monitor Blocklist 286
Blocklisting Appliances 287
Blocklisting Health Policy Modules 288
Health Monitor Alerts 288
Health Monitor Alert Information 288
Creating Health Monitor Alerts 289
Editing Health Monitor Alerts 290
Deleting Health Monitor Alerts 290
Using the Health Monitor 291
Health Monitor Status Categories 291
Viewing Appliance Health Monitors 292
Running All Modules for an Appliance 293
Running a Specific Health Module 294
Generating Health Module Alert Graphs 294
Health Event Views 295
Viewing Health Events 295
Viewing Health Events by Module and Appliance 296
Viewing the Health Events Table 296
Hardware Alert Details for 7000 and 8000 Series Devices 297
The Health Events Table 298

CHAPTER 15 **Monitoring the System** 301
About System Statistics 301
The Host Statistics Section 301
The Disk Usage Section 302
The Processes Section 302
Process Status Fields 302
System Daemons 304
Executables and System Utilities 305
The SFDataCorrelator Process Statistics Section 308
The Intrusion Event Information Section 308

Viewing System Statistics 309

CHAPTER 16 **Troubleshooting the System** 311

First Steps for Troubleshooting 311

System Messages 311

Message Types 312

Message Management 313

View Basic System Information 314

View Appliance Information 314

Managing System Messages 314

Viewing Deployment Messages 315

Viewing Health Messages 316

Viewing Task Messages 316

Managing Task Messages 317

Configuring Notification Behavior 317

Memory Usage Thresholds for Health Monitor Alerts 318

Disk Usage and Drain of Events Health Monitor Alerts 319

Health Monitor Reports for Troubleshooting 322

Producing Troubleshooting Files for Specific System Functions 323

Downloading Advanced Troubleshooting Files 324

Advanced Troubleshooting for the Firepower Threat Defense Device 325

Using the Firepower Threat Defense CLI from the Web Interface 325

Packet Tracer Overview 326

Use the Packet Tracer 326

Packet Capture Overview 327

Use the Capture Trace 329

Feature-Specific Troubleshooting 330

PART IV **Deployment Management** 333

CHAPTER 17 **Domain Management** 335

Introduction to Multitenancy Using Domains 335

Domains Terminology 336

Domain Properties 337

Requirements and Prerequisites for Domains 338
Managing Domains 338
Creating New Domains 339
Moving Data Between Domains 340
Moving Devices Between Domains 341

CHAPTER 18 **Policy Management** 343
Requirements and Prerequisites for Policy Management 343
Policy Deployment 344
Best Practices for Deploying Configuration Changes 344
Deploy Configuration Changes 346
Redeploy Existing Configurations to a Device 348
Snort® Restart Scenarios 348
Inspect Traffic During Policy Apply 349
Snort® Restart Traffic Behavior 350
Configurations that Restart the Snort Process When Deployed or Activated 352
Changes that Immediately Restart the Snort Process 354
Policy Comparison 355
Comparing Policies 355
Policy Reports 356
Generating Current Policy Reports 356
Out-of-Date Policies 357
Performance Considerations for Limited Deployments 358
Discovery Without Intrusion Prevention 358
Intrusion Prevention Without Discovery 359

CHAPTER 19 **Rule Management: Common Characteristics** 361
Requirements and Prerequisites for Rule Management 361
Introduction to Rules 361
Rule Condition Types 363
Rule Condition Mechanics 365
Interface Conditions 366
Configuring Interface Conditions 367
Network Conditions 368

Configuring Network Conditions 369
Tunnel Endpoint Conditions 370
Configuring Tunnel Endpoint Conditions 371
VLAN Conditions 372
Port and ICMP Code Conditions 373
Configuring Port Conditions 374
Encapsulation Conditions 375
Application Conditions (Application Control) 375
Configuring Application Conditions and Filters 376
Application Characteristics 378
Best Practices for Application Control 379
Best Practices for Configuring Application Control 381
Application-Specific Notes and Limitations 382
Troubleshoot Application Control Rules 383
URL Conditions (URL Filtering) 384
User, Realm, and ISE Attribute Conditions (User Control) 384
User Control Prerequisites 385
Configuring User and Realm Conditions 386
Configuring ISE Attribute Conditions 387
Troubleshoot User Control 387
Custom SGT Conditions 389
ISE SGT vs Custom SGT Rule Conditions 389
Autotransition from Custom SGTs to ISE SGTs 389
Configuring Custom SGT Conditions 390
Troubleshooting Custom SGT Conditions 390
Searching for Rules 390
Filtering Rules by Device 391
Rule and Other Policy Warnings 392

CHAPTER 20 **Reusable Objects** 393
Introduction to Reusable Objects 394
The Object Manager 396
Editing Objects 396
Viewing Objects and Their Usage 397

Filtering Objects or Object Groups 397
Object Groups 398
Grouping Reusable Objects 398
Object Overrides 399
Managing Object Overrides 401
Allowing Object Overrides 401
Adding Object Overrides 401
Editing Object Overrides 402
Network Objects 402
Creating Network Objects 403
Port Objects 404
Creating Port Objects 405
Tunnel Zones 406
Application Filters 406
VLAN Tag Objects 406
Creating VLAN Tag Objects 406
Security Group Tag Objects 407
Creating Security Group Tag Objects 407
URL Objects 408
Creating URL Objects 409
Geolocation Objects 409
Creating Geolocation Objects 409
Interface Objects: Interface Groups and Security Zones 410
Creating Security Zone and Interface Group Objects 411
Variable Sets 412
Variable Sets in Intrusion Policies 413
Variables 413
Predefined Default Variables 414
Network Variables 416
Port Variables 417
Advanced Variables 418
Variable Reset 419
Adding Variables to Sets 419
Nesting Variables 421

Managing Variable Sets 422
Creating Variable Sets 423
Managing Variables 424
Adding Variables 425
Editing Variables 426
Security Intelligence Lists and Feeds 427
How to Modify Security Intelligence Objects 428
Global and Domain Security Intelligence Lists 428
Security Intelligence Lists and Multitenancy 429
Add Entries to Global Security Intelligence Lists 430
Delete Entries from Global Security Intelligence Lists 431
List and Feed Updates for Security Intelligence 432
Changing the Update Frequency for Security Intelligence Feeds 432
Custom Security Intelligence Lists and Feeds 432
Custom Lists and Feeds: Requirements 432
URL Lists and Feeds: URL Syntax and Matching Criteria 433
Custom Security Intelligence Feeds 434
Custom Security Intelligence Lists 436
Sinkhole Objects 438
Creating Sinkhole Objects 438
File Lists 438
Source Files for File Lists 439
Adding Individual SHA-256 Values to File Lists 440
Uploading Individual Files to File Lists 440
Uploading Source Files to File Lists 441
Editing SHA-256 Values in File Lists 442
Downloading Source Files from File Lists 443
Cipher Suite Lists 443
Creating Cipher Suite Lists 444
Distinguished Name Objects 444
Creating Distinguished Name Objects 446
PKI Objects 447
Internal Certificate Authority Objects 448
CA Certificate and Private Key Import 448

Importing a CA Certificate and Private Key 449
Generating a New CA Certificate and Private Key 449
New Signed Certificates 450
Creating an Unsigned CA Certificate and CSR 450
Uploading a Signed Certificate Issued in Response to a CSR 450
CA Certificate and Private Key Downloads 451
Downloading a CA Certificate and Private Key 451
Trusted Certificate Authority Objects 452
Trusted CA Object 452
Adding a Trusted CA Object 453
Certificate Revocation Lists in Trusted CA Objects 453
Adding a Certificate Revocation List to a Trusted CA Object 454
External Certificate Objects 454
Adding External Certificate Objects 455
Internal Certificate Objects 455
Adding Internal Certificate Objects 456
Certificate Enrollment Objects 456
Adding Certificate Enrollment Objects 458
Certificate Enrollment Object SCEP Options 459
Certificate Enrollment Object Certificate Parameters 460
Certificate Enrollment Object Key Options 460
Certificate Enrollment Object Revocation Options 461
SLA Monitor Objects 461
Prefix Lists 463
Configure IPv6 Prefix List 463
Configure IPv4 Prefix List 464
Route Maps 464
Access List 467
Configure Extended ACL Objects 468
Configure Standard ACL Objects 469
AS Path Objects 470
Community Lists 470
Policy Lists 472
VPN Objects 473

Firepower Threat Defense IKE Policies 473
Configure IKEv1 Policy Objects 473
Configure IKEv2 Policy Objects 475
Firepower Threat Defense IPsec Proposals 476
Configure IKEv1 IPsec Proposal Objects 477
Configure IKEv2 IPsec Proposal Objects 477
Address Pools 478
FlexConfig Objects 479

CHAPTER 21 **Firepower Threat Defense Certificate-Based Authentication 481**
Requirements and Prerequisites for FTD Certificate-Based Authentication 481
Firepower Threat Defense VPN Certificate Guidelines and Limitations 482
Managing Firepower Threat Defense VPN Certificates 482
Installing a Certificate Using Self-Signed Enrollment 483
Installing a Certificate Using SCEP Enrollment 484
Installing a Certificate Using Manual Enrollment 485
Installing a Certificate by Importing a PKCS12 File 486
Troubleshooting Firepower Threat Defense VPN Certificates 487

PART V **Classic Device Configuration Basics 489**

CHAPTER 22 **Classic Device Management Basics 491**
Requirements and Prerequisites for Classic Device Management 491
Remote Management Configuration (Classic Devices) 491
Configuring Remote Management on a Managed Device 492
Editing Remote Management on a Managed Device 493
Changing the Management Port 493
Interface Configuration Settings 494
The Physical Hardware View 494
The Interfaces Page 494
Interface Icons 496
Using the Physical Hardware View 496
Configuring Sensing Interfaces 497
Configuring HA Link Interfaces 498

Disabling Interfaces 499
Managing Cisco ASA FirePOWER Interfaces 500
MTU Ranges for 7000 and 8000 Series Devices and NGIPSv 500
Synchronizing Security Zone Object Revisions 501

CHAPTER 23 **IPS Device Deployments and Configuration** 503
Introduction to IPS Device Deployment and Configuration 503
License Requirements for IPS Device Deployment 503
Requirements and Prerequisites for IPS Device Deployment 503
Passive IPS Deployments 504
Passive Interfaces on the Firepower System 504
Configuring Passive Interfaces 505
Inline IPS Deployments 506
Inline Interfaces on the Firepower System 507
Configuring Inline Interfaces 508
Inline Sets 509
Viewing Inline Sets 510
Adding Inline Sets 511
Advanced Inline Set Options 512
Configuring Advanced Inline Set Options 513
Deleting Inline Sets 514

PART VI **Classic Device High Availability and Scalability** 515

CHAPTER 24 **7000 and 8000 Series Device High Availability** 517
About 7000 and 8000 Series Device High Availability 517
Device High Availability Requirements 518
Device High Availability Failover and Maintenance Mode 518
Configuration Deployment and Upgrade Behavior for High-Availability Pairs 519
Deployment Types and Device High Availability 519
7000/8000 Series High Availability Configuration 521
Establishing Firepower 7000/8000 Series High Availability 521
Editing Device High Availability 522
Configuring Individual Devices in a High-Availability Pair 523

Configuring Individual Device Stacks in a High-Availability Pair 523
Configuring Interfaces on a Device in a High-Availability Pair 524
Switching the Active Peer in a Device High-Availability Pair 525
Placing a High-Availability Peer into Maintenance Mode 525
Replacing a Device in a Stack in a High-Availability Pair 526
Device High Availability State Sharing 526
Establishing Device High-Availability State Sharing 528
Device High Availability State Sharing Statistics for Troubleshooting 529
Viewing Device High Availability State Sharing Statistics 531
Separating Device High-Availability Pairs 532

CHAPTER 25 **8000 Series Device Stacking 533**
About Device Stacks 533
Device Stack Configuration 535
Establishing Device Stacks 536
Editing Device Stacks 537
Replacing a Device in a Stack 537
Replacing a Device in a Stack in a High-Availability Pair 538
Configuring Individual Devices in a Stack 539
Configuring Interfaces on a Stacked Device 539
Separating Stacked Devices 540
Replacing a Device in a Stack 541

PART VII **Firepower Threat Defense Getting Started 543**

CHAPTER 26 **Transparent or Routed Firewall Mode for Firepower Threat Defense 545**
About the Firewall Mode 545
About Routed Firewall Mode 545
About Transparent Firewall Mode 546
Using the Transparent Firewall in Your Network 546
Interface 546
Passing Traffic For Routed-Mode Features 547
About Bridge Groups 547
Bridge Virtual Interface (BVI) 547

Bridge Groups in Transparent Firewall Mode 547
Bridge Groups in Routed Firewall Mode 548
Allowing Layer 3 Traffic 549
Allowed MAC Addresses 549
BPDU Handling 549
MAC Address vs. Route Lookups 550
Unsupported Features for Bridge Groups in Transparent Mode 551
Unsupported Features for Bridge Groups in Routed Mode 552
Default Settings 553
Guidelines for Firewall Mode 553
Set the Firewall Mode 554

CHAPTER 27 **Logical Devices for the Firepower Threat Defense on the Firepower 4100/9300** 557
About Interfaces 557
Chassis Management Interface 557
Interface Types 558
FXOS Interfaces vs. Application Interfaces 559
Inline Set Link State Propagation for the FTD 559
About Logical Devices 559
Standalone and Clustered Logical Devices 560
Requirements and Prerequisites for Logical Devices 560
Requirements and Prerequisites for Hardware and Software Combinations 560
Requirements and Prerequisites for High Availability 561
Guidelines and Limitations for Logical Devices 561
Guidelines and Limitations for Interfaces 561
General Guidelines and Limitations 562
Configure Interfaces 562
Enable or Disable an Interface 562
Configure a Physical Interface 563
Add an EtherChannel (Port Channel) 564
Configure Logical Devices 565
Add a Standalone FTD 565
Add a High Availability Pair 568
Change an Interface on a FTD Logical Device 569

Connect to the Console of the Application 571
History for Firepower Threat Defense Logical Devices 572

CHAPTER 28 **FlexConfig Policies for FTD 575**
FlexConfig Policy Overview 575
Recommended Usage for FlexConfig Policies 575
CLI Commands in FlexConfig Objects 576
Determine the ASA Software Version and Current CLI Configuration 576
Prohibited CLI Commands 577
Template Scripts 579
FlexConfig Variables 579
How to Process Variables 580
How to See What a Variable Will Return for a Device 583
FlexConfig Policy Object Variables 584
FlexConfig System Variables 585
Predefined FlexConfig Objects 586
Predefined Text Objects 589
Requirements and Prerequisites for FlexConfig Policies 591
Guidelines and Limitations for FlexConfig 592
Customizing Device Configuration with FlexConfig Policies 592
Configure FlexConfig Objects 594
Add a Policy Object Variable to a FlexConfig Object 596
Configure FlexConfig Text Objects 597
Configure the FlexConfig Policy 599
Set Target Devices for a FlexConfig Policy 600
Preview the FlexConfig Policy 600
Verify the Deployed Configuration 601
Remove Features Configured Using FlexConfig 603
Convert from FlexConfig to Managed Feature 604
History for FlexConfig 605

PART VIII **Firepower Threat Defense Interfaces and Device Settings 607**

CHAPTER 29 **Interface Overview for Firepower Threat Defense 609**

Management/Diagnostic Interface 609
Management Interface 609
Diagnostic Interface 609
Interface Mode and Types 610
Security Zones and Interface Groups 611
Auto-MDI/MDIX Feature 612
Default Settings for Interfaces 612
Enable the Physical Interface and Configure Ethernet Settings 613
Sync Interface Changes with the Firepower Management Center 614

CHAPTER 30 **Regular Firewall Interfaces for Firepower Threat Defense** 617
Requirements and Prerequisites for Regular Firewall Interfaces 617
Configure EtherChannel and Redundant Interfaces 617
About EtherChannels 618
About Redundant Interfaces (ASA Platform Only) 618
About EtherChannels 618
Guidelines for EtherChannels 621
Configure a Redundant Interface (ASA Platform Only) 622
Configure an EtherChannel 623
Configure VLAN Subinterfaces and 802.1Q Trunking 625
Guidelines and Limitations for VLAN Subinterfaces 625
Maximum Number of VLAN Subinterfaces by Device Model 626
Add a Subinterface 626
Configure Routed and Transparent Mode Interfaces 627
About Routed and Transparent Mode Interfaces 627
Dual IP Stack (IPv4 and IPv6) 628
Guidelines and Limitations for Routed and Transparent Mode Interfaces 628
Configure Routed Mode Interfaces 629
Configure Bridge Group Interfaces 632
Configure General Bridge Group Member Interface Parameters 632
Configure the Bridge Virtual Interface (BVI) 633
Configure a Diagnostic (Management) Interface for Transparent Mode 635
Configure IPv6 Addressing 636
About IPv6 636

Configure a Global IPv6 Address 637
Configure IPv6 Neighbor Discovery 639
Configure Advanced Interface Settings 641
About Advanced Interface Configuration 641
About MAC Addresses 642
About the MTU 642
About the TCP MSS 643
ARP Inspection for Bridge Group Traffic 644
MAC Address Table 645
Default Settings 645
Guidelines for ARP Inspection and the MAC Address Table 645
Configure the MTU 646
Configure the MAC Address 646
Add a Static ARP Entry 647
Add a Static MAC Address and Disable MAC Learning for a Bridge Group 648
Set Security Configuration Parameters 649
History for Regular Firewall Interfaces for Firepower Threat Defense 651

CHAPTER 31 **Inline Sets and Passive Interfaces for Firepower Threat Defense** 653
About IPS Interfaces 653
IPS Interface Types 653
About Hardware Bypass for Inline Sets 654
Hardware Bypass Triggers 654
Hardware Bypass Switchover 655
Snort Fail Open vs. Hardware Bypass 655
Hardware Bypass Status 655
Requirements and Prerequisites for Inline Sets 655
Guidelines for Inline Sets and Passive Interfaces 656
Configure a Passive Interface 657
Configure an Inline Set 659
History for Inline Sets and Passive Interfaces for Firepower Threat Defense 662

CHAPTER 32 **DHCP and DDNS Services for Threat Defense** 663
About DHCP and DDNS Services 663

About the DHCPv4 Server 663
DHCP Options 663
About the DHCP Relay Agent 664
Requirements and Prerequisites for DHCP and DDNS 664
Guidelines for DHCP and DDNS Services 664
Configure the DHCP Server 666
Configure the DHCP Relay Agent 667
Configure Dynamic DNS 668

CHAPTER 33 **Quality of Service (QoS) for Firepower Threat Defense 673**
Introduction to QoS 673
About QoS Policies 673
Requirements and Prerequisites for QoS 674
Rate Limiting with QoS Policies 675
Creating a QoS Policy 676
Setting Target Devices for a QoS Policy 676
Configuring QoS Rules 677
QoS Rule Components 678

PART IX **Firepower Threat Defense High Availability and Scalability 681**

CHAPTER 34 **High Availability for FTD 683**
About Firepower Threat Defense High Availability 683
High Availability System Requirements 683
Hardware Requirements 683
Software Requirements 684
License Requirements for FTD Devices in a High Availability Pair 684
Failover and Stateful Failover Links 684
Failover Link 685
Stateful Failover Link 686
Avoiding Interrupted Failover and Data Links 686
MAC Addresses and IP Addresses in High Availability 688
Stateful Failover 689
Supported Features 689

Unsupported Features 691
Bridge Group Requirements for High Availability 691
Failover Health Monitoring 692
Unit Health Monitoring 692
Interface Monitoring 692
Failover Triggers and Detection Timing 694
About Active/Standby Failover 695
Primary/Secondary Roles and Active/Standby Status 695
Active Unit Determination at Startup 695
Failover Events 695
Requirements and Prerequisites for High Availability 696
Guidelines for High Availability 697
Add a Firepower Threat Defense High Availability Pair 699
Configure Optional High Availability Parameters 700
Configure Standby IP Addresses and Interface Monitoring 700
Edit High Availability Failover Criteria 701
Configure Virtual MAC addresses 702
Manage High Availability 702
Switch the Active Peer in a Firepower Threat Defense High Availability Pair 702
Suspend and Resume High Availability 703
Replace a Unit in an FTD High Availability Pair 704
Replace a Primary FTD HA Unit 704
Replace a Secondary FTD HA Unit 704
Separate Units in a High Availability Pair 705
Unregister a High Availability Pair 706
Monitoring High Availability 706
View Failover History 707
View Stateful Failover Statistics 707

CHAPTER 35 **Clustering for the Firepower Threat Defense** 709
About Clustering on the Firepower 4100/9300 Chassis 709
Bootstrap Configuration 710
Cluster Members 710
Cluster Control Link 710

Size the Cluster Control Link 711
Cluster Control Link Redundancy 711
Cluster Control Link Reliability for Inter-Chassis Clustering 711
Cluster Control Link Network 712
Management Network 712
Management Interface 712
Cluster Interfaces 712
Spanned EtherChannels 712
Connecting to a Redundant Switch System 713
Configuration Replication 713
Licenses for Clustering 713
Requirements and Prerequisites for Clustering 714
Clustering Guidelines and Limitations 715
Configure Clustering 718
FXOS: Add a FTD Cluster 718
Create a FTD Cluster 718
Add More Cluster Nodes 724
FMC: Add a Cluster 727
FMC: Configure Cluster, Data, and Diagnostic Interfaces 728
FXOS: Remove a Cluster Node 730
FMC: Manage Cluster Members 731
Add or Replace a Cluster Member 731
Deactivate a Member 732
Rejoin the Cluster 733
Delete a Data Unit 733
FMC: Monitoring the Cluster 734
Examples for Clustering 734
Firewall on a Stick 735
Traffic Segregation 736
Spanned EtherChannel with Backup Links (Traditional 8 Active/8 Standby) 736
Reference for Clustering 739
Firepower Threat Defense Features and Clustering 740
Unsupported Features with Clustering 740
Centralized Features for Clustering 740

Dynamic Routing and Clustering 741
Connection Settings 741
FTP and Clustering 741
NAT and Clustering 742
SIP Inspection and Clustering 742
SNMP and Clustering 743
Syslog and Clustering 743
TLS/SSL Connections and Clustering 743
Cisco TrustSec and Clustering 743
Performance Scaling Factor 743
Control Unit Election 743
High Availability Within the Cluster 744
Chassis-Application Monitoring 744
Unit Health Monitoring 744
Interface Monitoring 745
Decorator Application Monitoring 745
Status After Failure 745
Rejoining the Cluster 745
Data Path Connection State Replication 746
How the Cluster Manages Connections 746
Connection Roles 747
New Connection Ownership 748
Sample Data Flow for TCP 748
History for Clustering 749

PART X **Firepower Threat Defense Routing** 751

CHAPTER 36 **Routing Overview for Firepower Threat Defense** 753
Path Determination 753
Supported Route Types 753
Static Versus Dynamic 754
Single-Path Versus Multipath 754
Flat Versus Hierarchical 754
Link-State Versus Distance Vector 755

Supported Internet Protocols for Routing 755
Routing Table 755
How the Routing Table Is Populated 756
Administrative Distances for Routes 756
Backup Dynamic and Floating Static Routes 757
How Forwarding Decisions Are Made 758
Dynamic Routing and High Availability 758
Dynamic Routing in Clustering 758
Routing Table for Management Traffic 759
Equal-Cost Multi-Path (ECMP) Routing 759
About Route Maps 760
Permit and Deny Clauses 760
Match and Set Clause Values 760

CHAPTER 37 **Static and Default Routes for Firepower Threat Defense 763**
About Static and Default Routes 763
Default Route 763
Static Routes 764
Route to null0 Interface to Drop Unwanted Traffic 764
Route Priorities 764
Transparent Firewall Mode and Bridge Group Routes 764
Static Route Tracking 765
Requirements and Prerequisites for Static Routes 765
Guidelines for Static and Default Routes 766
Add a Static Route 766

CHAPTER 38 **OSPF for Firepower Threat Defense 769**
OSPF for Firepower Threat Defense 769
About OSPF 769
OSPF Support for Fast Hello Packets 771
Prerequisites for OSPF Support for Fast Hello Packets 771
OSPF Hello Interval and Dead Interval 771
OSPF Fast Hello Packets 771
Benefits of OSPF Fast Hello Packets 771

Implementation Differences Between OSPFv2 and OSPFv3 772
Requirements and Prerequisites for OSPF 772
Guidelines for OSPF 772
Configure OSPFv2 774
Configure OSPF Areas, Ranges, and Virtual Links 774
Configure OSPF Redistribution 777
Configure OSPF Inter-Area Filtering 778
Configure OSPF Filter Rules 779
Configure OSPF Summary Addresses 780
Configure OSPF Interfaces and Neighbors 781
Configure OSPF Advanced Properties 783
Configure OSPFv3 786
Configure OSPFv3 Areas, Route Summaries, and Virtual Links 786
Configure OSPFv3 Redistribution 788
Configure OSPFv3 Summary Prefixes 789
Configure OSPFv3 Interfaces, Authentication, and Neighbors 790
Configure OSPFv3 Advanced Properties 792

CHAPTER 39 **BGP for Firepower Threat Defense** 797
About BGP 797
Routing Table Changes 797
When to Use BGP 798
BGP Path Selection 798
BGP Multipath 799
Requirements and Prerequisites for BGP 800
Guidelines for BGP 800
Configure BGP 801
Configure BGP Basic Settings 801
Configure BGP General Settings 803
Configure BGP Neighbor Settings 804
Configure BGP Aggregate Address Settings 808
Configure BGPv4 Filtering Settings 808
Configure BGP Network Settings 809
Configure BGP Redistribution Settings 809

Configure BGP Route Injection Settings 810

CHAPTER 40 **RIP for Firepower Threat Defense 813**

About RIP 813

Routing Update Process 813

RIP Routing Metric 814

RIP Stability Features 814

RIP Timers 814

Requirements and Prerequisites for RIP 815

Guidelines for RIP 815

Configure RIP 816

CHAPTER 41 **Multicast Routing for Firepower Threat Defense 819**

About Multicast Routing 819

IGMP Protocol 819

Stub Multicast Routing 820

PIM Multicast Routing 820

PIM Source Specific Multicast Support 821

Multicast Bidirectional PIM 821

PIM Bootstrap Router (BSR) 822

PIM Bootstrap Router (BSR) Terminology 822

Multicast Group Concept 823

Multicast Addresses 823

Clustering 823

Requirements and Prerequisites for Multicast Routing 823

Guidelines for Multicast Routing 823

Configure IGMP Features 824

Enable Multicast Routing 825

Configure IGMP Protocol 825

Configure IGMP Access Groups 827

Configure IGMP Static Groups 827

Configure IGMP Join Groups 828

Configure PIM Features 829

Configure PIM Protocol 829

Configure PIM Neighbor Filters 830
Configure PIM Bidirectional Neighbor Filters 831
Configure PIM Rendezvous Points 832
Configure PIM Route Trees 832
Configure PIM Request Filters 833
Configure the Firepower Threat Defense Device as a Candidate Bootstrap Router 834
Configure Multicast Routes 835
Configure Multicast Boundary Filters 835

PART XI **Firepower Threat Defense VPN** 837

CHAPTER 42 **VPN Overview for Firepower Threat Defense** 839
VPN Types 839
VPN Basics 840
Internet Key Exchange (IKE) 840
IPsec 841
VPN Packet Flow 842
VPN Licensing 842
How Secure Should a VPN Connection Be? 842
Deciding Which Encryption Algorithm to Use 843
Deciding Which Hash Algorithms to Use 843
Deciding Which Diffie-Hellman Modulus Group to Use 844
Deciding Which Authentication Method to Use 845
Pre-shared Keys 845
PKI Infrastructure and Digital Certificates 845
VPN Topology Options 847
Point-to-Point VPN Topology 847
Hub and Spoke VPN Topology 847
Full Mesh VPN Topology 848
Implicit Topologies 849

CHAPTER 43 **Site-to-Site VPNs for Firepower Threat Defense** 851
About Firepower Threat Defense Site-to-site VPNs 851
Firepower Threat Defense Site-to-site VPN Guidelines and Limitations 852

Requirements and Prerequisites for Site-to-Site VPN 853
Managing Firepower Threat Defense Site-to-site VPNs 853
Configuring Firepower Threat Defense Site-to-site VPNs 854
Firepower Threat Defense VPN Endpoint Options 855
Firepower Threat Defense VPN IKE Options 856
Firepower Threat Defense VPN IPsec Options 858
Firepower Threat Defense Advanced Site-to-site VPN Deployment Options 860
Firepower Threat Defense VPN Advanced IKE Options 861
Firepower Threat Defense VPN Advanced IPsec Options 862
Firepower Threat Defense Advanced Site-to-site VPN Tunnel Options 862
Enabling Sysopt Connection Permit-vpn Option 863
Monitoring Firepower Threat Defense VPNs 863

PART XII **Appliance Platform Settings** 865

CHAPTER 44 **System Configuration** 867
Requirements and Prerequisites for the System Configuration 868
About System Configuration 868
Navigating the Firepower Management Center System Configuration 868
System Configuration Settings 869
Appliance Information 870
HTTPS Certificates 871
Default HTTPS Server Certificates 871
Custom HTTPS Server Certificates 872
HTTPS Server Certificate Requirements 872
HTTPS Client Certificates 873
Viewing the Current HTTPS Server Certificate 874
Generating an HTTPS Server Certificate Signing Request 874
Importing HTTPS Server Certificates 875
Requiring Valid HTTPS Client Certificates 876
External Database Access Settings 877
Enabling External Access to the Database 878
Database Event Limits 878
Configuring Database Event Limits 879

Database Event Limits 879
Management Interfaces 881
About FMC Management Interfaces 881
Management Interfaces on the FMC 881
Management Interface Support Per FMC Model 882
Network Routes on FMC Management Interfaces 882
NAT Environments 883
Management and Event Traffic Channel Examples 884
Modify FMC Management Interfaces 886
Shut Down or Restart 889
Shut Down or Restart the FMC 890
Remote Storage Management 890
Configuring Local Storage 890
Configuring NFS for Remote Storage 891
Configuring SMB for Remote Storage 891
Configuring SSH for Remote Storage 892
Remote Storage Management Advanced Options 893
Change Reconciliation 893
Configuring Change Reconciliation 894
Change Reconciliation Options 894
Policy Change Comments 895
Configuring Comments to Track Policy Changes 895
Access List 895
Configure an Access List 896
Audit Logs 897
Stream Audit Logs to Syslog 897
Stream Audit Logs to an HTTP Server 898
Audit Log Certificate 900
Securely Stream Audit Logs 900
Obtain a Signed Audit Log Client Certificate for the FMC 901
Import an Audit Log Client Certificate into the FMC 902
Require Valid Audit Log Server Certificates 903
View the Audit Log Client Certificate on the FMC 904
Dashboard Settings 904

Enabling Custom Analysis Widgets for Dashboards 905
DNS Cache 905
Configuring DNS Cache Properties 905
Email Notifications 906
Configuring a Mail Relay Host and Notification Address 906
Language Selection 907
Set the Language for the Web Interface 907
Login Banners 907
Customize the Login Banner 908
SNMP Polling 908
Configure SNMP Polling 908
Security Certifications Compliance 909
Security Certifications Compliance Characteristics 910
Security Certifications Compliance Recommendations 910
Enable Security Certifications Compliance 911
Time and Time Synchronization 913
Synchronize Time on the FMC with an NTP Server 913
Synchronize Time Without Access to a Network NTP Server 914
About Changing Time Synchronization Settings 915
View Current System Time, Source, and NTP Server Connection Status 915
NTP Server Status 916
Session Timeouts 917
Configure Session Timeouts 917
Vulnerability Mapping 918
Mapping Vulnerabilities for Servers 918
Remote Console Access Management 918
Configuring Remote Console Settings on the System 919
Lights-Out Management User Access Configuration 919
Enabling Lights-Out Management User Access 920
Serial Over LAN Connection Configuration 920
Configuring Serial Over LAN with IPMItool 921
Configuring Serial Over LAN with IPMIutil 922
Lights-Out Management Overview 922
Configuring Lights-Out Management with IPMItool 923

Configuring Lights-Out Management with IPMIutil 924
REST API Preferences 924
Enabling REST API Access 924
VMware Tools and Virtual Systems 925
Enabling VMware Tools on the Firepower Management Center for VMware 925

CHAPTER 45 **Platform Settings Policies** 927
Introduction to Platform Settings 927
Requirements and Prerequisites for Platform Settings Policies 928
Managing Platform Settings Policies 928
Create a Platform Settings Policy 929
Setting Target Devices for a Platform Settings Policy 929

CHAPTER 46 **Platform Settings for Classic Devices** 931
About Platform Settings for Classic Devices 931
Requirements for Platform Settings for Classic Devices 932
Configure Platform Settings for Classic Devices 933
Configure Access Lists for Classic Devices 933
Stream Audit Logs from Classic Devices 934
Require Valid Audit Log Server Certificates for Classic Devices 935
Enable External Authentication to 7000/8000 Series Devices 936
About External Authentication for 7000/8000 Series Devices 937
Set the Language for the 7000/8000 Series Web Interface 938
Customize the Login Banner for Classic Devices 938
Synchronize Time on Classic Devices with an NTP Server 939
Configure Session Timeouts for Classic Devices 940
Configure SNMP Polling on Classic Devices 940
Local System Configuration for 7000/8000 Series Devices 941
Obtain a Signed Client Certificate for Secure Audit Log Streaming on a 7000/8000 Series Device 942
Prohibit Packet Transfer to FMC 944
Configure Management Interfaces on a 7000/8000 Series Device 944
Shut Down or Restart a 7000/8000 Series Device 947
View System Time for 7000/8000 Series Devices 947

CHAPTER 47 **Platform Settings for Firepower Threat Defense** 949

Configure ARP Inspection 949
Configure Banners 950
Configure Fragment Handling 951
Configure HTTP 952
Configure ICMP Access Rules 953
Configure Secure Shell 954
Configure SMTP 956
Configure SNMP for Threat Defense 956
Add SNMPv3 Users 957
Add SNMP Hosts 959
Configure SNMP Traps 960
About Configuring Syslog 962
About Syslog 962
Severity Levels 962
Syslog Message Filtering 963
Syslog Message Classes 963
Guidelines for Logging 966
Configure Syslog Logging for FTD Devices 967
Enable Logging and Configure Basic Settings 968
Enable Logging Destinations 969
Send Syslog Messages to an E-mail Address 970
Create a Custom Event List 970
Limit the Rate of Syslog Message Generation 971
Configure Syslog Settings 972
Configure a Syslog Server 973
Configure Global Timeouts 974
Configure NTP Time Synchronization for Threat Defense 976
History for Firepower Threat Defense Platform Settings 977

PART XIII **Network Address Translation (NAT)** 979

CHAPTER 48 **NAT Policy Management** 981

Requirements and Prerequisites for NAT Policies 981
Managing NAT Policies 981
Creating NAT Policies 982
Configuring NAT Policies 983
Configuring NAT Policy Targets 984
Copying NAT Policies 985

CHAPTER 49 **NAT for 7000 and 8000 Series Devices** 987
NAT Policy Configuration 987
NAT Policies Configuration Guidelines 988
Rule Organization in a NAT Policy 988
Organizing NAT Rules 989
NAT Rule Warnings and Errors 990
Showing and Hiding NAT Rule Warnings 990
NAT Policy Rules Options 990
Creating and Editing NAT Rules 991
NAT Rule Types 992
NAT Rule Condition Types 994
NAT Rule Conditions and Condition Mechanics 994
NAT Rule Conditions 995
Adding Conditions to NAT Rules 995
Literal Conditions in NAT Rules 997
Objects in NAT Rule Conditions 997
Zone Conditions in NAT Rules 997
Adding Zone Conditions to NAT Rules 998
Source Network Conditions in Dynamic NAT Rules 999
Adding Network Conditions to a Dynamic NAT Rule 1000
Destination Network Conditions in NAT Rules 1001
Adding Destination Network Conditions to NAT Rules 1002
Port Conditions in NAT Rules 1003
Adding Port Conditions to NAT Rules 1003

CHAPTER 50 **Network Address Translation (NAT) for Firepower Threat Defense** 1005
Why Use NAT? 1005

NAT Basics 1006
NAT Terminology 1006
NAT Types 1006
NAT in Routed and Transparent Mode 1007
NAT in Routed Mode 1007
NAT in Transparent Mode or Within a Bridge Group 1008
Auto NAT and Manual NAT 1009
Auto NAT 1009
Manual NAT 1009
Comparing Auto NAT and Manual NAT 1010
NAT Rule Order 1010
NAT Interfaces 1012
Configuring Routing for NAT 1013
Addresses on the Same Network as the Mapped Interface 1013
Addresses on a Unique Network 1013
The Same Address as the Real Address (Identity NAT) 1013
Guidelines for NAT 1014
Firewall Mode Guidelines for NAT 1014
IPv6 NAT Guidelines 1014
IPv6 NAT Best Practices 1015
NAT Support for Inspected Protocols 1015
Additional Guidelines for NAT 1017
Configure NAT for Threat Defense 1019
Customizing NAT Rules for Multiple Devices 1021
Dynamic NAT 1023
About Dynamic NAT 1023
Dynamic NAT Disadvantages and Advantages 1024
Configure Dynamic Auto NAT 1024
Configure Dynamic Manual NAT 1026
Dynamic PAT 1028
About Dynamic PAT 1028
Dynamic PAT Disadvantages and Advantages 1029
PAT Pool Object Guidelines 1029
Configure Dynamic Auto PAT 1030

Configure Dynamic Manual PAT 1032
Static NAT 1035
About Static NAT 1035
Configure Static Auto NAT 1039
Configure Static Manual NAT 1040
Identity NAT 1043
Configure Identity Auto NAT 1044
Configure Identity Manual NAT 1045
NAT Rule Properties for Firepower Threat Defense 1048
Interface Objects NAT Properties 1048
Translation Properties for Auto NAT 1049
Translation Properties for Manual NAT 1050
PAT Pool NAT Properties 1051
Advanced NAT Properties 1052
Translating IPv6 Networks 1053
NAT64/46: Translating IPv6 Addresses to IPv4 1053
NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet 1054
NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet and DNS Translation 1056
NAT66: Translating IPv6 Addresses to Different IPv6 Addresses 1059
NAT66 Example, Static Translation between Networks 1059
NAT66 Example, Simple IPv6 Interface PAT 1061
Monitoring NAT 1063
Examples for NAT 1063
Providing Access to an Inside Web Server (Static Auto NAT) 1063
Dynamic Auto NAT for Inside Hosts and Static NAT for an Outside Web Server 1066
Inside Load Balancer with Multiple Mapped Addresses (Static Auto NAT, One-to-Many) 1069
Single Address for FTP, HTTP, and SMTP (Static Auto NAT-with-Port-Translation) 1072
Different Translation Depending on the Destination (Dynamic Manual PAT) 1077
Different Translation Depending on the Destination Address and Port (Dynamic Manual PAT) 1081
NAT and Site-to-Site VPN 1085
Rewriting DNS Queries and Responses Using NAT 1089
DNS64 Reply Modification 1090
DNS Reply Modification, DNS Server on Outside 1096
DNS Reply Modification, DNS Server on Host Network 1099

History for FTD NAT 1101

PART XIV 7000 and 8000 Series Advanced Deployment Options 1103

CHAPTER 51 Setting Up Virtual Switches 1105

Virtual Switches 1105

Switched Interface Configuration 1105

Switched Interface Configuration Notes 1106

Configuring Physical Switched Interfaces 1107

Adding Logical Switched Interfaces 1108

Deleting Logical Switched Interfaces 1109

Virtual Switch Configuration 1110

Virtual Switch Configuration Notes 1110

Adding Virtual Switches 1111

Advanced Virtual Switch Settings 1111

Configuring Advanced Virtual Switch Settings 1112

Deleting Virtual Switches 1113

CHAPTER 52 Setting Up Virtual Routers 1115

Virtual Routers 1115

Routed Interfaces 1116

Configuring Physical Routed Interfaces 1117

Adding Logical Routed Interfaces 1119

Deleting Logical Routed Interfaces 1121

Configuring SFRP 1122

Virtual Router Configuration 1123

Adding Virtual Routers 1124

DHCP Relay 1125

Setting Up DHCPv4 Relay 1125

Setting Up DHCPv6 Relay 1126

Static Routes 1127

Viewing the Static Routes Table 1127

Adding Static Routes 1128

Dynamic Routing 1129

RIP Configuration 1129
Adding Interfaces for RIP Configuration 1129
Configuring Authentication Settings for RIP Configuration 1130
Configuring Advanced Settings for RIP Configuration 1131
Adding Import Filters for RIP Configuration 1132
Adding Export Filters for RIP Configuration 1133
OSPF Configuration 1134
OSPF Routing Areas 1134
OSPF Area Interfaces 1135
Adding OSPF Area Interfaces 1137
Adding OSPF Area Vlinks 1138
Adding Import Filters for OSPF Configuration 1139
Adding Export Filters for OSPF Configuration 1140
Virtual Router Filters 1141
Viewing Virtual Router Filters 1142
Setting Up Virtual Router Filters 1142
Adding Virtual Router Authentication Profiles 1143
Viewing Virtual Router Statistics 1144
Deleting Virtual Routers 1145

CHAPTER 53 **Aggregate Interfaces and LACP** 1147
About Aggregate Interfaces 1147
LAG Configuration 1148
Aggregate Switched Interfaces 1148
Aggregate Routed Interfaces 1149
Logical Aggregate Interfaces 1149
Load-Balancing Algorithms 1150
Link Selection Policies 1151
Link Aggregation Control Protocol (LACP) 1152
LACP 1152
Adding Aggregate Switched Interfaces 1153
Adding Aggregate Routed Interfaces 1155
Adding Logical Aggregate Interfaces 1158
Viewing Aggregate Interface Statistics 1159

Deleting Aggregate Interfaces 1159

CHAPTER 54 **Hybrid Interfaces** 1161

About Hybrid Interfaces 1161

Logical Hybrid Interfaces 1161

Adding Logical Hybrid Interfaces 1162

Deleting Logical Hybrid Interfaces 1164

CHAPTER 55 **Gateway VPNs** 1165

Gateway VPN Basics 1165

IPsec 1165

IKE 1166

VPN Deployments 1166

Point-to-Point VPN Deployments 1166

Star VPN Deployments 1167

Mesh VPN Deployments 1168

VPN Deployment Management 1168

VPN Deployment Options 1169

Point-to-Point VPN Deployment Options 1169

Star VPN Deployment Options 1170

Mesh VPN Deployment Options 1172

Advanced VPN Deployment Options 1173

Managing VPN Deployments 1174

Configuring Point-to-Point VPN Deployments 1175

Configuring Star VPN Deployments 1175

Configuring Mesh VPN Deployments 1176

Configuring Advanced VPN Deployment Settings 1177

Editing VPN Deployments 1178

VPN Deployment Status 1179

Viewing VPN Status 1179

VPN Statistics and Logs 1179

Viewing VPN Statistics and Logs 1181

PART XV **Access Control** 1183

CHAPTER 56 **Understanding Access Control** 1185

Introduction to Access Control 1185

Access Control Policy Default Action 1185

Deep Inspection Using File and Intrusion Policies 1187

Access Control Traffic Handling with Intrusion and File Policies 1188

File and Intrusion Inspection Order 1189

Access Control Policy Inheritance 1191

CHAPTER 57 **Best Practices for Access Control** 1193

General Best Practices for Access Control 1193

Best Practices for Access Control Rules 1194

Best Practices for Ordering Rules 1194

Rule Preemption 1195

Rule Actions and Rule Order 1195

Content Restriction Rule Order 1196

Application Rule Order 1197

SSL Rule Order 1197

URL Rule Order 1198

Best Practices for Simplifying and Focusing Rules 1198

Maximum Number of Access Control Rules and Intrusion Policies 1199

CHAPTER 58 **Access Control Policies** 1201

Access Control Policy Components 1201

Requirements and Prerequisites for Access Control Policies 1203

Managing Access Control Policies 1203

System-Created Access Control Policies 1204

Creating a Basic Access Control Policy 1204

Editing an Access Control Policy 1205

Managing Access Control Policy Inheritance 1207

Choosing a Base Access Control Policy 1208

Inheriting Access Control Policy Settings from the Base Policy 1208

Locking Settings in Descendant Access Control Policies 1209

Requiring an Access Control Policy in a Domain 1209

Setting Target Devices for an Access Control Policy 1210
Access Control Policy Advanced Settings 1211
Associating Other Policies with Access Control 1213
History for Access Control Policies 1214

CHAPTER 59 **Access Control Rules** 1215
Introduction to Access Control Rules 1215
Access Control Rule Management 1217
Access Control Rule Components 1218
Access Control Rule Order 1219
Requirements and Prerequisites for Access Control Rules 1219
Adding an Access Control Rule Category 1220
Create and Edit Access Control Rules 1220
Enabling and Disabling Access Control Rules 1222
Positioning an Access Control Rule 1223
Access Control Rule Actions 1223
Access Control Rule Monitor Action 1223
Access Control Rule Trust Action 1224
Access Control Rule Blocking Actions 1224
Access Control Rule Interactive Blocking Actions 1225
Access Control Rule Allow Action 1226
Access Control Rule Comments 1226
Adding Comments to an Access Control Rule 1227

CHAPTER 60 **URL Filtering** 1229
URL Filtering Overview 1229
About URL Filtering with Category and Reputation 1229
Best Practices for URL Filtering 1230
Filtering HTTPS Traffic 1232
License Requirements for URL Filtering 1233
Requirements and Prerequisites for URL Filtering 1233
How to Configure URL Filtering with Category and Reputation 1234
Enable URL Filtering Using Category and Reputation 1235
URL Filtering Options 1235

Configuring URL Conditions 1236
Rules with URL Conditions 1238
URL Rule Order 1238
Manual URL Filtering 1238
Manual URL Filtering Options 1239
Supplement or Selectively Override Category and Reputation-Based URL Filtering 1240
Troubleshoot URL Filtering 1241

CHAPTER 61 **HTTP Response Pages and Interactive Blocking** 1243
About HTTP Response Pages 1243
Limitations to HTTP Response Pages 1243
Requirements and Prerequisites for HTTP Response Pages 1244
Choosing HTTP Response Pages 1245
Interactive Blocking with HTTP Response Pages 1245
Configuring Interactive Blocking 1246
Setting the User Bypass Timeout for a Blocked Website 1246

CHAPTER 62 **Blocking Traffic with Security Intelligence** 1249
About Security Intelligence 1249
Best Practices for Security Intelligence 1250
License Requirements for Security Intelligence 1250
Requirements and Prerequisites for Security Intelligence 1251
Security Intelligence Sources 1251
Configure Security Intelligence 1252
Security Intelligence Options 1254
Security Intelligence Categories 1255
Block List Icons 1256
Configuration Example: Security Intelligence Blocking 1257
Security Intelligence Monitoring 1258
Override Security Intelligence Blocking 1259
Troubleshooting Security Intelligence 1259
Security Intelligence Categories Are Missing from the Available Options List 1259
Troubleshooting Memory Use 1260
History for Security Intelligence Block Listing 1260

CHAPTER 63 **DNS Policies** 1261

DNS Policy Overview 1261
DNS Policy Components 1262
License Requirements for DNS Policies 1263
Requirements and Prerequisites for DNS Policies 1263
Managing DNS Policies 1263
Creating Basic DNS Policies 1264
Editing DNS Policies 1264
DNS Rules 1265
Creating and Editing DNS Rules 1266
DNS Rule Management 1266
Enabling and Disabling DNS Rules 1266
DNS Rule Order Evaluation 1267
DNS Rule Actions 1268
DNS Rule Conditions 1269
Controlling Traffic Based on DNS and Security Zone 1269
Controlling Traffic Based on DNS and Network 1270
Controlling Traffic Based on DNS and VLAN 1270
Controlling Traffic Based on DNS List, Feed, or Category 1271
DNS Policy Deploy 1272

CHAPTER 64 **Prefiltering and Prefilter Policies** 1273

About Prefiltering 1273
Prefiltering vs Access Control 1273
Passthrough Tunnels and Access Control 1275
Best Practices for Prefiltering 1276
Encapsulated Traffic Handling Best Practices 1277
Requirements and Prerequisites for Prefilter Policies 1277
Configure Prefiltering 1278
About Prefilter Policies 1279
Tunnel vs Prefilter Rules 1280
Tunnel and Prefilter Rule Components 1281
Tunnel Zones and Prefiltering 1282

Using Tunnel Zones 1283
Creating Tunnel Zones 1285
Large Flow Offloads 1286
Flow Offload Limitations 1286

CHAPTER 65 **Intelligent Application Bypass** 1289
Introduction to IAB 1289
IAB Options 1290
Requirements and Prerequisites for Intelligent Application Bypass 1292
Configuring Intelligent Application Bypass 1292
IAB Logging and Analysis 1293

CHAPTER 66 **Access Control Using Content Restriction** 1297
About Content Restriction 1297
Requirements and Prerequisites for Content Restriction 1298
Using Access Control Rules to Enforce Content Restriction 1299
Safe Search Options for Access Control Rules 1300
YouTube EDU Options for Access Control Rules 1300
Using a DNS Sinkhole to Enforce Content Restriction 1301

PART XVI **Encrypted Traffic Handling** 1303

CHAPTER 67 **Understanding Traffic Decryption** 1305
Traffic Decryption Explained 1305
TLS/SSL Handshake Processing 1307
ClientHello Message Handling 1307
ServerHello and Server Certificate Message Handling 1309
TLS/SSL Best Practices 1310
The Case for Decryption 1311
When to Decrypt Traffic, When Not to Decrypt 1312
Decrypt and Resign (Outgoing Traffic) 1313
Known Key Decryption (Incoming Traffic) 1313
Other TLS/SSL Rule Actions 1314
TLS/SSL Rule Examples 1314

Block Nonsecure Protocols 1314
TLS/SSL Rule Components 1316
TLS/SSL Rule Order Evaluation 1317
Multi-Rule Example 1317
How to Configure TLS/SSL Policies and Rules 1319
TLS/SSL Inspection Appliance Deployment Scenarios 1321
Traffic Decryption in a Passive Deployment 1321
Encrypted Traffic Monitoring in a Passive Deployment 1322
Undecrypted Encrypted Traffic in a Passive Deployment 1323
Encrypted Traffic Inspection with a Private Key in a Passive Deployment 1324
Traffic Decryption in an Inline Deployment 1326
Encrypted Traffic Monitoring in an Inline Deployment 1328
Undecrypted Encrypted Traffic in an Inline Deployment 1328
Encrypted Traffic Blocking in an Inline Deployment 1329
Encrypted Traffic Inspection with a Private Key in an Inline Deployment 1330
Encrypted Traffic Inspection with a Re-signed Certificate in an Inline Deployment 1332
History for TLS/SSL 1334

CHAPTER 68 **Start Creating SSL Policies** 1337
SSL Policies Overview 1337
SSL Policy Default Actions 1338
Default Handling Options for Undecryptable Traffic 1339
Requirements and Prerequisites for SSL Policies 1340
Manage SSL Policies 1340
Create Basic SSL Policies 1341
Set Default Handling for Undecryptable Traffic 1342
Editing an SSL Policy 1343

CHAPTER 69 **Get Started with TLS/SSL Rules** 1345
TLS/SSL Rules Overview 1345
TLS/SSL Rule Guidelines and Limitations 1345
Guideline for Using TLS/SSL Decryption 1346
TLS/SSL Rule Unsupported Features 1346
TLS/SSL Do Not Decrypt Guidelines 1347

TLS/SSL Decrypt - Resign Guidelines 1347
TLS/SSL Decrypt - Known Key Guidelines 1349
TLS/SSL Block Guidelines 1350
TLS/SSL Certificate Pinning Guidelines 1351
TLS/SSL Heartbeat Guidelines 1351
TLS/SSL Anonymous Cipher Suite Limitation 1351
TLS/SSL Normalizer Guidelines 1351
Other TLS/SSL Rule Guidelines 1352
Requirements and Prerequisites for TLS/SSL Rules 1352
Encrypted Traffic Inspection Configuration 1353
Creating and Modifying TLS/SSL Rules 1354
Adding a TLS/SSL Rule to a Rule Category 1354
Positioning a TLS/SSL Rule by Number 1355
TLS/SSL Rule Search 1355
Searching TLS/SSL Rules 1356
Enabling and Disabling TLS/SSL Rules 1356
Moving a TLS/SSL Rule 1356
Adding a New TLS/SSL Rule Category 1357
TLS/SSL Rule Conditions 1357
TLS/SSL Rule Condition Types 1358
TLS/SSL Rule Actions 1359
TLS/SSL Rule Monitor Action 1359
TLS/SSL Rule Do Not Decrypt Action 1360
TLS/SSL Rule Blocking Actions 1360
TLS/SSL Rule Decrypt Actions 1361
Configuring TLS/SSL Rule Actions 1361
Configuring a Decrypt - Resign Action 1362
Configuring a Decrypt - Known Key Action 1362
TLS/SSL Rules Management 1363

CHAPTER 70 **Decryption Tuning Using TLS/SSL Rules** 1365
TLS/SSL Rule Conditions Overview 1365
Requirements and Prerequisites for Decryption Tuning 1366
Network-Based TLS/SSL Rule Conditions 1366

Network Zone TLS/SSL Rule Conditions 1367
Controlling Encrypted Traffic by Network Zone 1367
Network or Geolocation TLS/SSL Rule Conditions 1368
Controlling Encrypted Traffic by Network or Geolocation 1369
VLAN TLS/SSL Rule Conditions 1370
Controlling Encrypted VLAN Traffic 1370
Port TLS/SSL Rule Conditions 1371
Controlling Encrypted Traffic by Port 1372
User-Based TLS/SSL Rule Conditions 1372
Controlling Encrypted Traffic Based on User 1373
Reputation-Based TLS/SSL Rule Conditions 1373
Selected Applications and Filters in TLS/SSL Rules 1373
Application Filters in TLS/SSL Rules 1374
Available Applications in TLS/SSL Rules 1375
Application-Based TLS/SSL Rule Condition Requirements 1376
Adding an Application Condition to a TLS/SSL Rule 1377
Limitations to Encrypted Application Control 1377
Reputation-Based URL Blocking in Encrypted Traffic 1378
Block Encrypted Traffic Based on URL Reputation 1378
Server Certificate-Based TLS/SSL Rule Conditions 1379
Certificate Distinguished Name TLS/SSL Rule Conditions 1380
Controlling Encrypted Traffic by Certificate Distinguished Name 1381
Certificate TLS/SSL Rule Conditions 1382
Controlling Encrypted Traffic by Certificate 1383
Certificate Status TLS/SSL Rule Conditions 1383
Trusting External Certificate Authorities 1386
Matching Traffic on Certificate Status 1387
Cipher Suite TLS/SSL Rule Conditions 1389
Controlling Encrypted Traffic by Cipher Suite 1391
Encryption Protocol Version TLS/SSL Rule Conditions 1392
Controlling Traffic by Encryption Protocol Version 1392

CHAPTER 71 **Troubleshoot TLS/SSL Rules** 1393
About TLS/SSL Pinning 1393

Troubleshoot TLS/SSL Pinning 1393
Troubleshoot Unknown or Bad Certificates or Certificate Authorities 1395
Verify TLS/SSL Cipher Suites 1397

PART XVII **Advanced Malware Protection (AMP) and File Control** 1399

CHAPTER 72 **File Policies and Malware Protection** 1401
About File Policies and Advanced Malware Protection 1401
File Policies 1401
Requirements and Prerequisites for File Policies 1402
License Requirements for File and Malware Policies 1403
Best Practices for File Policies and Malware Detection 1403
File Rule Best Practices 1403
File Detection Best Practices 1404
File Blocking Best Practices 1404
File Policy Best Practices 1405
How to Configure Malware Protection 1406
Plan and Prepare for Malware Protection 1406
Configure File Policies 1407
Add File Policies to Your Access Control Configuration 1408
Configuring an Access Control Rule to Perform Malware Protection 1409
Set Up Maintenance and Monitoring of Malware Protection 1410
Cloud Connections for Malware Protection 1410
AMP Cloud Connection Configurations 1411
Requirements and Best Practices for AMP Cloud Connections 1412
Choose an AMP Cloud 1412
Cisco AMP Private Cloud 1413
Managing Connections to the AMP Cloud (Public or Private) 1414
Change AMP Options 1415
Dynamic Analysis Connections 1416
Requirements for Dynamic Analysis 1416
Viewing the Default Dynamic Analysis Connection 1416
Dynamic Analysis On-Premises Appliance (Cisco Threat Grid) 1417
Enabling Access to Dynamic Analysis Results in the Public Cloud 1418

Maintain Your System: Update File Types Eligible for Dynamic Analysis 1419
File Policies and File Rules 1419
Create or Edit a File Policy 1419
Advanced and Archive File Inspection Options 1420
Managing File Policies 1423
File Rules 1424
File Rule Components 1425
File Rule Actions 1426
Creating File Rules 1433
Access Control Rule Logging for Malware Protection 1434
Retrospective Disposition Changes 1434
(Optional) Malware Protection with AMP for Endpoints 1434
Comparison of Malware Protection: Firepower vs. AMP for Endpoints 1435
About Integrating Firepower with AMP for Endpoints 1436
Benefits of Integrating Firepower and AMP for Endpoints 1436
AMP for Endpoints and AMP Private Cloud 1436
Integrate Firepower and AMP for Endpoints 1437
History for File Policies and Malware Protection 1439

CHAPTER 73 **File and Malware Inspection Performance and Storage Tuning** 1441
File and Malware Inspection Performance and Storage Options 1441
Tuning File and Malware Inspection Performance and Storage 1443

PART XVIII **Intrusion Detection and Prevention** 1445

CHAPTER 74 **An Overview of Intrusion Detection and Prevention** 1447
Network Analysis and Intrusion Policy Basics 1447
How Policies Examine Traffic For Intrusions 1448
Decoding, Normalizing, and Preprocessing: Network Analysis Policies 1449
Access Control Rules: Intrusion Policy Selection 1450
Intrusion Inspection: Intrusion Policies, Rules, and Variable Sets 1451
Intrusion Event Generation 1452
System-Provided and Custom Network Analysis and Intrusion Policies 1453
System-Provided Network Analysis and Intrusion Policies 1453

Benefits of Custom Network Analysis and Intrusion Policies 1455
Benefits of Custom Network Analysis Policies 1455
Benefits of Custom Intrusion Policies 1456
Limitations of Custom Policies 1457
License Requirements for Network Analysis and Intrusion Policies 1459
Requirements and Prerequisites for Network Analysis and Intrusion Policies 1459
The Navigation Panel: Network Analysis and Intrusion Policies 1459
Conflicts and Changes: Network Analysis and Intrusion Policies 1461
Exiting a Network Analysis or Intrusion Policy 1462

CHAPTER 75 **Layers in Intrusion and Network Analysis Policies** 1463
Layer Basics 1463
License Requirements for Network Analysis and Intrusion Policy Layers 1463
Requirements and Prerequisites for Network Analysis and Intrusion Policy Layers 1464
The Layer Stack 1464
The Base Layer 1465
System-Provided Base Policies 1465
Custom Base Policies 1465
The Effect of Rule Updates on Base Policies 1466
Changing the Base Policy 1467
The Firepower Recommendations Layer 1467
Layer Management 1468
Shared Layers 1469
Managing Layers 1470
Navigating Layers 1471
Intrusion Rules in Layers 1472
Configuring Intrusion Rules in Layers 1473
Removing Rule Settings from Multiple Layers 1473
Accepting Rule Changes from a Custom Base Policy 1475
Preprocessors and Advanced Settings in Layers 1475
Configuring Preprocessors and Advanced Settings in Layers 1476

CHAPTER 76 **Getting Started with Intrusion Policies** 1479
Intrusion Policy Basics 1479

License Requirements for Intrusion Policies 1480
Requirements and Prerequisites for Intrusion Policies 1481
Managing Intrusion Policies 1481
Custom Intrusion Policy Creation 1482
Creating a Custom Intrusion Policy 1482
Editing Snort 2 Intrusion Policies 1483
Intrusion Policy Changes 1484
Access Control Rule Configuration to Perform Intrusion Prevention 1484
Access Control Rule Configuration and Intrusion Policies 1485
Configuring an Access Control Rule to Perform Intrusion Prevention 1485
Drop Behavior in an Inline Deployment 1486
Setting Drop Behavior in an Inline Deployment 1486
Drop Behavior in a Dual System Deployment 1487
Intrusion Policy Advanced Settings 1487
Optimizing Performance for Intrusion Detection and Prevention 1488

CHAPTER 77 **Tuning Intrusion Policies Using Rules** 1489
Intrusion Rule Tuning Basics 1489
Intrusion Rule Types 1489
License Requirements for Intrusion Rules 1490
Requirements and Prerequisites for Intrusion Rules 1491
Viewing Intrusion Rules in an Intrusion Policy 1491
Intrusion Rules Page Columns 1491
Intrusion Rule Details 1492
Viewing Intrusion Rule Details 1493
Setting a Threshold for an Intrusion Rule 1494
Setting Suppression for an Intrusion Rule 1494
Setting a Dynamic Rule State from the Rule Details Page 1495
Setting an SNMP Alert for an Intrusion Rule 1496
Adding a Comment to an Intrusion Rule 1496
Intrusion Rule Filters in an Intrusion Policy 1497
Intrusion Rule Filters Notes 1497
Intrusion Policy Rule Filters Construction Guidelines 1497
Intrusion Rule Configuration Filters 1500

Intrusion Rule Content Filters 1500
Intrusion Rule Categories 1501
Intrusion Rule Filter Components 1501
Intrusion Rule Filter Usage 1502
Setting a Rule Filter in an Intrusion Policy 1502
Intrusion Rule States 1503
Intrusion Rule State Options 1503
Setting Intrusion Rule States 1504
Intrusion Event Notification Filters in an Intrusion Policy 1505
Intrusion Event Thresholds 1505
Intrusion Event Thresholds Configuration 1505
Adding and Modifying Intrusion Event Thresholds 1507
Viewing and Deleting Intrusion Event Thresholds 1508
Intrusion Policy Suppression Configuration 1509
Intrusion Policy Suppression Types 1509
Suppressing Intrusion Events for a Specific Rule 1509
Viewing and Deleting Suppression Conditions 1510
Dynamic Intrusion Rule States 1511
Dynamic Intrusion Rule State Configuration 1512
Setting a Dynamic Rule State from the Rules Page 1512
Adding Intrusion Rule Comments 1514

CHAPTER 78 **Tailoring Intrusion Protection to Your Network Assets** 1517
About Firepower Recommended Rules 1517
Default Settings for Firepower Recommendations 1518
Advanced Settings for Firepower Recommendations 1519
Generating and Applying Firepower Recommendations 1520

CHAPTER 79 **Sensitive Data Detection** 1523
Sensitive Data Detection Basics 1523
Global Sensitive Data Detection Options 1524
Individual Sensitive Data Type Options 1525
System-Provided Sensitive Data Types 1526
License Requirements for Sensitive Data Detection 1527

Requirements and Prerequisites for Sensitive Data Detection 1527
Configuring Sensitive Data Detection 1527
Monitored Application Protocols and Sensitive Data 1529
Selecting Application Protocols to Monitor 1529
Special Case: Sensitive Data Detection in FTP Traffic 1530
Custom Sensitive Data Types 1531
Data Patterns in Custom Sensitive Data Types 1531
Configuring Custom Sensitive Data Types 1533
Editing Custom Sensitive Data Types 1534

CHAPTER 80 **Globally Limiting Intrusion Event Logging** 1537
Global Rule Thresholding Basics 1537
Global Rule Thresholding Options 1538
License Requirements for Global Thresholds 1539
Requirements and Prerequisites for Global Thresholds 1540
Configuring Global Thresholds 1540
Disabling the Global Threshold 1541

CHAPTER 81 **The Intrusion Rules Editor** 1543
An Introduction to Intrusion Rule Editing 1543
License Requirements for the Intrusion Rule Editor 1544
Requirements and Prerequisites for the Intrusion Rule Editor 1544
Rule Anatomy 1544
The Intrusion Rule Header 1545
Intrusion Rule Header Action 1546
Intrusion Rule Header Protocol 1546
Intrusion Rule Header Direction 1547
Intrusion Rule Header Source and Destination IP Addresses 1547
Intrusion Rule Header Source and Destination Ports 1550
Intrusion Event Details 1551
Adding a Custom Classification 1554
Defining an Event Priority 1555
Defining an Event Reference 1555
Custom Rule Creation 1556

Writing New Rules 1557
Modifying Existing Rules 1558
Viewing Rule Documentation 1559
Adding Comments to Intrusion Rules 1559
Deleting Custom Rules 1560
Searching for Rules 1561
Search Criteria for Intrusion Rules 1561
Rule Filtering on the Intrusion Rules Editor Page 1562
Filtering Guidelines 1562
Keyword Filtering 1563
Character String Filtering 1564
Combination Keyword and Character String Filtering 1564
Filtering Rules 1565
Keywords and Arguments in Intrusion Rules 1565
The content and protected_content Keywords 1566
Basic content and protected_content Keyword Arguments 1567
content and protected_content Keyword Search Locations 1568
Overview: HTTP content and protected_content Keyword Arguments 1571
Overview: content Keyword Fast Pattern Matcher 1574
The replace Keyword 1577
The byte_jump Keyword 1578
The byte_test Keyword 1580
The byte_extract Keyword 1582
The byte_math Keyword 1585
Overview: The pcre Keyword 1587
pcre Syntax 1588
pcre Modifier Options 1590
pcre Example Keyword Values 1593
The metadata Keyword 1595
Service Metadata 1596
Metadata Search Guidelines 1601
IP Header Values 1602
ICMP Header Values 1604
TCP Header Values and Stream Size 1605

The stream_reassembly Keyword 1609
SSL Keywords 1609
The appid Keyword 1611
Application Layer Protocol Values 1612
The RPC Keyword 1612
The ASN.1 Keyword 1612
The urilen Keyword 1613
DCE/RPC Keywords 1614
SIP Keywords 1617
GTP Keywords 1619
SCADA Keywords 1631
Modbus Keywords 1631
DNP3 Keywords 1632
Packet Characteristics 1635
Active Response Keywords 1636
The resp Keyword 1637
The react Keyword 1638
The detection_filter Keyword 1638
The tag Keyword 1640
The flowbits Keyword 1641
flowbits Keyword Options 1641
Guidelines for Using the flowbits Keyword 1642
flowbits Keyword Examples 1643
The http_encode Keyword 1648
http_encode Keyword Syntax 1649
http_encode Keyword example: Using Two http_endcode Keywords to Search for Two Encodings 1649
Overview: The file_type and file_group Keywords 1649
The file_type and file_group Keywords 1650
The file_data Keyword 1651
The pkt_data Keyword 1652
The base64_decode and base64_data Keywords 1652

CHAPTER 82 **Intrusion Prevention Performance Tuning** 1655

About Intrusion Prevention Performance Tuning 1655
License Requirements for Intrusion Prevention Performance Tuning 1656
Requirements and Prerequisites for Intrusion Prevention Performance Tuning 1656
Limiting Pattern Matching for Intrusions 1656
Regular Expression Limits Overrides for Intrusion Rules 1657
Overriding Regular Expression Limits for Intrusion Rules 1658
Per Packet Intrusion Event Generation Limits 1658
Limiting Intrusion Events Generated Per Packet 1659
Packet and Intrusion Rule Latency Threshold Configuration 1659
Latency-Based Performance Settings 1660
Packet Latency Thresholding 1660
Packet Latency Thresholding Notes 1661
Configuring Packet Latency Thresholding 1661
Rule Latency Thresholding 1662
Rule Latency Thresholding Notes 1664
Configuring Rule Latency Thresholding 1664
Intrusion Performance Statistic Logging Configuration 1665
Configuring Intrusion Performance Statistic Logging 1666

PART XIX **Advanced Network Analysis and Preprocessing** 1667

CHAPTER 83 **Advanced Access Control Settings for Network Analysis and Intrusion Policies** 1669
About Advanced Access Control Settings for Network Analysis and Intrusion Policies 1669
Requirements and Prerequisites for Advanced Access Control Settings for Network Analysis and Intrusion Policies 1669
Inspection of Packets That Pass Before Traffic Is Identified 1670
Best Practices for Handling Packets That Pass Before Traffic Identification 1670
Specify a Policy to Handle Packets That Pass Before Traffic Identification 1670
Advanced Settings for Network Analysis Policies 1672
Setting the Default Network Analysis Policy 1672
Network Analysis Rules 1673
Configuring Network Analysis Rules 1674
Managing Network Analysis Rules 1674

CHAPTER 84 **Getting Started with Network Analysis Policies** 1677

Network Analysis Policy Basics 1677

License Requirements for Network Analysis Policies 1678

Requirements and Prerequisites for Network Analysis Policies 1678

Managing Network Analysis Policies 1678

Custom Network Analysis Policy Creation for Snort 2 1679

Creating a Custom Network Analysis Policy 1679

Network Analysis Policy Management for Snort 2 1680

Network Analysis Policy Settings and Cached Changes 1681

Editing Network Analysis Policies 1681

Preprocessor Configuration in a Network Analysis Policy for Snort 2 1682

Preprocessor Traffic Modification in Inline Deployments 1683

Preprocessor Configuration in a Network Analysis Policy Notes 1683

CHAPTER 85 **Application Layer Preprocessors** 1685

Introduction to Application Layer Preprocessors 1685

License Requirements for Application Layer Preprocessors 1686

Requirements and Prerequisites for Application Layer Preprocessors 1686

The DCE/RPC Preprocessor 1686

Connectionless and Connection-Oriented DCE/RPC Traffic 1687

DCE/RPC Target-Based Policies 1688

RPC over HTTP Transport 1688

DCE/RPC Global Options 1689

DCE/RPC Target-Based Policy Options 1691

Traffic-Associated DCE/RPC Rules 1695

Configuring the DCE/RPC Preprocessor 1695

The DNS Preprocessor 1697

DNS Preprocessor Options 1698

Configuring the DNS Preprocessor 1699

The FTP/Telnet Decoder 1700

Global FTP and Telnet Options 1700

Telnet Options 1701

Server-Level FTP Options 1702

FTP Command Validation Statements 1704
Client-Level FTP Options 1705
Configuring the FTP/Telnet Decoder 1706
The HTTP Inspect Preprocessor 1708
Global HTTP Normalization Options 1708
Server-Level HTTP Normalization Options 1709
Server-Level HTTP Normalization Encoding Options 1717
Configuring The HTTP Inspect Preprocessor 1720
Additional HTTP Inspect Preprocessor Rules 1721
The Sun RPC Preprocessor 1722
Sun RPC Preprocessor Options 1722
Configuring the Sun RPC Preprocessor 1723
The SIP Preprocessor 1724
SIP Preprocessor Options 1724
Configuring the SIP Preprocessor 1726
Additional SIP Preprocessor Rules 1727
The GTP Preprocessor 1728
GTP Preprocessor Rules 1729
Configuring the GTP Preprocessor 1729
The IMAP Preprocessor 1730
IMAP Preprocessor Options 1730
Configuring the IMAP Preprocessor 1731
Additional IMAP Preprocessor Rules 1732
The POP Preprocessor 1733
POP Preprocessor Options 1733
Configuring the POP Preprocessor 1735
Additional POP Preprocessor Rules 1736
The SMTP Preprocessor 1736
SMTP Preprocessor Options 1736
Configuring SMTP Decoding 1741
The SSH Preprocessor 1742
SSH Preprocessor Options 1742
Configuring the SSH Preprocessor 1745
The SSL Preprocessor 1746

How SSL Preprocessing Works 1746
SSL Preprocessor Options 1747
Configuring the SSL Preprocessor 1748
SSL Preprocessor Rules 1749

CHAPTER 86 SCADA Preprocessors 1751
Introduction to SCADA Preprocessors 1751
License Requirements for SCADA Preprocessors 1751
Requirements and Prerequisites for SCADA Preprocessors 1752
The Modbus Preprocessor 1752
Modbus Preprocessor Ports Option 1752
Configuring the Modbus Preprocessor 1752
Modbus Preprocessor Rules 1753
The DNP3 Preprocessor 1754
DNP3 Preprocessor Options 1754
Configuring the DNP3 Preprocessor 1754
DNP3 Preprocessor Rules 1755

CHAPTER 87 Transport & Network Layer Preprocessors 1757
Introduction to Transport and Network Layer Preprocessors 1757
License Requirements for Transport and Network Layer Preprocessors 1757
Requirements and Prerequisites for Transport and Network Layer Preprocessors 1758
Advanced Transport/Network Preprocessor Settings 1758
Ignored VLAN Headers 1758
Active Responses in Intrusion Drop Rules 1759
Advanced Transport/Network Preprocessor Options 1759
Configuring Advanced Transport/Network Preprocessor Settings 1760
Checksum Verification 1761
Checksum Verification Options 1761
Verifying Checksums 1762
The Inline Normalization Preprocessor 1762
Inline Normalization Options 1763
Configuring Inline Normalization 1768
The IP Defragmentation Preprocessor 1769

IP Fragmentation Exploits 1769
Target-Based Defragmentation Policies 1770
IP Defragmentation Options 1770
Configuring IP Defragmentation 1772
The Packet Decoder 1774
Packet Decoder Options 1774
Configuring Packet Decoding 1777
TCP Stream Preprocessing 1778
State-Related TCP Exploits 1778
Target-Based TCP Policies 1778
TCP Stream Reassembly 1779
TCP Stream Preprocessing Options 1780
Configuring TCP Stream Preprocessing 1787
UDP Stream Preprocessing 1788
UDP Stream Preprocessing Options 1789
Configuring UDP Stream Preprocessing 1789

CHAPTER 88 **Detecting Specific Threats** 1791
Introduction to Specific Threat Detection 1791
License Requirements for Specific Threat Detection 1791
Requirements and Prerequisites for Specific Threat Detection 1792
Back Orifice Detection 1792
Back Orifice Detection Preprocessor 1792
Detecting Back Orifice 1793
Portscan Detection 1793
Portscan Types, Protocols, and Filtered Sensitivity Levels 1794
Portscan Event Generation 1796
Portscan Event Packet View 1797
Configuring Portscan Detection 1798
Rate-Based Attack Prevention 1800
Rate-Based Attack Prevention Examples 1801
detection_filter Keyword Example 1801
Dynamic Rule State Thresholding or Suppression Example 1802
Policy-Wide Rate-Based Detection and Thresholding or Suppression Example 1803

Rate-Based Detection with Multiple Filtering Methods Example 1804
Rate-Based Attack Prevention Options and Configuration 1805
Rate-Based Attack Prevention, Detection Filtering, and Thresholding or Suppression 1807
Configuring Rate-Based Attack Prevention 1807

CHAPTER 89 **Adaptive Profiles** 1811
About Adaptive Profiles 1811
License Requirements for Adaptive Profiles 1812
Requirements and Prerequisites for Adaptive Profiles 1812
Adaptive Profile Updates 1812
Adaptive Profile Updates and Firepower Recommended Rules 1813
Adaptive Profile Options 1813
Configuring Adaptive Profiles 1814

PART XX **Discovery and Identity** 1817

CHAPTER 90 **Introduction to Network Discovery and Identity** 1819
Uses for Host, Application, and User Discovery and Identity Data 1819
Host and Application Detection Fundamentals 1820
Passive Detection of Operating System and Host Data 1820
Active Detection of Operating System and Host Data 1820
Current Identities for Applications and Operating Systems 1821
Current User Identities 1822
Application and Operating System Identity Conflicts 1822
Netflow Data in the Firepower System 1823
Requirements for Using NetFlow Data 1824
Differences between NetFlow and Managed Device Data 1824
About User Identity 1827
Identity Terminology 1827
Best Practices for User Identity 1828
Identity Deployments 1830
The User Activity Database 1830
The Users Database 1831
Firepower System Host and User Limits 1831

Firepower System Host Limit 1832
Firepower System User Limit 1833

CHAPTER 91 **Host Identity Sources** 1835
Overview: Host Data Collection 1835
Requirements and Prerequisites for Host Identity Sources 1836
Determining Which Host Operating Systems the System Can Detect 1836
Identifying Host Operating Systems 1836
Custom Fingerprinting 1837
Managing Fingerprints 1838
Activating and Deactivating Fingerprints 1838
Editing an Active Fingerprint 1839
Editing an Inactive Fingerprint 1839
Creating a Custom Fingerprint for Clients 1840
Creating a Custom Fingerprint for Servers 1842
Host Input Data 1845
Requirements for Using Third-Party Data 1845
Third-Party Product Mappings 1846
Mapping Third-Party Products 1846
Mapping Third-Party Product Fixes 1848
Mapping Third-Party Vulnerabilities 1849
Custom Product Mappings 1850
Creating Custom Product Mappings 1850
Editing Custom Product Mapping Lists 1851
Activating and Deactivating Custom Product Mappings 1851
eStreamer Server Streaming 1852
Choosing eStreamer Event Types 1852
Configuring eStreamer Client Communications 1853
Configuring the Host Input Client 1854
Nmap Scanning 1855
Nmap Remediation Options 1855
Nmap Scanning Guidelines 1859
Example: Using Nmap to Resolve Unknown Operating Systems 1860
Example: Using Nmap to Respond to New Hosts 1862

Managing Nmap Scanning 1863
Adding an Nmap Scan Instance 1863
Editing an Nmap Scan Instance 1864
Adding an Nmap Scan Target 1865
Editing an Nmap Scan Target 1866
Creating an Nmap Remediation 1867
Editing an Nmap Remediation 1869
Running an On-Demand Nmap Scan 1869
Nmap Scan Results 1870
Viewing Nmap Scan Results 1870
Nmap Scan Results Fields 1871
Importing Nmap Scan Results 1872

CHAPTER 92 **Application Detection** 1873
Overview: Application Detection 1873
Application Detector Fundamentals 1874
Identification of Application Protocols in the Web Interface 1875
Implied Application Protocol Detection from Client Detection 1876
Host Limits and Discovery Event Logging 1876
Special Considerations for Application Detection 1877
Requirements and Prerequisites for Application Detection 1878
Custom Application Detectors 1878
Custom Application Detector and User-Defined Application Fields 1879
Configuring Custom Application Detectors 1882
Creating a User-Defined Application 1883
Specifying Detection Patterns in Basic Detectors 1883
Specifying Detection Criteria in Advanced Detectors 1884
Testing a Custom Application Protocol Detector 1885
Viewing or Downloading Detector Details 1886
Sorting the Detector List 1887
Filtering the Detector List 1887
Filter Groups for the Detector List 1887
Navigating to Other Detector Pages 1888
Activating and Deactivating Detectors 1889

Editing Custom Application Detectors 1889
Deleting Detectors 1890

CHAPTER 93 User Identity Sources 1891
About User Identity Sources 1891
The User Agent Identity Source 1893
User Agent Guidelines 1893
Configure the User Agent for User Control 1894
Troubleshoot the User Agent Identity Source 1894
The ISE Identity Source 1895
How to Configure ISE for User Control 1895
ISE Guidelines and Limitations 1896
Configure ISE for User Control 1897
ISE Configuration Fields 1899
Troubleshoot ISE or Cisco TrustSec Issues 1900
The Terminal Services (TS) Agent Identity Source 1901
TS Agent Guidelines 1901
Configure the TS Agent for User Control 1902
Troubleshoot the TS Agent Identity Source 1902
The Captive Portal Identity Source 1903
Captive Portal Guidelines and Limitations 1903
How to Configure the Captive Portal for User Control 1905
Configure the Captive Portal Part 1: Create an Identity Policy 1907
Configure the Captive Portal Part 2: Create a TCP Port Access Control Rule 1908
Configure the Captive Portal Part 3: Create a User Access Control Rule 1909
Configure Captive Portal Part 4: Create an SSL Decrypt-Resign Policy 1910
Configure Captive Portal Part 5: Associate Identity and SSL Policies with the Access Control Policy 1911
Captive Portal Fields 1912
Exclude Applications from Captive Portal 1913
Troubleshoot the Captive Portal Identity Source 1914
The Traffic-Based Detection Identity Source 1915

CHAPTER 94 Network Discovery Policies 1917

Overview: Network Discovery Policies 1917
Requirements and Prerequisites for Network Discovery Policies 1918
Network Discovery Customization 1918
Configuring the Network Discovery Policy 1919
Network Discovery Rules 1919
Configuring Network Discovery Rules 1920
Actions and Discovered Assets 1921
Monitored Networks 1921
Port Exclusions 1924
Zones in Network Discovery Rules 1925
The Traffic-Based Detection Identity Source 1926
Configuring Advanced Network Discovery Options 1928
Network Discovery General Settings 1929
Configuring Network Discovery General Settings 1930
Network Discovery Identity Conflict Settings 1930
Configuring Network Discovery Identity Conflict Resolution 1931
Network Discovery Vulnerability Impact Assessment Options 1931
Enabling Network Discovery Vulnerability Impact Assessment 1932
Indications of Compromise 1932
Enabling Indications of Compromise Rules 1933
Adding NetFlow Exporters to a Network Discovery Policy 1934
Network Discovery Data Storage Settings 1934
Configuring Network Discovery Data Storage 1936
Configuring Network Discovery Event Logging 1936
Adding Network Discovery OS and Server Identity Sources 1937
Troubleshooting Your Network Discovery Strategy 1938

CHAPTER 95 **Realms and Identity Policies** 1941
About Realms and Identity Policies 1941
About Realms 1941
Realms and Trusted Domains 1943
Supported Servers for Realms 1943
Supported Server Object Class and Attribute Names 1944
Troubleshoot Realms and User Downloads 1945

About Identity Policies 1947
License Requirements for Realms 1948
Requirements and Prerequisites for Realms 1948
Create a Realm 1949
Realm Fields 1950
Realm Directory and Download fields 1952
Realms and Identity Policies 1954
Connect Securely to Active Directory 1954
Find the Active Directory Server's Name 1954
Export the Active Directory Server's Root Certificate 1955
Export the Active Directory Server's Root Certificate 1956
Find the Active Directory Server's Name 1958
Configure a Realm Directory 1959
Download Users and Groups 1960
Create an Identity Policy 1961
Create an Identity Rule 1962
Identity Rule Fields 1963
Manage a Realm 1965
Compare Realms 1966
Manage an Identity Policy 1966
Manage an Identity Rule 1967
History for Realms 1967

PART XXI **Correlation and Compliance** 1969

CHAPTER 96 **Compliance White Lists** 1971
Introduction to Compliance White Lists 1971
Compliance White List Target Networks 1972
Compliance White List Host Profiles 1973
Operating System-Specific Host Profiles 1974
Shared Host Profiles 1974
White Violation Triggers 1975
Requirements and Prerequisites for Compliance 1976
Creating a Compliance White List 1976

Setting Target Networks for a Compliance White List 1977
Building White List Host Profiles 1978
Adding an Application Protocol to a Compliance White List 1979
Adding a Client to a Compliance White List 1980
Adding a Web Application to a Compliance White List 1981
Adding a Protocol to a Compliance White List 1981
Managing Compliance White Lists 1982
Editing a Compliance White List 1982
Managing Shared Host Profiles 1984

CHAPTER 97 **Correlation Policies** 1985
Introduction to Correlation Policies and Rules 1985
Requirements and Prerequisites for Compliance 1986
Configuring Correlation Policies 1987
Adding Responses to Rules and White Lists 1987
Managing Correlation Policies 1988
Configuring Correlation Rules 1989
Syntax for Intrusion Event Trigger Criteria 1990
Syntax for Malware Event Trigger Criteria 1993
Syntax for Discovery Event Trigger Criteria 1994
Syntax for User Activity Event Trigger Criteria 1997
Syntax for Host Input Event Trigger Criteria 1998
Syntax for Connection Event Trigger Criteria 1999
Syntax for Traffic Profile Changes 2002
Syntax for Correlation Host Profile Qualifications 2004
Syntax for User Qualifications 2006
Connection Trackers 2007
Adding a Connection Tracker 2008
Syntax for Connection Trackers 2008
Syntax for Connection Tracker Events 2011
Sample Configuration for Excessive Connections From External Hosts 2011
Sample Configuration for Excessive BitTorrent Data Transfers 2013
Snooze and Inactive Periods 2015
Correlation Rule Building Mechanics 2015

Adding and Linking Conditions in Correlation Rules 2017
Using Multiple Values in Correlation Rule Conditions 2017
Managing Correlation Rules 2018
Configuring Correlation Response Groups 2019
Managing Correlation Response Groups 2019

CHAPTER 98 Traffic Profiling 2021
Introduction to Traffic Profiles 2021
Traffic Profile Conditions 2023
Requirements and Prerequisites for Traffic Profiles 2025
Managing Traffic Profiles 2025
Configuring Traffic Profiles 2026
Adding Traffic Profile Conditions 2027
Adding Host Profile Qualifications to a Traffic Profile 2027
Syntax for Traffic Profile Conditions 2028
Syntax for Host Profile Qualifications in a Traffic Profile 2029
Using Multiple Values in a Traffic Profile Condition 2031

CHAPTER 99 Remediations 2033
Requirements and Prerequisites for Remediations 2033
Introduction to Remediations 2033
Cisco ISE EPS Remediations 2034
Configuring ISE EPS Remediations 2035
Cisco IOS Null Route Remediations 2036
Configuring Remediations for Cisco IOS Routers 2037
Nmap Scan Remediations 2041
Set Attribute Value Remediations 2041
Configuring Set Attribute Remediations 2042
Managing Remediation Modules 2043
Managing Remediation Instances 2044
Managing Instances for a Single Remediation Module 2044

PART XXII Reporting and Alerting 2047

CHAPTER 100 **Working with Reports** 2049

Requirements and Prerequisites for Reports 2049

Introduction to Reports 2049

Risk Reports 2050

Generating, Viewing, and Printing Risk Reports 2050

Standard Reports 2051

About Designing Reports 2051

Report Templates 2051

Report Template Fields 2052

Report Template Creation 2053

Report Template Configuration 2056

Managing Report Templates 2067

About Generating Reports 2069

Generating Reports 2069

Report Generation Options 2070

Distributing Reports by Email at Generation Time 2071

Schedule Future Reports 2071

About Working with Generated Reports 2071

Viewing Reports 2072

Downloading Reports 2072

Storing Reports Remotely 2073

Moving Reports to Remote Storage 2073

Deleting Reports 2074

CHAPTER 101 **External Alerting with Alert Responses** 2075

Firepower Management Center Alert Responses 2075

Configurations Supporting Alert Responses 2076

Requirements and Prerequisites for Alert Responses 2076

Creating an SNMP Alert Response 2076

Creating a Syslog Alert Response 2078

Syslog Alert Facilities 2079

Syslog Severity Levels 2080

Creating an Email Alert Response 2081

Configuring Impact Flag Alerting 2081
Configuring Discovery Event Alerting 2082
Configuring AMP for Networks Alerting 2082

CHAPTER 102 External Alerting for Intrusion Events 2085
About External Alerting for Intrusion Events 2085
License Requirements for External Alerting for Intrusion Events 2086
Requirements and Prerequisites for External Alerting for Intrusion Events 2086
Configuring SNMP Alerting for Intrusion Events 2086
Intrusion SNMP Alert Options 2087
Configuring Syslog Alerting for Intrusion Events 2088
Facilities and Priorities for Intrusion Syslog Alerts 2089
Configuring Email Alerting for Intrusion Events 2090
Intrusion Email Alert Options 2090

PART XXIII Event and Asset Analysis Tools 2093

CHAPTER 103 Using the Context Explorer 2095
About the Context Explorer 2095
Differences Between the Dashboard and the Context Explorer 2096
The Traffic and Intrusion Event Counts Time Graph 2096
The Indications of Compromise Section 2097
The Hosts by Indication Graph 2097
The Indications by Host Graph 2097
The Network Information Section 2097
The Operating Systems Graph 2097
The Traffic by Source IP Graph 2098
The Traffic by Source User Graph 2098
The Connections by Access Control Action Graph 2098
The Traffic by Destination IP Graph 2099
The Traffic by Ingress/Egress Security Zone Graph 2099
The Application Information Section 2099
Focusing the Application Information Section 2100
The Traffic by Risk/Business Relevance and Application Graph 2100

The Intrusion Events by Risk/Business Relevance and Application Graph 2100
The Hosts by Risk/Business Relevance and Application Graph 2101
The Application Details List 2101
The Security Intelligence Section 2101
The Security Intelligence Traffic by Category Graph 2102
The Security Intelligence Traffic by Source IP Graph 2102
The Security Intelligence Traffic by Destination IP Graph 2102
The Intrusion Information Section 2102
The Intrusion Events by Impact Graph 2103
The Top Attackers Graph 2103
The Top Users Graph 2103
The Intrusion Events by Priority Graph 2103
The Top Targets Graph 2103
The Top Ingress/Egress Security Zones Graph 2103
The Intrusion Event Details List 2104
The Files Information Section 2104
The Top File Types Graph 2104
The Top File Names Graph 2104
The Files by Disposition Graph 2105
The Top Hosts Sending Files Graph 2105
The Top Hosts Receiving Files Graph 2105
The Top Malware Detections Graph 2106
The Geolocation Information Section 2106
The Connections by Initiator/Responder Country Graph 2106
The Intrusion Events by Source/Destination Country Graph 2106
The File Events by Sending/Receiving Country Graph 2107
The URL Information Section 2107
The Traffic by URL Graph 2107
The Traffic by URL Category Graph 2107
The Traffic by URL Reputation Graph 2108
Requirements and Prerequisites for the Context Explorer 2108
Refreshing the Context Explorer 2108
Setting the Context Explorer Time Range 2109
Minimizing and Maximizing Context Explorer Sections 2109

Drilling Down on Context Explorer Data 2110
Filters in the Context Explorer 2111
Data Type Field Options 2112
Creating a Filter from the Add Filter Window 2113
Creating a Quick Filter from the Context Menu 2114
Saving Filtered Context Explorer Views 2115
Viewing Filter Data 2115
Deleting a Filter 2115

CHAPTER 104 **Using the Network Map** 2117
Requirements and Prerequisites for the Network Map 2117
The Network Map 2117
The Hosts Network Map 2118
The Network Devices Network Map 2119
The Mobile Devices Network Map 2120
The Indications of Compromise Network Map 2120
The Application Protocols Network Map 2120
The Vulnerabilities Network Map 2121
The Host Attributes Network Map 2122
Viewing Network Maps 2122
Custom Network Topologies 2123
Creating Custom Topologies 2124
Importing Networks from the Network Discovery Policy 2124
Manually Adding Networks to Your Custom Topology 2125
Activating and Deactivating Custom Topologies 2125
Editing Custom Topologies 2126

CHAPTER 105 **Incidents** 2127
About Incident Handling 2127
Definition of an Incident 2127
Common Incident Handling Processes 2128
Incident Types in the Firepower System 2130
License Requirements for Incidents 2131
Requirements and Prerequisites for Incidents 2131

Creating Custom Incident Types 2131
Creating an Incident 2132
Editing an Incident 2132
Generating Incident Reports 2133

CHAPTER 106 Using Lookups 2135
Introduction to Lookups 2135
Performing Whois Lookups 2135
Finding URL Category and Reputation 2136
Finding Geolocation Information for an IP Address 2137

PART XXIV Workflows 2139

CHAPTER 107 Workflows 2141
Overview: Workflows 2141
Predefined Workflows 2141
Predefined Intrusion Event Workflows 2142
Predefined Malware Workflows 2143
Predefined File Workflows 2143
Predefined Captured File Workflows 2143
Predefined Connection Data Workflows 2144
Predefined Security Intelligence Workflows 2145
Predefined Host Workflows 2145
Predefined Indications of Compromise Workflows 2146
Predefined Applications Workflows 2146
Predefined Application Details Workflows 2147
Predefined Servers Workflows 2147
Predefined Host Attributes Workflows 2147
The Predefined Discovery Events Workflow 2148
Predefined User Workflows 2148
Predefined Vulnerabilities Workflows 2148
Predefined Third-Party Vulnerabilities Workflows 2148
Predefined Correlation and White List Workflows 2149
Predefined System Workflows 2149

Custom Table Workflows 2150
Using Workflows 2150
Workflow Access by User Role 2151
Workflow Selection 2152
Workflow Pages 2153
Workflow Page Navigation Tools 2155
Workflow Page Traversal Tools 2155
File Trajectory Icons 2155
Host Profile Icons 2156
Threat Score Icons 2156
User Icons 2157
The Workflow Toolbar 2157
Using Drill-Down Pages 2158
Using Table View Pages 2158
Geolocation 2159
Connection Event Graphs 2159
Using Connection Event Graphs 2160
Event Time Constraints 2165
Per-Session Time Window Customization for Events 2166
The Default Time Window for Events 2169
Event View Constraints 2171
Constraining Events 2172
Compound Event View Constraints 2172
Using Compound Event View Constraints 2173
Inter-Workflow Navigation 2173
Bookmarks 2174
Creating Bookmarks 2175
Viewing Bookmarks 2175

CHAPTER 108 **Searching for Events** 2177
Event Searches 2177
Search Constraints 2177
General Search Constraints 2178
Wildcards and Symbols in Searches 2178

Objects and Application Filters in Searches 2179
Time Constraints in Searches 2179
IP Addresses in Searches 2180
Managed Devices in Searches 2181
Ports in Searches 2181
Event Fields in Searches 2181
Performing a Search 2182
Saving a Search 2183
Loading a Saved Search 2184
Query Overrides Via the Shell 2185
Shell-Based Query Management Syntax 2185
Stopping Long-Running Queries 2185

CHAPTER 109 **Custom Workflows** 2187
Introduction to Custom Workflows 2187
Saved Custom Workflows 2187
Custom Workflow Creation 2188
Creating Custom Workflows Based on Non-Connection Data 2189
Creating Custom Connection Data Workflows 2190
Custom Workflow Use and Management 2191
Viewing Custom Workflows Based on Predefined Tables 2192
Viewing Custom Workflows Based on Custom Tables 2192
Editing Custom Workflows 2192

CHAPTER 110 **Custom Tables** 2195
Introduction to Custom Tables 2195
Predefined Custom Tables 2195
Possible Table Combinations 2196
User-Defined Custom Tables 2200
Creating a Custom Table 2200
Modifying a Custom Table 2201
Deleting a Custom Table 2202
Viewing a Workflow Based on a Custom Table 2202
Searching Custom Tables 2202

PART XXV Events and Assets 2205

CHAPTER 111 Connection Logging 2207

About Connection Logging 2207
Connections That Are Always Logged 2208
Other Connections You Can Log 2209
How Rules and Policy Actions Affect Logging 2209
Logging for Fastpathed Connections 2210
Logging for Monitored Connections 2210
Logging for Trusted Connections 2210
Logging for Blocked Connections 2211
Logging for Allowed Connections 2212
Beginning vs End-of-Connection Logging 2213
Firepower Management Center vs External Logging 2214
Limitations of Connection Logging 2215
When Events Appear in the Event Viewer 2215
Best Practices for Connection Logging 2216
Requirements and Prerequisites for Connection Logging 2218
Configure Connection Logging 2218
Logging Connections with Tunnel and Prefilter Rules 2218
Logging Decryptable Connections with SSL Rules 2219
Logging Connections with Security Intelligence 2220
Logging Connections with Access Control Rules 2220
Logging Connections with a Policy Default Action 2221
Limiting Logging of Long URLs 2222

CHAPTER 112 Connection and Security Intelligence Events 2223

About Connection Events 2223
Connection vs. Security Intelligence Events 2223
NetFlow Connections 2224
Connection Summaries (Aggregated Data for Graphs) 2224
Long-Running Connections 2225
Combined Connection Summaries from External Responders 2225

Connection and Security Intelligence Event Fields 2225
About Connection and Security Intelligence Event Fields 2238
A Note About Initiator/Responder, Source/Destination, and Sender/Receiver Fields 2238
Connection Event Reasons 2239
Requirements for Populating Connection Event Fields 2240
Information Available in Connection Event Fields 2242
Using Connection and Security Intelligence Event Tables 2245
Viewing Files and Malware Detected in a Connection 2247
Viewing Intrusion Events Associated with a Connection 2248
Encrypted Connection Certificate Details 2249
Viewing the Connection Summary Page 2249

CHAPTER 113 **Working with Intrusion Events** 2251
About Intrusion Events 2251
Tools for Reviewing and Evaluating Intrusion Events 2251
License Requirements for Intrusion Events 2252
Requirements and Prerequisites for Intrusion Events 2252
Viewing Intrusion Events 2252
About Intrusion Event Fields 2253
Intrusion Event Fields 2254
Intrusion Event Impact Levels 2263
Viewing Connection Data Associated with Intrusion Events 2264
Marking Intrusion Events Reviewed 2265
Viewing Previously Reviewed Intrusion Events 2266
Marking Reviewed Intrusion Events Unreviewed 2266
Preprocessor Events 2266
Preprocessor Generator IDs 2267
Intrusion Event Workflow Pages 2268
Using Intrusion Event Workflows 2269
Intrusion Event Drill-Down Page Constraints 2271
Intrusion Event Table View Constraints 2271
Using the Intrusion Event Packet View 2272
Event Information Fields 2273
Frame Information Fields 2279

Data Link Layer Information Fields 2280
Viewing Network Layer Information 2281
Viewing Transport Layer Information 2283
Viewing Packet Byte Information 2286
Internally Sourced Intrusion Events 2286
The Intrusion Events Clipboard 2286
Generating Clipboard Reports 2286
Deleting Events from the Clipboard 2287
Viewing Intrusion Event Statistics 2288
Host Statistics 2288
Event Overview 2289
Event Statistics 2289
Viewing Intrusion Event Performance Graphs 2290
Intrusion Event Performance Statistics Graph Types 2290
Viewing Intrusion Event Graphs 2294

CHAPTER 114 **File/Malware Events and Network File Trajectory** 2295
About File/Malware Events and Network File Trajectory 2295
File and Malware Events 2296
File and Malware Event Types 2296
File Events 2296
Malware Events 2296
Retrospective Malware Events 2298
Malware Events Generated by AMP for Endpoints 2298
Using File and Malware Event Workflows 2300
File and Malware Event Fields 2300
Malware Event Sub-Types 2308
Information Available in File and Malware Event Fields 2309
View Details About Analyzed Files 2312
File Composition Report 2312
View File Details in AMP Private Cloud 2312
Threat Scores and Dynamic Analysis Summary Reports 2312
Viewing Dynamic Analysis Results in the Cisco Threat Grid Public Cloud 2313
Using Captured File Workflows 2314

Captured File Fields 2314
Stored Files Download 2317
Manually Submit Files for Analysis 2317
Network File Trajectory 2318
Recently Detected Malware and Analyzed Trajectories 2318
Network File Trajectory Detailed View 2319
Network File Trajectory Summary Information 2319
Network File Trajectory Map and Related Events List 2321
Using a Network File Trajectory 2322
Work with Event Data in the AMP for Endpoints Console 2323

CHAPTER 115 **Using Host Profiles** 2325
Requirements and Prerequisites for Host Profiles 2325
Host Profiles 2326
Host Profile Limitations 2327
Viewing Host Profiles 2327
Basic Host Information in the Host Profile 2327
Operating Systems in the Host Profile 2329
Viewing Operating System Identities 2331
Setting the Current Operating System Identity 2332
Operating System Identity Conflicts 2332
Making a Conflicting Operating System Identity Current 2333
Resolving an Operating System Identity Conflict 2333
Servers in the Host Profile 2333
Server Details in the Host Profile 2335
Viewing Server Details 2336
Editing Server Identities 2336
Resolving Server Identity Conflicts 2337
Web Applications in the Host Profile 2338
Deleting Web Applications from the Host Profile 2339
Host Protocols in the Host Profile 2339
Deleting a Protocol From the Host Profile 2340
Indications of Compromise in the Host Profile 2340
VLAN Tags in the Host Profile 2340

User History in the Host Profile 2341
Host Attributes in the Host Profile 2341
Predefined Host Attributes 2341
White List Host Attributes 2342
User-Defined Host Attributes 2342
Creating Text- or URL-Based Host Attributes 2343
Creating Integer-Based Host Attributes 2343
Creating List-Based Host Attributes 2344
Setting Host Attribute Values 2344
White List Violations in the Host Profile 2345
Creating Shared White List Host Profiles 2345
Malware Detections in the Host Profile 2346
Vulnerabilities in the Host Profile 2347
Downloading Patches for Vulnerabilities 2347
Deactivating Vulnerabilities for Individual Hosts 2348
Deactivating Individual Vulnerabilities 2348
Scan Results in the Host Profile 2349
Scanning a Host from the Host Profile 2349

CHAPTER 116 **Working with Discovery Events** 2351
Requirements and Prerequisites for Discovery Events 2351
Discovery and Identity Data in Discovery Events 2351
Viewing Discovery Event Statistics 2352
The Statistics Summary Section 2353
The Event Breakdown Section 2354
The Protocol Breakdown Section 2354
The Application Protocol Breakdown Section 2355
The OS Breakdown Section 2355
Viewing Discovery Performance Graphs 2355
Discovery Performance Graph Types 2356
Using Discovery and Identity Workflows 2356
Discovery and Host Input Events 2358
Discovery Event Types 2358
Host Input Event Types 2362

Viewing Discovery and Host Input Events 2364
Discovery Event Fields 2364
Host Data 2365
Viewing Host Data 2366
Host Data Fields 2366
Creating a Traffic Profile for Selected Hosts 2370
Creating a Compliance White List Based on Selected Hosts 2371
Host Attribute Data 2371
Viewing Host Attributes 2372
Host Attribute Data Fields 2372
Setting Host Attributes for Selected Hosts 2373
Indications of Compromise Data 2373
View and Work with Indications of Compromise Data 2374
Indications of Compromise Data Fields 2376
Editing Indication of Compromise Rule States for a Single Host or User 2376
Viewing Source Events for Indication of Compromise Tags 2377
Resolving Indication of Compromise Tags 2377
Server Data 2378
Viewing Server Data 2378
Server Data Fields 2379
Application and Application Details Data 2381
Viewing Application Data 2381
Application Data Fields 2382
Viewing Application Detail Data 2383
Application Detail Data Fields 2384
Vulnerability Data 2385
Vulnerability Data Fields 2386
Vulnerability Deactivation 2387
Viewing Vulnerability Data 2388
Viewing Vulnerability Details 2389
Deactivating Multiple Vulnerabilities 2389
Third-Party Vulnerability Data 2390
Viewing Third-Party Vulnerability Data 2390
Third-Party Vulnerability Data Fields 2390

Users and User Activity Data 2392
User-Related Fields 2392
User Data 2395
User Activity Data 2398
User Profile and Host History 2400
History for Working with Discovery Events 2402

CHAPTER 117 Correlation and Compliance Events 2403
Viewing Correlation Events 2403
Correlation Event Fields 2404
Using Compliance White List Workflows 2406
Viewing White List Events 2407
White List Event Fields 2408
Viewing White List Violations 2409
White List Violation Fields 2410
Remediation Status Events 2411
Viewing Remediation Status Events 2411
Remediation Status Table Fields 2412
Using the Remediation Status Events Table 2413

CHAPTER 118 Auditing the System 2415
The System Log 2415
Viewing the System Log 2415
Syntax for System Log Filters 2416
About System Auditing 2417
Audit Records 2417
Viewing Audit Records 2417
Suppressing Audit Records 2420
About Sending Audit Logs to an External Location 2423

APPENDIX A Security, Internet Access, and Communication Ports 2425
Security Requirements 2425
Cisco Clouds 2425
Internet Access Requirements 2426

Communication Port Requirements 2427

APPENDIX B **Classic Device Command Line Reference 2431**

About the Classic Device CLI 2431

Classic Device CLI Modes 2432

Classic Device CLI Access Levels 2432

Classic Device CLI Management Commands 2432

configure password 2432

exit 2433

expert 2433

history 2434

logout 2434

? (question mark) 2434

Classic Device CLI Show Commands 2435

access-control-config 2435

alarms 2436

arp-tables 2436

audit-log 2436

audit_cert 2437

bypass 2437

high-availability Commands 2438

config 2438

high-availability ha-statistics 2438

cpu 2438

database Commands 2439

processes 2439

slow-query-log 2440

device-settings 2440

disk 2440

disk-manager 2441

dns 2441

fan-status 2442

fastpath-rules 2442

gui 2442

hostname 2443
hosts 2443
hyperthreading 2443
inline-sets 2444
interfaces 2444
ifconfig 2445
lcd 2445
link-aggregation Commands 2445
configuration 2445
statistics 2446
link-state 2446
log-ips-connection 2447
managers 2447
memory 2447
model 2448
mpls-depth 2448
NAT Commands 2448
active-dynamic 2449
active-static 2449
allocators 2449
config 2449
dynamic-rules 2450
flows 2450
static-rules 2450
netstat 2451
network 2451
network-modules 2451
network-static-routes 2452
ntp 2452
perfstats 2452
portstats 2453
power-supply-status 2453
process-tree 2454
processes 2454

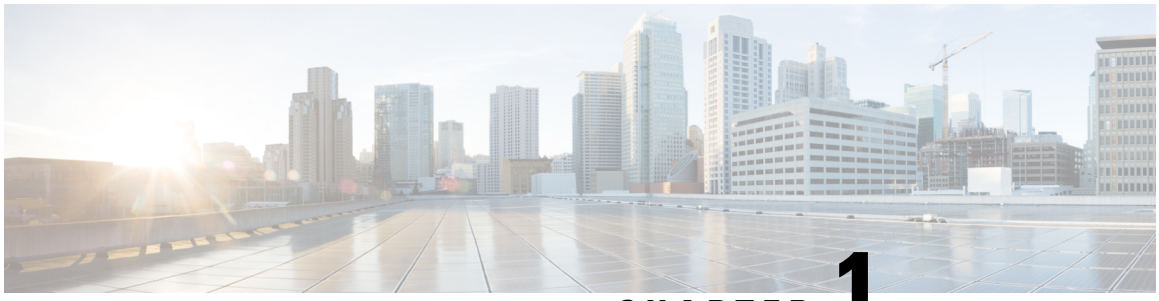
route 2455
routing-table 2455
serial-number 2455
ssl-policy-config 2456
stacking 2456
summary 2457
syslog 2457
time 2457
traffic-statistics 2458
user 2458
users 2459
version 2460
virtual-routers 2460
virtual-switches 2461
vmware-tools 2461
VPN Commands 2462
config 2462
config by virtual router 2462
status 2463
status by virtual router 2463
counters 2463
counters by virtual router 2463
Classic Device CLI Configuration Commands 2464
audit_cert Commands 2464
delete 2464
import 2464
bypass 2465
high-availability 2465
gui 2466
lcd 2466
log-ips-connections 2466
manager Commands 2467
add 2467
delete 2468

mpls-depth 2468
network Commands 2468
dns searchdomains 2468
dns servers 2469
hostname 2469
http-proxy 2469
http-proxy-disable 2470
ipv4 delete 2470
ipv4 dhcp 2470
ipv4 manual 2471
ipv6 delete 2471
ipv6 dhcp 2472
ipv6 manual 2472
ipv6 router 2472
management-interface disable 2473
management-interface disable-event-channel 2473
management-interface disable-management-channel 2474
management-interface enable 2474
management-interface enable-event-channel 2475
management-interface enable-management-channel 2475
management-interface tcpport 2476
management-port 2476
static-routes ipv4 add 2476
static-routes ipv4 delete 2477
static-routes ipv6 add 2477
static-routes ipv6 delete 2477
password 2478
stacking disable 2478
user Commands 2479
access 2479
add 2479
aging 2480
delete 2480
disable 2480

enable 2480
forcereset 2481
maxfailedlogins 2481
minpasswdlen 2481
password 2482
strengthcheck 2482
unlock 2482
vmware-tools 2483
Classic Device CLI System Commands 2483
access-control Commands 2483
archive 2484
clear-rule-counts 2484
rollback 2484
compliance Commands 2484
enable cc 2485
enable ucapl 2485
show 2485
disable-http-user-cert 2486
file Commands 2486
copy 2486
delete 2487
list 2487
secure-copy 2487
generate-troubleshoot 2488
ldapsearch 2489
lockdown-sensor 2489
nat rollback 2490
reboot 2490
restart 2490
support Commands 2491
ssl-client-hello-display 2491
ssl-client-hello-enabled 2491
ssl-client-hello-force-reset 2493
ssl-client-hello-reset 2493

ssl-client-hello-tuning 2494

shutdown 2495



CHAPTER 1

Getting Started With Firepower

Cisco Firepower is an integrated suite of network security and traffic management products, deployed either on purpose-built platforms or as a software solution. The system is designed to help you handle network traffic in a way that complies with your organization's security policy—your guidelines for protecting your network.

In a typical deployment, multiple traffic-sensing *managed devices* installed on network segments monitor traffic for analysis and report to a *manager*:

- Firepower Management Center
- Firepower Device Manager
- Adaptive Security Device Manager (ASDM)

Managers provide a centralized management console with graphical user interface that you can use to perform administrative, management, analysis, and reporting tasks.

This guide focuses on the *Firepower Management Center* managing appliance. For information about the Firepower Device Manager or ASA with FirePOWER Services managed via ASDM, see the guides for those management methods.

- *Cisco Firepower Threat Defense Configuration Guide for Firepower Device Manager*
- *ASA with FirePOWER Services Local Management Configuration Guide*

- Quick Start: Basic Setup, on page 1
- Firepower Devices, on page 5
- Firepower Features, on page 6
- Switching Domains on the Firepower Management Center, on page 10
- The Context Menu, on page 11
- Firepower Online Help and Documentation, on page 12
- Firepower System IP Address Conventions, on page 15
- Additional Resources, on page 15

Quick Start: Basic Setup

The Firepower feature set is powerful and flexible enough to support basic and advanced configurations. Use the following sections to quickly set up a Firepower Management Center and its managed devices to begin controlling and analyzing traffic.

Installing and Performing Initial Setup on Physical Appliances

Procedure

Install and perform initial setup on all physical appliances using the documentation for your appliance:

- **Firepower Management Center**
  - *Cisco Firepower Management Center Getting Started Guide* for your hardware model, available from

    http://www.cisco.com/go/firepower-mc-install

- **Firepower Threat Defense managed devices**

  Important Ignore Firepower Device Manager documents on these pages.

  - Cisco Firepower 4100 Getting Started Guide
  - Cisco Firepower 9300 Getting Started Guide
  - Cisco Firepower Threat Defense for the ASA 5508-X and ASA 5516-X Using Firepower Management Center Quick Start Guide
  - Cisco Firepower Threat Defense for the ASA 5506-X Series Using Firepower Management Center Quick Start Guide
  - Cisco Firepower Threat Defense for the ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, and ASA 5555-X Using Firepower Management Center Quick Start Guide

- **Classic managed devices**
  - Cisco ASA FirePOWER Module Quick Start Guide
  - Cisco Firepower 8000 Series Getting Started Guide
  - Cisco Firepower 7000 Series Getting Started Guide

Deploying Virtual Appliances

Follow these steps if your deployment includes virtual appliances. Use the documentation roadmap to locate the documents listed below: http://www.cisco.com/c/en/us/td/docs/security/firepower/roadmap/firepower-roadmap.html.

Procedure

Step 1 Determine the supported virtual platforms you will use for the Management Center and devices (these may not be the same). See the Cisco Firepower Compatibility Guide.

Step 2 Deploy virtual Firepower Management Centers on the supported Public and Private cloud environment. See, Cisco Secure Firewall Management Center Virtual Getting Started Guide.

Step 3 Deploy virtual devices for your appliance on the supported Public and Private cloud environment. For details, see the following documentation.

- NGIPSv running on VMware: Cisco Firepower NGIPSv Quick Start Guide for VMware
- Cisco Firepower Threat Defense for the ASA 5508-X and ASA 5516-X Using Firepower Management Center Quick Start Guide
- Firepower Threat Defense Virtual running on Public and Private cloud environments, see Cisco Secure Firewall Threat Defense Virtual Getting Started Guide, Version 7.3.

Logging In for the First Time

Before you begin

- Prepare your appliances as described in Installing and Performing Initial Setup on Physical Appliances, on page 2 or Deploying Virtual Appliances, on page 2.

Procedure

Step 1 Log in to the Firepower Management Center web interface with **`admin`** as the username and **`Admin123`** as the password. Change the password for this account as described in the *Quick Start Guide* for your appliance.

Step 2 Set a time zone for this account as described in Setting Your Default Time Zone, on page 37.

Step 3 Add licenses as described in Licensing the Firepower System, on page 105.

Step 4 Register managed devices as described in Add a Device to the FMC, on page 229.

Step 5 Configure your managed devices as described in:

- Introduction to IPS Device Deployment and Configuration, on page 503, to configure passive or inline interfaces on 7000 Series or 8000 Series devices
- Interface Overview for Firepower Threat Defense, on page 609, to configure transparent or routed mode on Firepower Threat Defense devices
- Interface Overview for Firepower Threat Defense, on page 609, to configure interfaces on Firepower Threat Defense devices

What to do next

- Begin controlling and analyzing traffic by configuring basic policies as described in Setting Up Basic Policies and Configurations, on page 4.

Setting Up Basic Policies and Configurations

You must configure and deploy basic policies in order to see data in the dashboard, Context Explorer, and event tables.

Note This is not a full discussion of policy or feature capabilities. For guidance on other features and more advanced configurations, see the rest of this guide.

Before you begin

- Log into the web interface, set your time zone, add licenses, register devices, and configure devices as described in Logging In for the First Time, on page 3.

Procedure

Step 1 Configure an access control policy as described in Creating a Basic Access Control Policy, on page 1204.

- In most cases, Cisco suggests setting the Balanced Security and Connectivity intrusion policy as your default action. For more information, see Access Control Policy Default Action, on page 1185 and System-Provided Network Analysis and Intrusion Policies, on page 1453.
- In most cases, Cisco suggests enabling connection logging to meet the security and compliance needs of your organization. Consider the traffic on your network when deciding which connections to log so that you do not clutter your displays or overwhelm your system. For more information, see About Connection Logging, on page 2207.

Step 2 Apply the system-provided default health policy as described in Applying Health Policies, on page 284.

Step 3 Customize a few of your system configuration settings:

- If you want to allow inbound connections for a service (for example, SNMP or the syslog), modify the ports in the access list as described in Configure an Access List, on page 896.
- Understand and consider editing your database event limits as described in Configuring Database Event Limits, on page 879.
- If you want to change the display language, edit the language setting as described in Set the Language for the Web Interface, on page 907.
- If your organization restricts network access using a proxy serverand you did not configure proxy settings during initial configuration, edit your proxy settings as described in Modify FMC Management Interfaces, on page 886.

Step 4 Customize your network discovery policy as described in Configuring the Network Discovery Policy, on page 1919. By default, the network discovery policy analyzes all traffic on your network. In most cases, Cisco suggests restricting discovery to the addresses in RFC 1918.

Step 5 Consider customizing these other common settings:

- If you do not want to display message center pop-ups, disable notifications as described in Configuring Notification Behavior, on page 317.

- If you want to customize the default values for system variables, understand their use as described in Variable Sets, on page 412.
- If you want to update the Geolocation Database, update manually or on a scheduled basis as described in Update the Geolocation Database, on page 141.
- If you want to create additional locally authenticated user accounts to access the FMC, see Creating a User Account, on page 62.
- If you want to use LDAP or RADIUS external authentication to allow access to the FMC, see External Authentication, on page 71.

Step 6 Deploy configuration changes; see Deploy Configuration Changes, on page 346.

What to do next

- Review and consider configuring other features described in Firepower Features, on page 6 and the rest of this guide.

Firepower Devices

In a typical deployment, multiple traffic-handling devices report to one Firepower Management Center, which you use to perform administrative, management, analysis, and reporting tasks.

Classic Devices

Classic devices run next-generation IPS (NGIPS) software. They include:

- Firepower 7000 series and Firepower 8000 series physical devices.
- NGIPSv, hosted on VMware.
- ASA with FirePOWER Services, available on select ASA 5500-X series devices. The ASA provides the first-line system policy, and then passes traffic to an ASA FirePOWER module for discovery and access control.

  Note that you must use the ASA CLI or ASDM to configure the ASA-based features on an ASA FirePOWER device. This includes device high availability, switching, routing, VPN, NAT, and so on. You cannot use the FMC to configure ASA FirePOWER interfaces, and the FMC GUI does not display ASA interfaces when the ASA FirePOWER is deployed in SPAN port mode. Also, you cannot use the FMC to shut down, restart, or otherwise manage ASA FirePOWER processes.

Firepower Threat Defense Devices

A Firepower Threat Defense (FTD) device is a next-generation firewall (NGFW) that also has NGIPS capabilities. NGFW and platform features include site-to-site VPN, robust routing, NAT, clustering, and other optimizations in application inspection and access control.

FTD is available on a wide range of physical and virtual platforms.

Compatibility

For details on manager-device compatibility, including the software compatible with specific device models, virtual hosting environments, operating systems, and so on, see the Cisco Firepower Release Notes and Cisco Firepower Compatibility Guide.

Firepower Features

These tables list some commonly used Firepower features.

Appliance and System Management Features

To locate unfamiliar documents, see: http://www.cisco.com/c/en/us/td/docs/security/firepower/roadmap/firepower-roadmap.html.

If you want to...	Configure...	As described in...
Manage user accounts for logging in to your Firepower appliances	Firepower authentication	Firepower System User Authentication, on page 69
Monitor the health of system hardware and software	Health monitoring policy	About Health Monitoring, on page 275
Back up data on your appliance	Backup and restore	Backup and Restore, on page 155
Upgrade to a new Firepower version	System updates	Cisco Firepower Management Center Upgrade Guide, Version 6.0–7.0 Firepower Release Notes
Baseline your physical appliance	Restore to factory defaults (reimage)	The Cisco Firepower Management Center Upgrade Guide, Version 6.0–7.0, for a list of links to instructions on performing fresh installations.
Update the VDB, intrusion rule updates, or GeoDB on your appliance	Vulnerability Database (VDB) updates, intrusion rule updates, or Geolocation Database (GeoDB) updates	System Updates, on page 137
Apply licenses in order to take advantage of license-controlled functionality	Classic or Smart licensing	About Firepower Licenses, on page 105

If you want to...	Configure...	As described in...
Ensure continuity of appliance operations	Managed device high availability and/or Firepower Management Center high availability	About 7000 and 8000 Series Device High Availability, on page 517 About Firepower Threat Defense High Availability, on page 683 About Firepower Management Center High Availability, on page 201
Combine processing resources of multiple 8000 Series devices	Device stacking	About Device Stacks, on page 533
Configure a device to route traffic between two or more interfaces	Routing	Virtual Routers, on page 1115 Routing Overview for Firepower Threat Defense, on page 753
Configure packet switching between two or more networks	Device switching	Virtual Switches, on page 1105 Configure Bridge Group Interfaces, on page 632
Translate private addresses into public addresses for internet connections	Network Address Translation (NAT)	NAT Policy Configuration, on page 987 Network Address Translation (NAT) for Firepower Threat Defense, on page 1005
Establish a secure tunnel between managed Firepower Threat Defense or 7000/8000 Series devices	Site-to-Site virtual private network (VPN)	VPN Overview for Firepower Threat Defense, on page 839
Segment user access to managed devices, configurations, and events	Multitenancy using domains	Introduction to Multitenancy Using Domains, on page 335
View and manage appliance configuration using a REST API client	REST API and REST API Explorer	REST API Preferences, on page 924 *Firepower REST API Quick Start Guide*
Troubleshoot issues	N/A	Troubleshooting the System, on page 311

High Availability and Scalability Features by Platform

High availability configurations (sometimes called failover) ensure continuity of operations. Clustered and stacked configurations group multiple devices together as a single logical device, achieving increased throughput and redundancy.

Platform	High Availability	Clustering	Stacking
Firepower Management Center	Yes Except MC750	—	—
Firepower Management Center Virtual	—	—	—
Firepower Threat Defense Virtual: • VMware • KVM	Yes	—	—
Firepower Threat Defense Virtual (public cloud): • AWS • Public Cloud: Azure	—	—	—
Firepower Threat Defense: • ASA 5500-X series	Yes	—	—
Firepower Threat Defense: • Firepower 4100/9300 chassis	Yes	Yes	—
Firepower NGIPS: • Firepower 7010, 7020, 7030, 7050 • Firepower 7110, 7115, 7120, 7125 • Firepower 8120, 8130 • AMP 7150, 8050, 8150	Yes	—	—
Firepower NGIPS: • Firepower 8140 • Firepower 8250, 8260, 8270, 8290 • Firepower 8350, 8360, 8370, 8390 • AMP 8350	Yes	—	Yes
Firepower NGIPS: • ASA FirePOWER • NGIPSv	—	—	—

Related Topics

About 7000 and 8000 Series Device High Availability, on page 517

About Firepower Threat Defense High Availability, on page 683
About Firepower Management Center High Availability, on page 201

Features for Detecting, Preventing, and Processing Potential Threats

To locate unfamiliar documents, see: http://www.cisco.com/c/en/us/td/docs/security/firepower/roadmap/firepower-roadmap.html.

If you want to...	Configure...	As described in...
Inspect, log, and take action on network traffic	Access control policy, the parent of several other policies	Introduction to Access Control, on page 1185
Block or monitor connections to or from IP addresses, URLs, and/or domain names	Security Intelligence within your access control policy	About Security Intelligence, on page 1249
Control the websites that users on your network can access	URL filtering within your policy rules	URL Filtering, on page 1229
Monitor malicious traffic and intrusions on your network	Intrusion policy	Intrusion Policy Basics, on page 1479
Block encrypted traffic without inspection Inspect encrypted or decrypted traffic	SSL policy	SSL Policies Overview, on page 1337
Tailor deep inspection to encapsulated traffic and improve performance with fastpathing	Prefilter policy	About Prefiltering, on page 1273
Rate limit network traffic that is allowed or trusted by access control	Quality of Service (QoS) policy	About QoS Policies, on page 673
Allow or block files (including malware) on your network	File/malware policy	File Policies and Malware Protection, on page 1401
Configure passive or active user authentication to perform user awareness and user control	User awareness, user identity, identity policies	About Realms and Identity Policies, on page 1941 About Identity Policies, on page 1947
Collect host, application, and user data from traffic on your network to perform user awareness	Network Discovery policies	Overview: Network Discovery Policies, on page 1917
Perform application detection and control	Application detectors	Overview: Application Detection, on page 1873
Troubleshoot issues	N/A	Troubleshooting the System, on page 311

Integration with External Tools

To locate unfamiliar documents, see: http://www.cisco.com/c/en/us/td/docs/security/firepower/roadmap/firepower-roadmap.html.

If you want to...	Configure...	As described in...
Automatically launch remediations when conditions on your network violate an associated policy	Remediations	Introduction to Remediations, on page 2033 *Firepower System Remediation API Guide*
Stream event data from a Firepower Management Center to a custom-developed client application	eStreamer integration	eStreamer Server Streaming, on page 1852 *Firepower System eStreamer Integration Guide*
Query database tables on a Firepower Management Center using a third-party client	External database access	External Database Access Settings, on page 877 *Firepower System Database Access Guide*
Augment discovery data by importing data from third-party sources	Host input	Host Input Data, on page 1845 *Firepower System Host Input API Guide*
Troubleshoot issues	N/A	Troubleshooting the System, on page 311

Switching Domains on the Firepower Management Center

In a multidomain deployment, user role privileges determine which domains a user can access and which privileges the user has within each of those domains. You can associate a single user account with multiple domains and assign different privileges for that user in each domain. For example, you can assign a user read-only privileges in the Global domain, but Administrator privileges in a descendant domain.

Users associated with multiple domains can switch between domains within the same web interface session.

Under your user name in the toolbar, the system displays a tree of available domains. The tree:

- Displays ancestor domains, but may disable access to them based on the privileges assigned to your user account.
- Hides any other domain your user account cannot access, including sibling and descendant domains.

When you switch to a domain, the system displays:

- Data that is relevant to that domain only.
- Menu options determined by the user role assigned to you for that domain.

Procedure

From the drop-down list under your user name, choose the domain you want to access.

The Context Menu

Certain pages in the Firepower System web interface support a right-click (most common) or left-click context menu that you can use as a shortcut for accessing other features in the Firepower System. The contents of the context menu depend where you access it—not only the page but also the specific data.

For example:

- IP address hotspots provide information about the host associated with that address, including any available whois and host profile information.
- SHA-256 hash value hotspots allow you to add a file's SHA-256 hash value to the clean list or custom detection list, or view the entire hash value for copying.

On pages or locations that do not support the Firepower System context menu, the normal context menu for your browser appears.

Policy Editors

Many policy editors contain hotspots over each rule. You can insert new rules and categories; cut, copy, and paste rules; set the rule state; and edit the rule.

Intrusion Rules Editor

The intrusion rules editor contains hotspots over each intrusion rule. You can edit the rule, set the rule state, configure thresholding and suppression options, and view rule documentation.

Event Viewer

Event pages (the drill-down pages and table views available under the Analysis menu) contain hotspots over each event, IP address, URL, DNS query, and certain files' SHA-256 hash values. While viewing most event types, you can:

- View related information in the Context Explorer.
- Drill down into event information in a new window.
- View the full text in places where an event field contains text too long to fully display in the event view, such as a file's SHA-256 hash value, a vulnerability description, or a URL.

While viewing connection events, you can add items to the default Security Intelligence Block and Do Not Block lists:

- An IP address, from an IP address hotspot.
- A URL or domain name, from a URL hotspot.
- A DNS query, from a DNS query hotspot.

While viewing captured files, file events, and malware events, you can:

- Add a file to or remove a file from the clean list or custom detection list.
- Download a copy of the file.
- View nested files inside an archive file.
- Download the parent archive file for a nested file.
- View the file composition.
- Submit the file for local malware and dynamic analysis.

While viewing intrusion events, you can perform similar tasks to those in the intrusion rules editor or an intrusion policy:

- Edit the triggering rule.
- Set the rule state, including disabling the rule.
- Configure thresholding and suppression options.
- View rule documentation.

Intrusion Event Packet View

Intrusion event packet views contain IP address hotspots. The packet view uses a left-click context menu.

Dashboard

Many dashboard widgets contain hotspots to view related information in the Context Explorer. Dashboard widgets can also contain IP address and SHA-256 hash value hotspots.

Context Explorer

The Context Explorer contains hotspots over its charts, tables, and graphs. If you want to examine data from graphs or lists in more detail than the Context Explorer allows, you can drill down to the table views of the relevant data. You can also view related host, user, application, file, and intrusion rule information.

The Context Explorer uses a left-click context menu, which also contains filtering and other options unique to the Context Explorer.

Related Topics

Security Intelligence Lists and Feeds, on page 427

Firepower Online Help and Documentation

You can reach the online help from the web interface:

- By clicking the context-sensitive help link on each page
- By choosing **Help** > **Online**

You can find additional documentation related to the Firepower system using the documentation roadmap: http://www.cisco.com/c/en/us/td/docs/security/firepower/roadmap/firepower-roadmap.html.

Top-Level Documentation Listing Pages for FMC Deployments

The following documents may be helpful when configuring Firepower Management Center deployments, Version 6.0+.

Note Some of the linked documents are not applicable to Firepower Management Center deployments. For example, some links on Firepower Threat Defense pages are specific to deployments managed by Firepower Device Manager, and some links on hardware pages are unrelated to FMC. To avoid confusion, pay careful attention to document titles. Also, some documents cover multiple products and therefore may appear on multiple product pages.

Firepower Management Center

- Firepower Management Center hardware appliances:

  http://www.cisco.com/c/en/us/support/security/defense-center/tsd-products-support-series-home.html

- Firepower Management Center Virtual appliances:
  - http://www.cisco.com/c/en/us/support/security/defense-center-virtual-appliance/tsd-products-support-series-home.html
  - http://www.cisco.com/c/en/us/support/security/defense-center/tsd-products-support-series-home.html

Firepower Threat Defense, also called NGFW (Next Generation Firewall) devices

- Firepower Threat Defense software:

  http://www.cisco.com/c/en/us/support/security/firepower-ngfw/tsd-products-support-series-home.html

- Firepower Threat Defense Virtual:

  http://www.cisco.com/c/en/us/support/security/firepower-ngfw-virtual/tsd-products-support-series-home.html

- Firepower 4100 series:

  https://www.cisco.com/c/en/us/support/security/firepower-4100-series/tsd-products-support-series-home.html

- Firepower 9300:

  https://www.cisco.com/c/en/us/support/security/firepower-9000-series/tsd-products-support-series-home.html

Classic devices, also called NGIPS (Next Generation Intrusion Prevention System) devices

- ASA with FirePOWER Services:
  - ASA 5500-X with FirePOWER Services:
    - https://www.cisco.com/c/en/us/support/security/asa-firepower-services/tsd-products-support-series-home.html

- https://www.cisco.com/c/en/us/support/security/asa-5500-series-next-generation-firewalls/tsd-products-support-series-home.html

- Firepower 8000 series:

  https://www.cisco.com/c/en/us/support/security/firepower-8000-series-appliances/tsd-products-support-series-home.html

- Firepower 7000 series:

  https://www.cisco.com/c/en/us/support/security/firepower-7000-series-appliances/tsd-products-support-series-home.html

- AMP for Networks:

  https://www.cisco.com/c/en/us/support/security/amp-appliances/tsd-products-support-series-home.html

- NGIPSv (virtual device):

  https://www.cisco.com/c/en/us/support/security/ngips-virtual-appliance/tsd-products-support-series-home.html

License Statements in the Documentation

The License statement at the beginning of a section indicates which Classic or Smart license you must assign to a managed device in the Firepower System to enable the feature described in the section.

Because licensed capabilities are often additive, the license statement provides only the highest required license for each feature.

An "or" statement in a License statement indicates that you must assign a particular license to the managed device to enable the feature described in the section, but an additional license can add functionality. For example, within a file policy, some file rule actions require that you assign a Protection license to the device while others require that you assign a Malware license.

For more information about licenses, see About Firepower Licenses, on page 105.

Related Topics

About Firepower Licenses, on page 105

Supported Devices Statements in the Documentation

The Supported Devices statement at the beginning of a chapter or topic indicates that a feature is supported only on the specified device series, family, or model. For example, many features are supported only on Firepower Threat Defense devices.

For more information on platforms supported by this release, see the release notes.

Access Statements in the Documentation

The Access statement at the beginning of each procedure in this documentation indicates the predefined user roles required to perform the procedure. Any of the listed roles can perform the procedure.

Users with custom roles may have permission sets that differ from those of the predefined roles. When a predefined role is used to indicate access requirements for a procedure, a custom role with similar permissions also has access. Some users with custom roles may use slightly different menu paths to reach configuration pages. For example, users who have a custom role with only intrusion policy privileges access the network analysis policy via the intrusion policy instead of the standard path through the access control policy.

For more information about user roles, see Predefined User Roles, on page 42 and Custom User Roles, on page 43.

Firepower System IP Address Conventions

You can use IPv4 Classless Inter-Domain Routing (CIDR) notation and the similar IPv6 prefix length notation to define address blocks in many places in the Firepower System.

When you use CIDR or prefix length notation to specify a block of IP addresses, the Firepower System uses **only** the portion of the network IP address specified by the mask or prefix length. For example, if you type 10.1.2.3/8, the Firepower System uses 10.0.0.0/8.

In other words, although Cisco recommends the standard method of using a network IP address on the bit boundary when using CIDR or prefix length notation, the Firepower System does not require it.

Additional Resources

The Firewalls Community is an exhaustive repository of reference material that complements our extensive documentation. This includes links to 3D models of our hardware, hardware configuration selector, product collateral, configuration examples, troubleshooting tech notes, training videos, lab and Cisco Live sessions, social media channels, Cisco Blogs and all the documentation published by the Technical Publications team.

Some of the individuals posting to community sites or video sharing sites, including the moderators, work for Cisco Systems. Opinions expressed on those sites and in any corresponding comments are the personal opinions of the original authors, not of Cisco. The content is provided for informational purposes only and is not meant to be an endorsement or representation by Cisco or any other party.

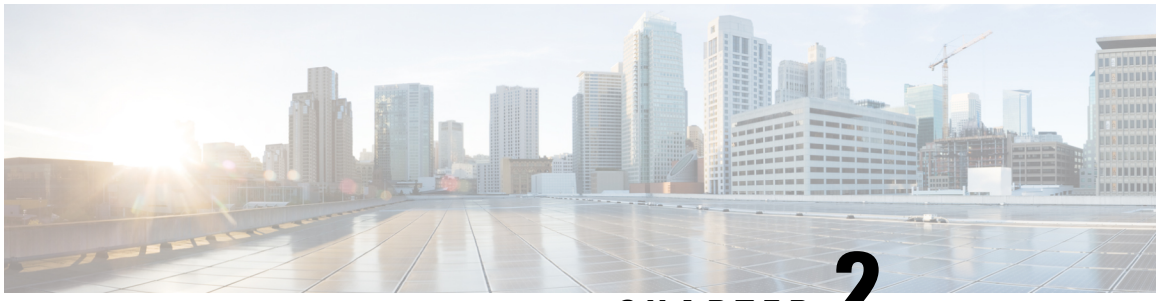
Note Some of the videos, technical notes, and reference material in the Firewalls Community points to older versions of the FMC. Your version of the FMC and the version referenced in the videos or technical notes might have differences in the user interface that cause the procedures not to be identical.



PART I

Your User Account

- Logging into the Firepower System, on page 19
- Specifying User Preferences, on page 31



CHAPTER 2

Logging into the Firepower System

The following topics describe how to log into the Firepower System:

- Firepower System User Accounts, on page 19
- Firepower System User Interfaces, on page 21
- Logging Into the Firepower Management Center Web Interface, on page 24
- Logging Into the Web Interface of a 7000 or 8000 Series Device, on page 25
- Logging Into the Firepower Management Center with CAC Credentials, on page 26
- Logging Into a 7000 or 8000 Series Device with CAC Credentials, on page 26
- Logging Into the CLI on 7000/8000 Series, ASA FirePOWER, and NGIPSv Devices, on page 27
- Logging Into the Command Line Interface on Firepower Threat Defense Devices, on page 28
- Logging Out of a Firepower System Web Interface, on page 29

Firepower System User Accounts

You must provide a username and password to obtain local access to the web interface, shell, or CLI on an FMC or managed device. On managed devices, CLI users with Config level access can use the `expert` command to access the Linux shell. On the FMC, all CLI users can use the `expert` command. The FTD and FMC can be configured to use external authentication, storing user credentials on an external LDAP or RADIUS server; you can withhold or provide CLI/shell access rights to external users.

The FMC CLI provides a single **admin** user who has access to all commands. The features FMC web interface users can access are controlled by the privileges an administrator grants to the user account. On managed devices, the features that users can access for both the CLI and the web interface are controlled by the privileges an administrator grants to the user account.

Note The system audits user activity based on user accounts; make sure that users log into the system with the correct account.

Caution All FMC CLI users and, on managed devices, users with Config level CLI access can obtain root privileges in the Linux shell, which can present a security risk. For system security reasons, we strongly recommend:

- If you establish external authentication, make sure that you restrict the list of users with CLI/shell access appropriately.
- When granting CLI access privileges on managed devices, restrict the list of internal users with Config level CLI access.
- Do not establish Linux shell users; use only the pre-defined **admin** user and users created by the **admin** user within the CLI.

Caution We strongly recommend that you do not use the Linux shell unless directed by Cisco TAC or explicit instructions in the Firepower user documentation.

Different appliances support different types of user accounts, each with different capabilities.

Firepower Management Centers

Firepower Management Centers support the following user account types:

- A pre-defined **admin** account for web interface access, which has the administrator role and can be managed through the web interface.
- Custom user accounts, which provide web interface access and which **admin** users and users with administrator privileges can create and manage.
- A pre-defined **admin** account for shell access, which can obtain root privileges.

Caution For system security reasons, Cisco strongly recommends that you not establish additional Linux shell users on any appliance.

7000 & 8000 Series Devices

7000 & 8000 Series devices support the following user account types:

- A pre-defined **admin** account which can be used for all forms of access to the device.
- Custom user accounts, which **admin** users and users with the administrator role can create and manage.

The 7000 & 8000 Series supports external authentication for users.

NGIPSv Devices

NGIPSv devices support the following user account types:

- A pre-defined **admin** account which can be used for all forms of access to the device.
- Custom user accounts, which **admin** users and users with Config access can create and manage.

The NGIPSv does not support external authentication for users.

Firepower Threat Defense and Firepower Threat Defense Virtual Devices

Firepower Threat Defense and Firepower Threat Defense Virtual devices support the following user account types:

- A pre-defined **admin**account which can be used for all forms of access to the device.
- Custom user accounts, which **admin** users and users with Config access can create and manage.

The Firepower Threat Defense does not support external authentication for SSH or HTTP users.

ASA FirePOWER Devices

The ASA FirePOWER module supports the following user account types:

- A pre-defined **admin** account.
- Custom user accounts, which **admin** users and users with Config access can create and manage.

The ASA FirePOWER module does not support external authentication for users. Accessing ASA devices via the ASA CLI and ASDM is described in the *Cisco ASA Series General Operations CLI Configuration Guide* and the *Cisco ASA Series General Operations ASDM Configuration Guide*.

Firepower System User Interfaces

Depending on appliance type, you can interact with Firepower appliances using a web-based GUI, auxiliary CLI, or the Linux shell. In a Firepower Management Center deployment, you perform most configuration tasks from the FMC GUI. Only a few tasks require that you access the appliance directly using the CLI or Linux shell. We strongly discourage using the Linux shell unless directed by Cisco TAC or explicit instructions in the Firepower user documentation.

For information on browser requirements, see the Firepower Release Notes.

Appliance	Web-Based GUI	Auxiliary CLI	Linux Shell
Firepower Management Center	• Supported for predefined **admin** user and custom user accounts. • Can be used for administrative, management, and analysis tasks.	—	• Supported for predefined **admin** user and custom external user accounts. • Accessible using an SSH, serial, or keyboard and monitor connection. • Should be used only for administration and troubleshooting directed by Cisco TAC or by explicit instructions in the FMC documentation.

Appliance	Web-Based GUI	Auxiliary CLI	Linux Shell
7000 & 8000 Series devices	• Supported for predefined **admin** user and custom user accounts. • Can be used for initial setup, basic analysis, and configuration tasks only.	• Supported for predefined **admin** user and custom user accounts. • Accessible using an SSH, serial, or keyboard and monitor connection. • Can be used for setup and troubleshooting directed by Cisco TAC.	• Supported for predefined **admin** user and custom user accounts. • Accessible by CLI users with Config access using the `expert` command. • Should be used only for administration and troubleshooting directed by Cisco TAC or by explicit instructions in the FMC documentation.
Firepower Threat Defense Firepower Threat Defense Virtual	—	• Supported for predefined **admin** user and custom user accounts. • Accessible in physical devices using an SSH, serial, or keyboard and monitor connection. Accessible in virtual devices via SSH or VM console. • Can be used for setup and troubleshooting directed by Cisco TAC.	• Supported for predefined **admin** user and custom user accounts. • Accessible by CLI users with Config access using the `expert` command. • Should be used only for administration and troubleshooting directed by Cisco TAC or by explicit instructions in the FMC documentation..
NGIPSv	—	• Supported for predefined **admin** user and custom user accounts • Accessible using an SSH connection or VM console • Can be used for setup and troubleshooting directed by Cisco TAC.	• Supported for predefined **admin** user and custom user accounts • Accessible by CLI users with Config access using the `expert` command • Should be used only for administration and troubleshooting directed by Cisco TAC or explicit instructions in the FMC documentation..

Appliance	Web-Based GUI	Auxiliary CLI	Linux Shell
ASA FirePOWER module	—	• Supported for predefined **admin** user and custom user accounts. • Accessible using an SSH connection. Also accessible using a keyboard and monitor connection for ASA 5585-X devices (hardware module), or the console port for other ASA 5500-X series devices (software modules). • Can be used for configuration and management tasks.	• Supported for predefined **admin** user and custom user accounts • Accessible by CLI users with Config access using the `expert` command • Should be used only for administration and troubleshooting directed by Cisco TAC or by explicit instructions in the FMC documentation..

Related Topics

Managing User Accounts, on page 61

Web Interface Considerations

- If your organization uses Common Access Cards (CACs) for authentication, external users authenticated with LDAP can use CAC credentials to obtain access to the web interface of an appliance.
- The first time you visit the appliance home page during a web session, you can view information about your last login session for that appliance. You can see the following information about your last login:
  - the day of the week, month, date, and year of the login
  - the appliance-local time of the login in 24-hour notation
  - the host and domain name last used to access the appliance
- The menus and menu options listed at the top of the default home page are based on the privileges for your user account. However, the links on the default home page include options that span the range of user account privileges. If you click a link that requires different privileges from those granted to your account, the system displays a warning message and logs the activity.
- Some processes that take a significant amount of time may cause your web browser to display a message that a script has become unresponsive. If this occurs, make sure you allow the script to continue until it finishes.

Related Topics

Specifying Your Home Page, on page 32

Session Timeout

By default, the Firepower System automatically logs you out of a session after 1 hour of inactivity, unless you are otherwise configured to be exempt from session timeout.

Users with the Administrator role can change the session timeout interval for an appliance via the following settings:

Appliance	Setting
Firepower Management Center	**System** > **Configuration** > **Shell Timeout**
7000 & 8000 Series devices	**Devices** > **Platform Settings** > **Shell Timeout**

Related Topics

Configure Session Timeouts, on page 917

Logging Into the Firepower Management Center Web Interface

Users are restricted to a single active session. If you try to log in with a user account that already has an active session, the system prompts you to terminate the other session or log in as a different user.

In a NAT environment where multiple FMCs share the same IP address:

- Each FMC can support only one login session at a time.
- To access different FMCs, use a different browser for each login (for example Firefox and Chrome), or set the browser to incognito or private mode.

Before you begin

- If you do not have access to the web interface, contact your system administrator to modify your account privileges, or log in as a user with Administrator access and modify the privileges for the account.
- Create user accounts as described in Creating a User Account, on page 62.

Procedure

Step 1 Direct your browser to **https://***ipaddress_or_hostname*/, where *ipaddress* or *hostname* corresponds to your FMC.

Step 2 In the **Username** and **Password** fields, enter your user name and password. Pay attention to the following guidelines:

- User names are *not* case-sensitive.
- In a multidomain deployment, prepend the user name with the domain where your user account was created. You are not required to prepend any ancestor domains. For example, if your user account was created in SubdomainB, which has an ancestor DomainA, enter your user name in the following format:

```
SubdomainB\username
```

- If your organization uses SecurID® tokens when logging in, append the token to your SecurID PIN and use that as your password to log in. For example, if your PIN is `1111` and the SecurID token is `222222`, enter `1111222222`. You must have already generated your SecurID PIN before you can log into the Firepower System.

Step 3 Click **Login**.

Related Topics

Session Timeout, on page 24

Logging Into the Web Interface of a 7000 or 8000 Series Device

Users are restricted to a single active session on a 7000 & 8000 Series device. If you try to log in with a user account that already has an active session, the system prompts you to terminate the other session or log in as a different user.

Before you begin

- If you do not have access to the web interface, contact your system administrator to modify your account privileges, or log in as a user with Administrator access and modify the privileges for the account.
- Complete the initial setup process and create user accounts as described in the Firepower getting started guide appropriate to the device, and Creating a User Account, on page 62.

Procedure

Step 1 Direct your browser to `https://hostname/`, where *`hostname`* corresponds to the host name of the managed device you want to access.

Step 2 In the **Username** and **Password** fields, enter your user name and password. Pay attention to the following guidelines:

- User names are *not* case-sensitive.
- If your organization uses SecurID® tokens when logging in, append the token to your SecurID PIN and use that as your password to log in. For example, if your PIN is `1111` and the SecurID token is `222222`, enter `1111222222`. You must have already generated your SecurID PIN before you can log into the Firepower System.

Step 3 Click **Login**.

Related Topics

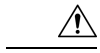
Session Timeout, on page 24

Logging Into the Firepower Management Center with CAC Credentials

Users are restricted to a single active session. If you try to log in with a user account that already has an active session, the system prompts you to terminate the other session or log in as a different user.

In a NAT environment where multiple FMCs share the same IP address:

- Each FMC can support only one login session at a time.
- To access different FMCs, use a different browser for each login (for example Firefox and Chrome), or set the browser to incognito or private mode.



Caution Do **not** remove a CAC during an active browsing session. If you remove or replace a CAC during a session, your web browser terminates the session and the system logs you out of the web interface.

Before you begin

- If you do not have access to the web interface, contact your system administrator to modify your account privileges, or log in as a user with Administrator access and modify the privileges for the account.
- Create user accounts as described in the Creating a User Account, on page 62.
- Configure CAC authentication and authorization as described in Configuring CAC Authentication, on page 74.

Procedure

Step 1 Insert a CAC as instructed by your organization.

Step 2 Direct your browser to **https://***ipaddress_or_hostname/*, where *ipaddress* or *hostname* corresponds to your FMC.

Step 3 If prompted, enter the PIN associated with the CAC you inserted in step 1.

Step 4 If prompted, choose the appropriate certificate from the drop-down list.

Step 5 Click **Continue**.

Related Topics

CAC Authentication, on page 74
Session Timeout, on page 24

Logging Into a 7000 or 8000 Series Device with CAC Credentials

Users are restricted to a single active session on a 7000 & 8000 Series device.

Caution Do **not** remove a CAC during an active browsing session. If you remove or replace a CAC during a session, your web browser terminates the session and the system logs you out of the web interface.

Before you begin

- If you do not have access to the web interface, contact your system administrator to modify your account privileges, or log in as a user with Administrator access and modify the privileges for the account.
- Create user accounts as described in Creating a User Account, on page 62.
- Configure CAC authentication and authorization as described in Configuring CAC Authentication, on page 74.

Procedure

Step 1 Insert a CAC as instructed by your organization.

Step 2 Direct your browser to `https://hostname/`, where *hostname* corresponds to the host name of the appliance you want to access.

Step 3 If prompted, enter the PIN associated with the CAC you inserted in step 1.

Step 4 If prompted, choose the appropriate certificate from the drop-down list.

Step 5 Click **Continue**.

Related Topics

CAC Authentication, on page 74
Session Timeout, on page 24

Logging Into the CLI on 7000/8000 Series, ASA FirePOWER, and NGIPSv Devices

With a minimum of basic CLI configuration access, you can log directly into Classic managed devices.

Before you begin

- Complete the initial setup process using the default **admin** user for the initial login.
- Create additional user accounts that can log into the CLI using the **configure user add** command.
- For the 7000 & 8000 Series devices, create user accounts at the web interface as described in Creating a User Account, on page 62.

Procedure

Step 1 SSH to the device's management interface (hostname or IP address) or use the console.

With the exception of ASA 5585-X devices, which have dedicated ASA FirePOWER console port, ASA FirePOWER devices accessed via the console default to the operating system CLI. This requires an extra step to access the Firepower CLI: **`session sfr`**.

If your organization uses SecurID® tokens when logging in, append the token to your SecurID PIN and use that as your password to log in. For example, if your PIN is `1111` and the SecurID token is `222222`, enter `1111222222`. You must have already generated your SecurID PIN before you can log in.

Step 2 At the CLI prompt, use any of the commands allowed by your level of command line access.

Logging Into the Command Line Interface on Firepower Threat Defense Devices

You can log directly into the command line interface on Firepower Threat Defense managed devices.

Before you begin

Complete the initial setup process using the default **admin** user for the initial login. Create additional user accounts that can log into the CLI using the **configure user add** command.

Procedure

Step 1 Connect to the Firepower Threat Defense CLI, either from the console port or using SSH.

You can SSH to the management interface of the Firepower Threat Defense device. You can also connect to the address on a data interface if you open the interface for SSH connections. SSH access to data interfaces is disabled by default. See Configure Secure Shell, on page 954 to allow SSH connections to specific data interfaces.

You can directly connect to the Console port on the device. Use the console cable included with the device to connect your PC to the console using a terminal emulator set for 9600 baud, 8 data bits, no parity, 1 stop bit, no flow control. See the hardware guide for your device for more information about the console cable.

The initial CLI you access on the Console port differs by device type.

- ASA Series devices—The CLI on the Console port is the regular Firepower Threat Defense CLI.
- Firepower Series devices—The CLI on the Console port is FXOS. You can get to the Firepower Threat Defense CLI using the **connect ftd** command. Use the FXOS CLI for chassis-level configuration and troubleshooting only. Use the Firepower Threat Defense CLI for basic configuration, monitoring, and normal system troubleshooting. See the FXOS documentation for information on FXOS commands.

Step 2 Log in with the **admin** username and password.

Step 3 At the CLI prompt (>), use any of the commands allowed by your level of command line access.

Step 4 (Optional) Access the diagnostic CLI:

system support diagnostic-cli

Use this CLI for advanced troubleshooting. This CLI includes additional **show** and other commands, including the **session wlan console** command needed to enter the CLI for the wireless access point on an ASA 5506W-X.

This CLI has two sub-modes: user EXEC and privileged EXEC mode. More commands are available in privileged EXEC mode. To enter privileged EXEC mode, enter the **enable** command; press enter without entering a password when prompted.

Example:

```
> system support diagnostic-cli
firepower> enable
Password:
firepower#
```

To return to the regular CLI, type **Ctrl-a**, **d**.

Logging Out of a Firepower System Web Interface

When you are no longer actively using a Firepower System web interface, Cisco recommends that you log out, even if you are only stepping away from your web browser for a short period of time. Logging out ends your web session and ensures that no one can use the interface with your credentials.



Note If you are logging out of an SSO session at the FMC, when you log out the system redirects your browser to the SSO IdP for your organization. To ensure FMC security and prevent others from accessing the FMC using your SSO account, we recommend you log out of the SSO federation at the IdP.

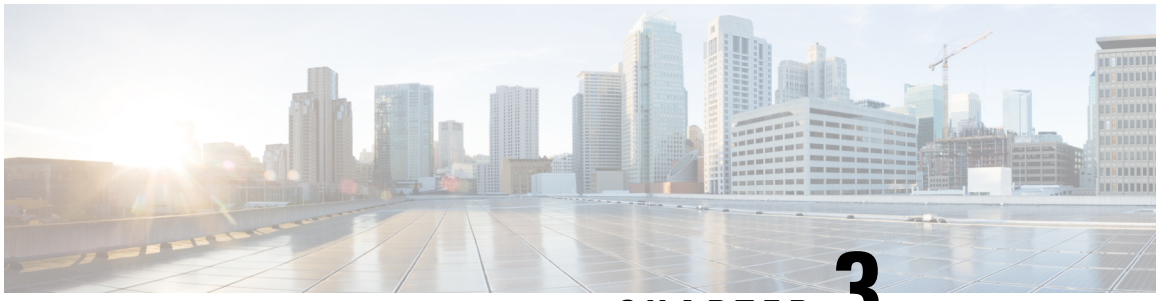
Procedure

Step 1 From the drop-down list under your user name, choose **Logout**.

Step 2 If you are logging out of an SSO session at the FMC, the system redirects you to the SSO IdP for your organization. Log out at the IdP to ensure FMC security.

Related Topics

Session Timeout, on page 24



CHAPTER 3

Specifying User Preferences

The following topics describe how to specify user preferences:

- User Preferences Introduction, on page 31
- Changing Your Password, on page 31
- Changing an Expired Password, on page 32
- Specifying Your Home Page, on page 32
- Configuring Event View Settings, on page 33
- Setting Your Default Time Zone, on page 37
- Specifying Your Default Dashboard, on page 37

User Preferences Introduction

Depending on your user role, you can specify certain preferences for your user account.

In a multidomain deployment, user preferences apply to all domains where your account has access. When specifying home page and dashboard preferences, keep in mind that certain pages and dashboard widgets are constrained by domain.

Changing Your Password

All user accounts are protected with a password. You can change your password at any time, and depending on the settings for your user account, you may have to change your password periodically.

If password strength checking is enabled, passwords must be at least eight alphanumeric characters of mixed case and must include at least one numeric character. Passwords cannot be a word that appears in a dictionary or include consecutive repeating characters.

If you are an LDAP or a RADIUS user, you cannot change your password through the web interface.

Procedure

Step 1 From the drop-down list under your user name, choose **User Preferences**.

Step 2 Enter your **Current Password**, and click **Change**.

Step 3 In the **New Password** and **Confirm** fields, enter your new password.

Step 4 Click **Change**.

Changing an Expired Password

Depending on the settings for your user account, your password may expire. The password expiration time period is set when your account is created. If your password has expired, the Password Expiration Warning page appears.

Procedure

On the Password Expiration Warning page, you have two choices:

- Click **Change Password** to change your password now. If you have zero warning days left, you **must** change your password.

  Tip If password strength checking is enabled, passwords must be at least eight alphanumeric characters of mixed case and must include at least one numeric character. Passwords cannot be a word that appears in a dictionary or include consecutive repeating characters.

- Click **Skip** to change your password later.

Specifying Your Home Page

You can specify the page within the web interface to use as your home page for the appliance. The default home page is the default dashboard (**Overview** > **Dashboards**), except for user accounts with no dashboard access, such as External Database users. (See Specifying Your Default Dashboard, on page 37 to set the default dashboard.)

In a multidomain deployment, the home page you choose applies to all domains where your user account has access. When choosing a home page for an account that frequently accesses multiple domains, keep in mind that certain pages are constrained to the Global domain.

Procedure

Step 1 From the drop-down list under your user name, choose **User Preferences**.

Step 2 Click **Home Page**.

Step 3 Choose the page you want to use as your home page from the drop-down list.

The options in the drop-down list are based on the access privileges for your user account. For more information, see User Account Privileges, on page 44.

Step 4 Click **Save**.

Configuring Event View Settings

Use the Event View Settings page to configure characteristics of event views on the Firepower Management Center. Note that some event view configurations are available only for specific user roles. Users with the External Database User role can view parts of the event view settings user interface, but changing those settings has no meaningful result.

Procedure

Step 1 From the drop-down list under your user name, choose **User Preferences**.

Step 2 Click **Event View Settings**.

Step 3 In the **Event Preferences** section, configure the basic characteristics of event views; see Event View Preferences, on page 33.

Step 4 In the **File Preferences** section, configure file download preferences; see File Download Preferences, on page 34.

Step 5 In the **Default Time Windows** section, configure the default time window or windows; see Default Time Windows, on page 35.

Step 6 In the **Default Workflow** sections, configure default workflows; see Default Workflows, on page 37.

Step 7 Click **Save**.

Event View Preferences

Use the Event Preferences section of the Event View Settings page to configure basic characteristics of event views in the Firepower System. This section is available for all user roles, although it has little to no significance for users who cannot view events.

The following fields appear in the Event Preferences section:

- The **Confirm "All" Actions** field controls whether the appliance forces you to confirm actions that affect all events in an event view.

  For example, if this setting is enabled and you click **Delete All** on an event view, you must confirm that you want to delete all the events that meet the current constraints (including events not displayed on the current page) before the appliance will delete them from the database.

- The **Resolve IP Addresses** field allows the appliance, whenever possible, to display host names instead of IP addresses in event views.

  Note that an event view may be slow to display if it contains a large number of IP addresses and you have enabled this option. Note also that for this setting to take effect, you must use management interfaces configuration to establish a DNS server in the system settings.

- The **Expand Packet View** field allows you to configure how the packet view for intrusion events appears. By default, the appliance displays a collapsed version of the packet view:
  - **None** - collapse all subsections of the Packet Information section of the packet view
  - **Packet Text** - expand only the Packet Text subsection

- **Packet Bytes** - expand only the Packet Bytes subsection
- **All** - expand all sections

Regardless of the default setting, you can always manually expand the sections in the packet view to view detailed information about a captured packet.

- The **Rows Per Page** field controls how many rows of events per page you want to appear in drill-down pages and table views.
- The **Refresh Interval** field sets the refresh interval for event views in minutes. Entering `0` disables the refresh option. Note that this interval does not apply to dashboards.
- The **Statistics Refresh Interval** controls the refresh interval for event summary pages such as the Intrusion Event Statistics and Discovery Statistics pages. Entering `0` disables the refresh option. Note that this interval does not apply to dashboards.
- The **Deactivate Rules** field controls which links appear on the packet view of intrusion events generated by standard text rules:
  - **All Policies** - a single link that deactivates the standard text rule in all the locally defined custom intrusion policies
  - **Current Policy** - a single link that deactivates the standard text rule in only the currently deployed intrusion policy. Note that you cannot deactivate rules in the default policies.
  - **Ask** - links for each of these options

To see these links on the packet view, your user account must have either Administrator or Intrusion Admin access.

Related Topics

Management Interfaces, on page 881

File Download Preferences

Use the File Preferences section of the Event View Settings page to configure basic characteristics of local file downloads. This section is only available to users with the Administrator, Security Analyst, or Security Analyst (Read Only) user roles.

Note that if your appliance does not support downloading captured files, these options are disabled.

The following fields appear in the File Preferences section:

- The **Confirm 'Download File' Actions** check box controls whether a File Download pop-up window appears each time you download a file, displaying a warning and prompting you to continue or cancel.



Caution Cisco strongly recommends you do **not** download malware, as it can cause adverse consequences. Exercise caution when downloading any file, as it may contain malware. Ensure you have taken any necessary precautions to secure the download destination before downloading files.

Note that you can disable this option any time you download a file.

- When you download a captured file, the system creates a password-protected .zip archive containing the file. The **Zip File Password** field defines the password you want to use to restrict access to the .zip file. If you leave this field blank, the system creates archive files without passwords.
- The **Show Zip File Password** check box toggles displaying plain text or obfuscated characters in the **Zip File Password** field. When this field is cleared, the **Zip File Password** displays obfuscated characters.

Default Time Windows

The time window, sometimes called the time range, imposes a time constraint on the events in any event view. Use the Default Time Windows section of the Event View Settings page to control the default behavior of the time window.

User role access to this section is as follows:

- Administrators and Maintenance Users can access the full section.
- Security Analysts and Security Analysts (Read Only) can access all options except **Audit Log Time Window**.
- Access Admins, Discovery Admins, External Database Users, Intrusion Admins, Network Admins, and Security Approvers can access only the **Events Time Window** option.

Note that, regardless of the default time window setting, you can always manually change the time window for individual event views during your event analysis. Also, keep in mind that time window settings are valid for only the current session. When you log out and then log back in, time windows are reset to the defaults you configured on this page.

There are three types of events for which you can set the default time window:

- The **Events Time Window** sets a single default time window for most events that can be constrained by time.
- The **Audit Log Time Window** sets the default time window for the audit log.
- The **Health Monitoring Time Window** sets the default time window for health events.

You can only set time windows for event types your user account can access. All user types can set event time windows. Administrators, Maintenance Users, and Security Analysts can set health monitoring time windows. Administrators and Maintenance Users can set audit log time windows.

Note that because not all event views can be constrained by time, time window settings have no effect on event views that display hosts, host attributes, applications, clients, vulnerabilities, user identity, or compliance white list violations.

You can either use **Multiple** time windows, one for each of these types of events, or you can use a **Single** time window that applies to all events. If you use a single time window, the settings for the three types of time window disappear and a new **Global Time Window** setting appears.

There are three types of time window:

- *static*, which displays all the events generated from a specific start time to a specific end time
- *expanding*, which displays all the events generated from a specific start time to the present; as time moves forward, the time window expands and new events are added to the event view

- *sliding*, which displays all the events generated from a specific start time (for example, one day ago) to the present; as time moves forward, the time window "slides" so that you see only the events for the range you configured (in this example, for the last day)

The maximum time range for all time windows is from midnight on January 1, 1970 (UTC) to 3:14:07 AM on January 19, 2038 (UTC).

The following options appear in the **Time Window Settings** drop-down list:

- The **Show the Last - Sliding** option allows you configure a sliding default time window of the length you specify.

  The appliance displays all the events generated from a specific start time (for example, 1 hour ago) to the present. As you change event views, the time window "slides" so that you always see events from the last hour.

- The **Show the Last - Static/Expanding** option allows you to configure either a static or expanding default time window of the length you specify.

  For **static** time windows, enable the **Use End Time** check box. The appliance displays all the events generated from a specific start time (for example, 1 hour ago) to the time when you first viewed the events. As you change event views, the time window stays fixed so that you see only the events that occurred during the static time window.

  For **expanding** time windows, disable the **Use End Time** check box. The appliance displays all the events generated from a specific start time (for example, 1 hour ago) to the present. As you change event views, the time window expands to the present time.

- The **Current Day - Static/Expanding** option allows you to configure either a static or expanding default time window for the current day. The current day begins at midnight, based on the time zone setting for your current session.

  For **static** time windows, enable the **Use End Time** check box. The appliance displays all the events generated from midnight to the time when you first viewed the events. As you change event views, the time window stays fixed so that you see only the events that occurred during the static time window.

  For **expanding** time windows, disable the **Use End Time** check box. The appliance displays all the events generated from midnight to the present. As you change event views, the time window expands to the present time. Note that if your analysis continues for over 24 hours before you log out, this time window can be more than 24 hours.

- The **Current Week - Static/Expanding** option allows you to configure either a static or expanding default time window for the current week. The current week begins at midnight on the previous Sunday, based on the time zone setting for your current session.

  For **static** time windows, enable the **Use End Time** check box. The appliance displays all the events generated from midnight to the time when you first viewed the events. As you change event views, the time window stays fixed so that you see only the events that occurred during the static time window.

  For **expanding** time windows, disable the **Use End Time** check box. The appliance displays all the events generated from midnight Sunday to the present. As you change event views, the time window expands to the present time. Note that if your analysis continues for over 1 week before you log out, this time window can be more than 1 week.

Default Workflows

A workflow is a series of pages displaying data that analysts use to evaluate events. For each event type, the appliance ships with at least one predefined workflow. For example, as a Security Analyst, depending on the type of analysis you are performing, you can choose among ten different intrusion event workflows, each of which presents intrusion event data in a different way.

The appliance is configured with a default workflow for each event type. For example, the Events by Priority and Classification workflow is the default for intrusion events. This means whenever you view intrusion events (including reviewed intrusion events), the appliance displays the Events by Priority and Classification workflow.

You can, however, change the default workflow for each event type. The default workflows you are able to configure depend on your user role. For example, intrusion event analysts cannot set default discovery event workflows.

Setting Your Default Time Zone

This setting determines the times displayed in the web interface for your user account only, for things like task scheduling and viewing dashboards. This setting does not change the system time or affect any other user, and does not affect data stored in the system, which generally uses UTC.

Warning The Time Zone function (in User Preferences) assumes that the system clock is set to UTC time. DO NOT ATTEMPT TO CHANGE THE SYSTEM TIME. Changing the system time from UTC is NOT supported, and doing so will require you to reimage the device to recover from an unsupported state.

Procedure

Step 1 From the drop-down list under your user name, choose **User Preferences**.

Step 2 Click **Time Zone**.

Step 3 Choose the continent or area that contains the time zone you want to use.

Step 4 Choose the country and state name that corresponds with the time zone you want to use.

Specifying Your Default Dashboard

The default dashboard appears when you choose **Overview** > **Dashboards**. Unless changed, the default dashboard for all users is the Summary dashboard. You can change the default dashboard if your user role is Administrator, Maintenance, or Security Analyst.

In a multidomain deployment, the default dashboard you choose applies to all domains where your user account has access. When choosing a dashboard for an account that frequently accesses multiple domains, keep in mind that certain dashboard widgets are constrained by domain.

Procedure

Step 1 From the drop-down list under your user name, choose **User Preferences**.

Step 2 Click **Dashboard Settings**.

Step 3 Choose the dashboard you want to use as your default from the drop-down list. If you choose **None**, when you select **Overview** > **Dashboards**, you can then choose a dashboard to view.

Step 4 Click **Save**.

Related Topics

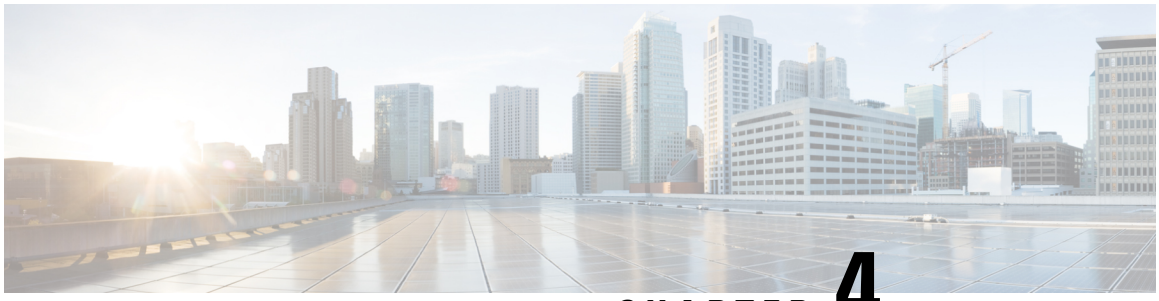
Viewing Dashboards, on page 274



PART II

Firepower System Management

- Firepower System User Management, on page 41
- Licensing the Firepower System, on page 105
- System Updates, on page 137
- Backup and Restore, on page 155
- Configuration Import and Export, on page 173
- Task Scheduling, on page 179
- Data Storage, on page 197
- Firepower Management Center High Availability, on page 201
- Device Management Basics, on page 219



CHAPTER 4

Firepower System User Management

The following topics describe how a user with Administrator access can manage user accounts in the Firepower System:

- User Roles, on page 41
- User Accounts, on page 61
- Firepower System User Authentication, on page 69
- LDAP Authentication, on page 72
- RADIUS Authentication, on page 93
- Single Sign-on (SSO), on page 102

User Roles

The Firepower System lets you allocate user privileges based on the user's role. For example, you can grant analysts predefined roles such as Security Analyst and Discovery Admin and reserve the Administrator role for the security administrator managing the Firepower System. You can also create custom user roles with access privileges tailored to your organization's needs.

In the platform settings policy for a managed device, you set a default access role for all users from that device who are externally authenticated. After an externally authenticated user logs in for the first time, you can add or remove access rights for that user on the User Management page. If you do not modify the user's rights, the user has only the rights granted by default. Because you create internally authenticated users manually, you set the access rights when you create them.

If you configured management of access rights through LDAP groups, the access rights for users are based on their membership in LDAP groups. They receive the default access rights for the group that they belong to that has the highest level of access. If they do not belong to any groups and you have configured group access, they receive the default user access rights configured in the authentication object for the LDAP server. If you configure group access, those settings override the default access setting in the platform settings policy.

Similarly, if you assign a user to specific user role lists in a RADIUS authentication object, the user receives all assigned roles, unless one or more of those roles are mutually incompatible. If a user is on the lists for two mutually incompatible roles, the user receives the role that has the highest level of access. If the user does not belong to any lists and you have configured a default access role in the authentication object, the user receives that role. If you configure default access in the authentication object, those settings override the default access setting in the platform settings policy.

In a multidomain deployment, you can assign users roles in multiple domains. For example, you can assign a user read-only privileges in the Global domain, but Administrator privileges in a subdomain.

Predefined User Roles

The Firepower System includes ten predefined user roles that provide a range of access privilege sets to meet the needs of your organization. Note that 7000 and 8000 Series devices have access to only three of the ten predefined user roles: Administrator, Maintenance User, and Security Analyst.

Although you cannot edit predefined user roles, you can use their access privilege sets as the basis for custom user roles. In addition, you cannot configure them to escalate to another user role.

The following table briefly describes the predefined roles available to you.

Access Admin

Provides access to access control policy and associated features in the **Policies** menu. Access Admins cannot deploy policies.

Administrator

Administrators have access to all functionality; their sessions present a higher security risk if compromised, so you cannot make them exempt from login session timeouts.

You should limit use of the Administrator role for security reasons.

Discovery Admin

Provides access to network discovery, application detection, and correlation features in the **Policies** menu. Discovery Admins cannot deploy policies.

External Database User

Provides read-only access to the Firepower System database using an application that supports JDBC SSL connections. For the third-party application to authenticate to the Firepower System appliance, you must enable database access in the system settings. On the web interface, External Database Users have access only to online help-related options in the **Help** menu. Because this role's function does not involve the web interface, access is provided only for ease of support and password changes.

Intrusion Admin

Provides access to all intrusion policy, intrusion rule, and network analysis policy features in the **Policies** and **Objects** menus. Intrusion Admins cannot deploy policies.

Maintenance User

Provides access to monitoring and maintenance features. Maintenance Users have access to maintenance-related options in the **Health** and **System** menus.

Network Admin

Provides access to access control, SSL inspection, DNS policy, and identity policy features in the **Policies** menu, as well as device configuration features in the **Devices** menus. Network Admins can deploy configuration changes to devices.

Security Analyst

Provides access to security event analysis features, and read-only access to health events, in the **Overview**, **Analysis**, **Health**, and **System** menus.

Security Analyst (Read Only)

Provides read-only access to security event analysis features and health event features in the **Overview**, **Analysis**, **Health**, and **System** menus.

Security Approver

Provides limited access to access control and associated policies and network discovery policies in the **Policies** menu. Security Approvers can view and deploy these policies, but cannot make policy changes.

Externally authenticated users, if assigned no other roles, have minimum access rights based on the settings in LDAP or RADIUS authentication objects and in platform settings. You can assign additional rights to these users, but to remove or change minimum access rights, you must perform the following tasks:

- Move the user from one list to another in the authentication object or change the user's attribute value or group membership on the external authentication server.
- Update platform settings.
- Use the User Management page to remove the access from that user account.

Related Topics

User Account Privileges, on page 44

Custom User Roles

In addition to the predefined user roles, you can also create custom user roles with specialized access privileges. Custom user roles can have any set of menu-based and system permissions, and may be completely original or based on a predefined user role. Like predefined user roles, custom roles can serve as the default role for externally authenticated users. Unlike predefined roles, you can modify and delete custom roles.

Selectable permissions are hierarchical, and are based on the Firepower System menu layout. Permissions are expandable if they have sub-pages or if they have more fine-grained permissions available beyond simple page access. In that case, the parent permission grants page view access and the children granular access to related features of that page. Permissions that contain the word “Manage” grant the ability to edit and delete information that other users create.



Tip For pages or features not included in the menu structure, privileges are granted by parent or related pages. For example, the Modify Intrusion Policy privilege also allows you to modify network analysis policies.

You can apply restricted searches to a custom user role. These constrain the data a user may see in the event viewer. You can configure a restricted search by first creating a private saved search and selecting it from the **Restricted Search** drop-down menu under the appropriate menu-based permission.

When you configure a custom user role on a Firepower Management Center, all menu-based permissions are available for you to grant. When you configure a custom user role on a managed device, only some permissions are available — those relevant to device functions.

The selectable options under System Permissions allow you to create a user role that can make queries to the external database or escalate to the permissions of a target user role.

Optionally, instead of creating a new custom user role, you can export a custom user role from another appliance, then import it onto your appliance. You can then edit the imported role to suit your needs before you apply it.

Related Topics

User Account Privileges, on page 44
External Database Access Settings, on page 877

Example: Custom User Roles and Access Control

You can create custom user roles for access control-related features to designate whether Firepower System users can view and modify access control and associated policies.

The following table lists custom roles that you could create and user permissions granted for each example. The table lists the privileges required for each custom role. In this example, Policy Approvers can view (but not modify) access control and intrusion policies. They can also deploy configuration changes to devices.

Table 1: Example Access Control Custom Roles

Custom Role Permission	Example: Access Control Editor	Example: Intrusion & Network Analysis Editor	Example: Policy Approver
Access Control	yes	no	yes
Access Control Policy	yes	no	yes
Modify Access Control Policy	yes	no	no
Intrusion Policy	no	yes	yes
Modify Intrusion Policy	no	yes	no
Deploy Configuration to Devices	no	no	yes

User Account Privileges

The following sections provide a list of the configurable user permissions in the Firepower System and the predefined user roles that can access them. Not all permissions are available on managed devices; permissions available only on the Firepower Management Center are marked accordingly.

Overview Menu

The following table lists, in order, the user role privileges required to access each option in the Overview menu and whether the user role has access to the sub-permissions within. The Security Approver, Discovery Admin, Intrusion Admin, Access Admin, Network Admin, and External Database User roles have no permissions in the Overview menu.

Table 2: Overview Menu

Permission	Admin	Maint User	Security Analyst	Security Analyst (RO)
Dashboards	yes	yes	yes	yes
Manage Dashboards	yes	no	no	no
Appliance Information Widget	yes	yes	yes	yes
Appliance Status Widget (*FMC only*)	yes	yes	yes	yes

Permission	Admin	Maint User	Security Analyst	Security Analyst (RO)
Correlation Events Widget	yes	no	yes	yes
Current Interface Status Widget	yes	yes	yes	yes
Current Sessions Widget	yes	no	no	no
Custom Analysis Widget (*FMC only*)	yes	no	yes	yes
Disk Usage Widget	yes	yes	yes	yes
Interface Traffic Widget	yes	yes	yes	yes
Intrusion Events Widget (*FMConly*)	yes	no	yes	yes
Network Correlation Widget (*FMC only*)	yes	no	yes	yes
Product Licensing Widget (*FMC only*)	yes	yes	no	no
Product Updates Widget	yes	yes	no	no
RSS Feed Widget	yes	yes	yes	yes
System Load Widget	yes	yes	yes	yes
System Time Widget	yes	yes	yes	yes
White List Events Widget (*FMC only*)	yes	no	yes	yes
Reporting (*FMC only*)	yes	no	yes	yes
Manage Report Templates (*FMC only*)	yes	no	yes	yes
Summary	yes	no	yes	yes
Intrusion Event Statistics (*FMC only*)	yes	no	yes	yes
Intrusion Event Performance	yes	no	no	no
Intrusion Event Graphs (*FMC only*)	yes	no	yes	yes
Discovery Statistics (*FMC only*)	yes	no	yes	yes
Discovery Performance (*FMConly*)	yes	no	no	no
Connection Summary (*FMC only*)	yes	no	yes	yes

Analysis Menu

The following table lists, in order, the user role privileges required to access each option in the Analysis menu and whether the user role has access to the sub-permissions within. Permissions that appear multiple times under different headings will be listed on the table only where they first appear, except to indicate submenu headings. The Security Approver, Intrusion Admin, Access Admin, Network Admin, and External Database

User roles have no permissions in the Analysis menu. The Analysis menu is only available on the Firepower Management Center.

Table 3: Analysis Menu

Menu	Admin	Discovery Admin	Maint User	Security Analyst	Security Analyst (RO)
Context Explorer	yes	no	no	yes	yes
Connection Events	yes	no	no	yes	yes
Modify Connection Events	yes	no	no	yes	no
Connection Summary Events	yes	no	no	yes	yes
Modify Connection Summary Events	yes	no	no	yes	no
Security Intelligence Events	yes	no	no	yes	yes
Modify Security Intelligence Events	yes	no	no	yes	no
Intrusion	yes	no	no	yes	yes
Intrusion Events	yes	no	no	yes	yes
Modify Intrusion Events	yes	no	no	yes	no
View Local Rules	yes	no	no	yes	yes
Reviewed Events	yes	no	no	yes	yes
Clipboard	yes	no	no	yes	yes
Incidents	yes	no	no	yes	yes
Modify Incidents	yes	no	no	yes	no
Files	yes	no	no	yes	yes
Malware Events	yes	no	no	yes	yes
Modify Malware Events	yes	no	no	yes	no
File Events	yes	no	no	yes	yes
Modify File Events	yes	no	no	yes	no
Captured Files	yes	no	no	yes	yes
Modify Captured Files	yes	no	no	yes	no
File Trajectory	yes	no	no	yes	yes
File Download	yes	no	no	yes	yes
Dynamic File Analysis	yes	no	no	yes	no

Menu	Admin	Discovery Admin	Maint User	Security Analyst	Security Analyst (RO)
Hosts	yes	no	no	yes	yes
Network Map	yes	no	no	yes	yes
Hosts	yes	no	no	yes	yes
Modify Hosts	yes	no	no	yes	no
Indications of Compromise	yes	no	no	yes	yes
Modify Indications of Compromise	yes	no	no	yes	no
Servers	yes	no	no	yes	yes
Modify Servers	yes	no	no	yes	no
Vulnerabilities	yes	no	no	yes	yes
Modify Vulnerabilities	yes	no	no	yes	no
Host Attributes	yes	no	no	yes	yes
Modify Host Attributes	yes	no	no	yes	no
Applications	yes	no	no	yes	yes
Application Details	yes	no	no	yes	yes
Modify Application Details	yes	no	no	yes	no
Host Attribute Management	yes	no	no	no	no
Discovery Events	yes	no	no	yes	yes
Modify Discovery Events	yes	no	no	yes	no
Users	yes	yes	no	yes	yes
User Activity	yes	yes	no	yes	yes
Modify User Activity Events	yes	yes	no	yes	no
Users	yes	yes	no	yes	yes
Modify Users	yes	yes	no	yes	no
Indications of Compromise	yes	no	no	yes	yes
Modify Indications of Compromise	yes	no	no	yes	no
Vulnerabilities	yes	no	no	yes	yes
Third-party Vulnerabilities	yes	no	no	yes	yes

Menu	Admin	Discovery Admin	Maint User	Security Analyst	Security Analyst (RO)
Modify Third-party Vulnerabilities	yes	no	no	yes	no
Correlation	yes	yes	no	yes	yes
Correlation Events	yes	yes	no	yes	yes
Modify Correlation Events	yes	yes	no	yes	no
White List Events	yes	yes	no	yes	yes
Modify White List Events	yes	yes	no	yes	no
White List Violations	yes	yes	no	yes	yes
Remediation Status	yes	yes	no	no	no
Modify Remediation Status	yes	yes	no	no	no
Custom	yes	no	no	yes	yes
Custom Workflows	yes	no	no	yes	yes
Manage Custom Workflows	yes	no	no	yes	yes
Custom Tables	yes	no	no	yes	yes
Manage Custom Tables	yes	no	no	yes	yes
Search	yes	no	yes	yes	yes
Manage Search	yes	no	no	no	no
Bookmarks	yes	no	no	yes	yes
Manage Bookmarks	yes	no	no	yes	yes
Application Statistics	yes	no	no	yes	yes
Geolocation Statistics	yes	no	no	yes	yes
User Statistics	yes	no	no	yes	yes
URL Category Statistics	yes	no	no	yes	yes
URL Reputation Statistics	yes	no	no	yes	yes
DNS Queries by Record Types	yes	no	no	yes	yes
SSL Statistics	yes	no	no	yes	yes
Intrusion Event Statistics by Application	yes	no	no	yes	yes
Intrusion Event Statistics by User	yes	no	no	yes	yes

Menu	Admin	Discovery Admin	Maint User	Security Analyst	Security Analyst (RO)
Security Intelligence Category Statistics	yes	no	no	yes	yes
File Storage Statistics by Disposition	yes	no	no	yes	yes
File Storage Statistics by Type	yes	no	no	yes	yes
Dynamic File Analysis Statistics	yes	no	no	yes	yes

Policies Menu

The following table lists, in order, the user role privileges required to access each option in the Policies menu and whether the user roles has access to the sub-permissions within. The External Database User, Maintenance User, Security Analyst, and Security Analyst (Read Only) roles have no permissions in the Policies menu. The Policies menu is only available on the Firepower Management Center.

Note that the Intrusion Policy and Modify Intrusion Policy privileges also allow you to create and modify network analysis policies.

Table 4: Policies Menu

Menu	Access Admin	Admin	Discovery Admin	Intrusion Admin	Network Admin	Security Approver
Access Control	yes	yes	no	no	yes	yes
Access Control Policy	yes	yes	no	no	yes	yes
Modify Access Control Policy	yes	yes	no	no	yes	no
Modify Administrator Rules	yes	yes	no	no	yes	no
Modify Root Rules	yes	yes	no	no	yes	no
Intrusion Policy	no	yes	no	yes	no	yes
Modify Intrusion Policy	no	yes	no	yes	no	no
Malware & File Policy	yes	yes	no	no	no	yes
Modify Malware & File Policy	yes	yes	no	no	no	no
DNS Policy	yes	yes	no	no	yes	yes
Modify DNS Policy	yes	yes	no	no	yes	no
Identity Policy	yes	yes	no	no	yes	no
Modify Identity Policy	yes	yes	no	no	yes	no
Modify Administrator Rules	yes	yes	no	no	yes	no
Modify Root Rules	yes	yes	no	no	yes	no

Menu	Access Admin	Admin	Discovery Admin	Intrusion Admin	Network Admin	Security Approver
SSL Policy	yes	yes	no	no	yes	yes
Modify SSL Policy	yes	yes	no	no	yes	no
Modify Administrator Rules	yes	yes	no	no	yes	no
Modify Root Rules	yes	yes	no	no	yes	no
Prefilter Policy	yes	yes	no	no	yes	yes
Modify Prefilter Policy	yes	yes	no	no	yes	no
Network Discovery	no	yes	yes	no	no	yes
Custom Fingerprinting	no	yes	yes	no	no	no
Modify Custom Fingerprinting	no	yes	yes	no	no	no
Custom Topology	no	yes	yes	no	no	no
Modify Custom Topology	no	yes	no	no	no	no
Modify Network Discovery	no	yes	yes	no	no	no
Application Detectors	no	yes	yes	no	no	no
Modify Application Detectors	no	yes	yes	no	no	no
User 3rd Party Mappings	no	yes	yes	no	no	no
Modify User 3rd Party Mappings	no	yes	no	no	no	no
Custom Product Mappings	no	yes	yes	no	no	no
Modify Custom Product Mappings	no	yes	no	no	no	no
Correlation	no	yes	no	no	no	no
Policy Management	no	yes	no	no	no	no
Modify Policy Management	no	yes	yes	no	no	no
Rule Management	no	yes	no	no	no	no
Modify Rule Management	no	yes	yes	no	no	no
White List	no	yes	no	no	no	no
Modify White List	no	yes	yes	no	no	no
Traffic Profiles	no	yes	no	no	no	no
Modify Traffic Profiles	no	yes	yes	no	no	no

Menu	Access Admin	Admin	Discovery Admin	Intrusion Admin	Network Admin	Security Approver
Actions	no	yes	yes	no	no	yes
Alerts	no	yes	yes	no	no	yes
Impact Flag Alerts	no	yes	yes	no	no	no
Modify Impact Flag Alerts	no	yes	yes	no	no	no
Discovery Event Alerts	no	yes	yes	no	no	no
Modify Discovery Event Alerts	no	yes	yes	no	no	no
Email	no	yes	no	yes	no	no
Modify Email	no	yes	no	yes	no	no
Modify Alerts	no	yes	yes	no	no	no
Scanners	no	yes	yes	no	no	no
Scan Results	no	yes	yes	no	no	no
Modify Scan Results	no	yes	yes	no	no	no
Modify Scanners	no	yes	yes	no	no	no
Groups	no	yes	no	no	no	no
Modify Groups	no	yes	yes	no	no	no
Modules	no	yes	no	no	no	no
Modify Modules	no	yes	yes	no	no	no
Instances	no	yes	no	no	no	no
Modify Instances	no	yes	yes	no	no	no

Devices Menu

The **Devices** menu table lists, in order, the user role privileges required to access each option in the Devices menu and the sub-permissions within. The Discovery Admin, External Database User, Intrusion Admin, Maintenance User, Security Analyst, and Security Analyst (Read Only) have no permissions in the Devices menu. The Devices menu is only available on the Firepower Management Center.

Table 5: Devices Menu

Menu	Access Admin	Admin	Network Admin	Security Approver
Device Management	no	yes	yes	yes

Menu	Access Admin	Admin	Network Admin	Security Approver
Modify Devices	no	yes	yes	no
NAT	yes	yes	yes	yes
NAT List	yes	yes	yes	yes
Modify NAT Policy	yes	yes	yes	no
VPN	no	yes	yes	yes
Modify VPN	no	yes	yes	no
Certificates	no	yes	yes	yes
Modify Certificates	no	yes	yes	no
QoS	yes	yes	yes	no
Modify QoS Policy	yes	yes	yes	no
FlexConfig Policy	no	yes	no	no
Modify FlexConfig Policy	no	yes	no	no
Device Management	no	yes	yes	no
Modify Devices	no	yes	yes	no

Object Manager Menu

The Object Manager menu table lists, in order, the user role privileges required to access each option in the Object Manager menu and the sub-permission within. The Discovery Admin, Security Approver, Maintenance User, External Database User, Security Analyst, and Security Analyst (Read Only) have no permissions in the Object Manager menu. The Object Manager menu is available only on theFirepower Management Center.

Table 6: Object Manager Menu

Menu	Access Admin	Admin	Intrusion Admin	Network Admin
Object Manager	yes	yes	no	yes
Rule Editor	no	yes	yes	no
Modify Rule Editor	no	yes	yes	no
NAT List	yes	yes	no	yes
Modify Object Manager	no	yes	no	no

Cisco AMP

The Cisco AMP permission is available only to the Administrator user role. This permission is available only on the Firepower Management Center.

Deploy Configuration to Devices

The Deploy Configuration to Devices permission is available to the Administrator, Network Admin, and Security Approver roles. This permission is available only on the Firepower Management Center.

System Menu

The following table lists, in order, the user role privileges required to access each option in the System menu and whether the user role has access to the sub-permissions within. The External Database User role has no permissions in the System Menu.

Table 7: System Menu

Menu	Access Admin	Admin	Discovery Admin	Intrusion Admin	Maint User	Network Admin	Security Approver	Se Ar
Configuration	no	yes	no	no	no	no	no	no
Domains	no	yes	no	no	no	no	no	no
Integration	no	yes	no	no	no	yes	yes	no
Cisco CSI	yes	yes	no	no	no	yes	yes	no
Identity Realms (*FMC only*)	yes	yes	no	no	no	yes	yes	no
Modify Identity Realms (*FMC only*)	yes	yes	no	no	no	yes	no	no
Identity Sources (*FMC only*)	yes	yes	no	no	no	yes	yes	no
Modify Identity Sources (*FMC only*)	yes	yes	no	no	no	yes	no	no
eStreamer	no	yes	no	no	no	no	no	no
Host Input Client (*FMC only*)	no	yes	no	no	no	no	no	no
Smart Software Satellite (*FMC only*)	yes	yes	no	no	no	yes	yes	no
Modify Smart Software Satellite (*FMC only*)	yes	yes	no	no	no	yes	no	no
User Management	no	yes	no	no	no	no	no	no
Users	no	yes	no	no	no	no	no	no
User Roles	no	yes	no	no	no	no	no	no
External Authentication (*FMC only*)	no	yes	yes	no	no	no	no	no
Updates	no	yes	no	no	no	no	no	no

Menu	Access Admin	Admin	Discovery Admin	Intrusion Admin	Maint User	Network Admin	Security Approver	Secu Analy
Rule Updates (*FMC only*)	no	yes	no	yes	no	no	no	no
Rule Update Import Log (*FMC only*)	no	yes	no	no	no	no	no	no
Licenses	no	yes	no	no	no	no	no	no
Smart Licences	no	yes	no	no	no	no	no	no
Modify Smart Licenses	no	yes	no	no	no	no	no	no
Classic Licenses	no	yes	no	no	no	no	no	no
Health (*FMC only*)	no	yes	no	no	yes	no	no	yes
Health Policy (*FMC only*)	no	yes	no	no	yes	no	no	yes
Modify Health Policy (*FMC only*)	no	yes	no	no	yes	no	no	yes
Apply Health Policy (*FMC only*)	no	yes	no	no	yes	no	no	yes
Health Events (*FMC only*)	no	yes	no	no	yes	no	no	yes
Modify Health Events (*FMC only*)	no	yes	no	no	yes	no	no	yes
Monitoring	no	yes	no	no	yes	yes	yes	yes
Audit	no	yes	no	no	yes	no	no	no
Modify Audit Log	no	yes	no	no	yes	no	no	no
Syslog	no	yes	no	no	yes	no	no	no
Statistics	no	yes	no	no	yes	no	no	no
Tools	no	yes	no	no	yes	no	no	yes
Backup Management	no	yes	no	no	yes	no	no	no
Restore Backup	no	yes	no	no	yes	no	no	no
Scheduling	no	yes	no	no	yes	no	no	no
Delete Other Users' Scheduled Tasks	no	yes	no	no	no	no	no	no
Import/Export	no	yes	no	no	no	no	no	no
Discovery Data Purge (*FMC only*)	no	yes	no	no	no	no	no	yes
Whois (*FMC only*)	no	yes	no	no	yes	no	no	yes

REST VDI Menu

The **REST VDI** menu table lists, in order, the user role privileges required to access each option in the REST VDI menu and the sub-permissions within. REST VDI permissions are required to use the TS Agent for user awareness and user control. For more information about the TS Agent, see:

- The Terminal Services (TS) Agent Identity Source, on page 1901
- *Cisco Terminal Services Agent (TS Agent) Guide*

Table 8: REST VDI Menu



Help Menu

The Help menu and its permissions are accessible to all user roles. You cannot restrict Help menu options.

Managing User Roles

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

Each Firepower System user is associated with a user access role or roles. These user roles are assigned permissions that determine access to menus and other options in the system. For example, an analyst needs access to event data to analyze the security of your network, but might not require access to administrative functions for the Firepower System itself. You can grant Security Analyst access to analysts while reserving the Administrator role for the user or users managing the Firepower System.

The Firepower System includes ten predefined user roles designed for a variety of administrators and analysts. These predefined user roles have a set of predetermined access privileges.

You can also create custom user roles with more granular access privileges.

You can also restrict the data that a user role can view in the event viewer by applying a restricted search to that role. To create a custom role with restricted access, you must choose the tables you want to restrict from the Menu Based Permissions list, then choose private saved searches from the Restrictive Search drop-down lists.

You cannot delete predefined user roles, but you can delete custom roles that are no longer necessary. If you want to disable a custom role without removing it entirely, you can deactivate it instead. Note that you cannot delete your own user role or a role that is set as a default user role in a platform settings policy.

Procedure

Step 1 Choose **System** > **Users**.

Step 2 Click the **User Roles** tab.

Step 3 Manage user roles:

- Activate — Activate or deactivate a predefined user role as described in Activating and Deactivating User Roles, on page 56.
- Create — Create custom user roles as described in Creating Custom User Roles, on page 57
- Copy — Copy an existing user role to create a new custom user role as described in Copying User Roles, on page 57.
- Edit — Edit a custom user role as described in Editing Custom User Roles, on page 58.
- Delete — Click **Delete** () next to the custom role you want to delete. If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- **Note** If a deleted role is the only role assigned to a given user, that user can log in and access the User Preferences menu, but is otherwise unable to access the Firepower System.

Activating and Deactivating User Roles

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

You cannot delete predefined user roles, but you can deactivate them. Deactivating a role removes that role and all associated permissions from any user who is assigned that role.

In a multidomain deployment, the system displays custom user roles created in the current domain, which you can edit. It also displays custom user roles created in ancestor domains, which you cannot edit. To view and edit custom user roles in a lower domain, switch to that domain.



Caution If a deactivated role is the only role assigned to a given user, that user can log in and access the User Preferences menu, but is otherwise unable to access the Firepower System.

Procedure

Step 1 Choose **System** > **Users**.

Step 2 Click the **User Roles** tab.

Step 3 Click the slider next to the user role you want to activate or deactivate.

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

If you deactivate, then reactivate, a role with Lights-Out Management while a user with that role is logged in, or restore a user or user role from a backup during that user's login session, that user must log back into the web interface to regain access to IPMItool commands.

Creating Custom User Roles

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin



Caution Users with menu-based User Management permissions have the ability to elevate their own privileges or create new user accounts with extensive privileges, including the Administrator user role. For system security reasons we strongly recommend you restrict the list of users with User Management permissions appropriately.

Procedure

Step 1 Choose **System** > **Users**.

Step 2 Click the **User Roles** tab.

Step 3 Click **Create User Role**.

Step 4 In the **Name** field, enter a name for the new user role. User role names are case sensitive.

Step 5 Optionally, add a **Description**.

Step 6 Choose menu-based permissions for the new role.

When you choose a permission, all of its children are chosen, and the multi-value permissions use the first value. If you clear a high-level permission, all of its children are cleared also. If you choose a permission but not its children, it appears in italic text.

Copying a predefined user role to use as the base for your custom role preselects the permissions associated with that predefined role.

Step 7 Optionally, set database access permissions for the new role by checking or unchecking the **External Database Access** checkbox.

Step 8 Optionally, on Firepower Management Centers, set escalation permissions for the new user role as described in Configuring a Custom User Role for Escalation, on page 60.

Step 9 Click **Save**.

Copying User Roles

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

You can copy an existing role to use as the basis for a new custom role. This preselects the existing role's permissions in the User Role Editor so you can model one role on another.

You can copy any existing role, including predefined user roles and custom user roles inherited from ancestor domains.

Procedure

Step 1 Choose **System** > **Users**.

Step 2 Click the **User Roles** tab.

Step 3 Click **Copy** () next to the user role you want to copy.

Step 4 Enter a new **Name**.

The system creates a default name for the new user role by combining the name of the original user role and the `(copy)` suffix.

Step 5 Enter a new **Description**.

The system retains the description of the original user role if you do not overwrite it.

Step 6 Optionally, modify the menu-based permissions inherited from the original user role.

When you choose a permission, all of its children are chosen, and the multi-value permissions use the first value. If you clear a high-level permission, all of its children are cleared also. If you choose a permission but not its children, the permission appears in italic text.

Step 7 Optionally, set the database access permissions for the new role by checking or unchecking the **External Database Access** checkbox.

Step 8 Optionally, set escalation permissions for the new user role as described in Configuring a Custom User Role for Escalation, on page 60.

Step 9 Click **Save**.

Editing Custom User Roles

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

You cannot edit predefined user roles.

In a multidomain deployment, the system displays custom user roles created in the current domain, which you can edit. It also displays custom user roles created in ancestor domains, which you cannot edit. To view and edit custom user roles in a lower domain, switch to that domain.

Procedure

Step 1 Choose **System** > **Users**.

Step 2 Click the **User Roles** tab.

Step 3 Click **Edit** ( ) next to the custom user role you want to modify.If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Modify the **Name** and **Description** fields. User role names are case sensitive.

Step 5 Choose menu-based permissions for the user role.

When you choose a permission, all of its children are chosen, and the multi-value permissions use the first value. If you clear a high-level permission, all of its children are cleared also. If you choose a permission but not its children, the permission appears in italic text.

Step 6 Optionally, set the database access permissions for the role by checking or unchecking the **External Database Access** checkbox.

Step 7 Optionally, on Firepower Management Centers, set escalation permissions for the user role as described in Configuring a Custom User Role for Escalation, on page 60.

Step 8 Click **Save**.

User Role Escalation

You can give custom user roles the permission, with a password, to temporarily gain the privileges of another, targeted user role in addition to those of the base role. This allows you to easily substitute one user for another during an absence, or to more closely track the use of advanced user privileges.

For example, a user whose base role has very limited privileges may escalate to the Administrator role to perform administrative actions. You can configure this feature so that users can use their own passwords, or so they use the password of another user that you specify. The second option allows you to easily manage one escalation password for all applicable users.

Note that only one user role at a time can be the escalation target role. You can use a custom or predefined user role. Each escalation lasts for the duration of a login session and is recorded in the audit log.

Setting the Escalation Target Role

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

You can assign any of your user roles, predefined or custom, to act as the system-wide escalation target role. This is the role to which any other role may escalate, if it has the ability.

Procedure

Step 1 Choose **System** > **Users**.

Step 2 Click **User Roles**.

Step 3 Click **Configure Permission Escalation**.

Step 4 Choose a user role from the drop-down list.

Step 5 Click **OK** to save your changes.

Note Changing the escalation target role is effective immediately. Users in escalated sessions now have the permissions of the new escalation target.

Configuring a Custom User Role for Escalation

Smart License	Classic License	Supported Device	Supported Domains	Access
Any	Any	Any	Any	Admin

Consider the needs of your organization when you configure the escalation password for a custom role. If you want to easily manage many escalating users, you may want to choose another user whose password serves as the escalation password. If you change that user's password or deactivate that user, all escalating users who require that password are affected. This allows you to manage user role escalation more efficiently, especially if you choose an externally authenticated user that you can manage centrally.

Procedure

Step 1 Begin configuring your custom user role as described in Creating Custom User Roles, on page 57.

Step 2 In System Permissions, choose the **Set this role to escalate to:** check box.

The current escalation target role is listed beside the check box.

Step 3 Choose the password that this role uses to escalate. You have two options:

- If you want users with this role to use their own passwords when they escalate, choose **Authenticate with the assigned user's password**.
- If you want users with this role to use the password of another user, choose **Authenticate with the specified user's password** and enter that username.

Note When authenticating with another user's password, you can enter any username, even that of a deactivated or nonexistent user. Deactivating the user whose password is used for escalation makes escalation impossible for users with the role that requires it. You can use this feature to quickly remove escalation powers if necessary.

Step 4 Click **Save**.
Users with this role can now escalate to the target user role.

Escalating Your User Role

Smart License	Classic License	Supported Device	Supported Domains	Access
Any	Any	FMC	Any	Any

When a user has an assigned custom user role with permission to escalate, that user may escalate to the target role's permissions at any time. Note that escalation has no effect on user preferences.

Before you begin

- Confirm that a system administrator configured the escalation target role or custom user role for escalation as described in Setting the Escalation Target Role, on page 59 or Configuring a Custom User Role for Escalation, on page 60.

Procedure

Step 1 From the drop-down list under your user name, choose **Escalate Permissions**.

Step 2 Enter the authentication password.

Step 3 Click **Escalate**. You now have all permissions of the escalation target role in addition to your current role.

Note Escalation lasts for the remainder of your login session. To return to the privileges of your base role only, you must log out, then begin a new session.

User Accounts

The admin account and optional, custom user accounts on a Firepower Management Center or Firepower 7000 and 8000 Series device allow users to log into these. For internally-authenticated users, accounts must be created manually. For externally-authenticated users, accounts are created automatically.

For Firepower Threat Defense, you can create separate CLI users. These users can access the device through SSH to do additional troubleshooting and system monitoring. However, you must create these users in the CLI, you cannot create them in Firepower Management Center.

Related Topics

Firepower System User Accounts
Firepower System User Interfaces

Managing User Accounts

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

Procedure

Step 1 Choose **System** > **Users**.

Step 2 Manage user accounts:

- Activate/Deactivate — Click the slider next to a user to reactivate a deactivated user, or to disable an active user account without deleting it. Only internally authenticated users can be activated and deactivated.
- Create — Create a new user account; see Creating a User Account, on page 62.
- Edit — Edit an existing user account; see Editing a User Account, on page 63.

- Delete — If you want to delete a user, click **Delete** (). You can delete user accounts from the system at any time, with the exception of the admin account, which cannot be deleted.

Related Topics

Lights-Out Management User Access Configuration, on page 919
Predefined User Roles, on page 42
Custom User Roles, on page 43

Creating a User Account

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	FMC 7000 & 8000 Series	Any	Admin

When you set up a new user account, you can control which parts of the system the account can access. You can set password expiration and strength settings for the user account during creation. For a local account on a 7000 or 8000 Series device, you can also configure the level of command line access the user will have.

In a multidomain deployment, you can create user accounts in any domain in which you have been assigned Admin access. You can also create accounts in a higher-level domain and assign the users lower-level access only. For example, you might want a single user to be an administrator of two domains, but deny them access to the ancestor domain. This kind of user account can only be modified by switching to a subdomain in which access is assigned.

Procedure

Step 1 Choose **System** > **Users**.

Step 2 Click **Create User**.

Step 3 Enter a **User Name**.

Step 4 Modify the login options; see User Account Login Options, on page 64.

Step 5 Enter values in **Password** and **Confirm Password**.

The values you construct must be based on the password options you set earlier.

Step 6 If you are creating a user account on a 7000 or 8000 Series device, assign the appropriate level of **Command-Line Interface Access** as described in Command Line Access Levels, on page 66.

Step 7 Assign user roles:

- Check or uncheck the check box next to the user role(s) you want to assign the user.
- In a multidomain deployment, if you are adding a user account to a domain with descendant domains, click the **Add Domains** button that displays instead of the user role check boxes. Continue as described in Assigning User Roles in Multiple Domains, on page 63.

Note User roles determine the user's access rights. For more information, see Managing User Roles, on page 55.

Step 8 Click **Save**.

Editing a User Account

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

After adding user accounts to the system, you can modify access privileges, account options, or passwords at any time. Note that password management options do not apply to users who authenticate to an external directory server. You manage those settings on the external server. However, you must configure access rights for all accounts, including those that are externally authenticated.

Note For externally authenticated users, you cannot remove the minimum access rights through the Firepower System user management page for users assigned an access role because of LDAP group or RADIUS list membership or attribute values. You can, however, assign additional rights. When you modify the access rights for an externally authenticated user, the Authentication Method column on the User Management page provides a status of **External - Locally Modified**.

If you change the authentication for a user from externally authenticated to internally authenticated, you must supply a new password for the user.

Procedure

Step 1 Choose **System** > **Users**.

Step 2 Click **Edit** ( ) next to the user you want to modify.

Step 3 Modify settings described in Creating a User Account, on page 62.

Step 4 Click **Save**.

Assigning User Roles in Multiple Domains

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

In a multidomain deployment, you can assign users roles in ancestor and descendant domains. For example, you can assign a user read-only privileges in the Global domain, but Admin privileges in a descendant domain.

Procedure

Step 1 In the user account editor, click **Add Domain**.

Step 2 Choose a domain from the **Domain** drop-down list.

Step 3 Check the user roles you want to assign the user.

Step 4 Click **Save**.

Converting a User from Internal to External Authentication

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

Note When you convert a user from internal to external authentication, the user account retains the permissions already present in that account. The existing permissions override any permissions associated with the associated authentication object group or the default user role set in the platform settings policy.

Before you begin

- A user record with the same user name must be present on the external authentication server.

Procedure

Step 1 Enable LDAP (with or without CAC) or RADIUS authentication. For more information, see LDAP Authentication, on page 72 or RADIUS Authentication, on page 93.

Step 2 Instruct the user to log in with the password stored for that user on the external server.

User Account Login Options

The following table describes some of the options you can use to regulate passwords and account access for Firepower System users.

Note

- Password management options do not apply to users who authenticate to an external directory server. You manage those settings on the external authentication server. After you enable **Use External Authentication Method**, the system removes password management options from the display.
- If you enable security certifications compliance or Lights-Out Management (LOM) on an appliance, different password restrictions apply. For more information on security certifications compliance, see Security Certifications Compliance, on page 909 .

Table 9: User Account Login Options

Option	Description
Use External Authentication Method	Select this check box if you want this user's credentials to be externally authenticated. If you enable this option, the password management options are no longer displayed. **Note** • For users to authenticate to an external directory server, you must also create an authentication object for the server you want to use, and deploy a platform settings policy with authentication enabled. • Note that for externally authenticated users, if the authentication object for the server is disabled, the **Authentication Method** column in the Users list displays **External (Disabled)**. • If you select this option for the user and the external authentication server is unavailable, that user can log into the web interface but cannot access any functionality.
Maximum Number of Failed Logins	Enter an integer, without spaces, that determines the maximum number of times each user can try to log in after a failed login attempt before the account is locked. The default setting is five tries; use `0` to allow an unlimited number of failed logins.
Minimum Password Length	Enter an integer, without spaces, that determines the minimum required length, in characters, of a user's password. The default setting is `8`. A value of `0` indicates that no minimum length is required. If you enable the **Check Password Strength** option, and set a value for **Minimum Password Length** that exceeds 8 characters, the higher value applies.
Days Until Password Expiration	Enter the number of days after which the user's password expires. The default setting is `0`, which indicates that the password never expires. If you set this option, the **Password Lifetime** column of the Users list indicates the days remaining on each user's password.
Days Before Password Expiration Warning	Enter the number of warning days users have to change their password before their password actually expires. The default setting is `0` days. **Note** The number of warning days must be **less than** the number of days before the password expires.
Force Password Reset on Login	Select this option to force users to change their passwords the next time they log in.
Check Password Strength	Select this option to require strong passwords. A strong password must be at least eight alphanumeric characters of mixed case and must include at least one numeric character and one special character. It cannot be a word that appears in a dictionary or include consecutive repeating characters.
Exempt from Browser Session Timeout	Select this option if you do not want a user's login sessions to terminate due to inactivity. Users with the Administrator role cannot be made exempt.

Command Line Access Levels

You can use the local web interface on a 7000 or 8000 Series device to assign command line interface access to local device users. Note that you can also assign command line access for users on an NGIPSv, but you use commands from the command line interface.

The commands a user can run depend on the level of access you assign to the user. Possible values for the **Command-Line Interface Access** setting include:

None

The user cannot log into the appliance on the command line. Any session the user starts will close when the user provides credentials. The access level defaults to **None** on user creation.

Configuration

The user can access any of the command line options. Exercise caution in assigning this level of access to users.



Caution Command line access granted to externally authenticated users defaults to the **Configuration** level of command line access, granting rights to all command line utilities.

Basic

A specific set of commands can be run by the user, listed below.

Table 10: Basic Command Line Commands

`configure password`	`interfaces`
`end`	`lcd`
`exit`	`link-state`
`help`	`log-ips-connection`
`history`	`managers`
`logout`	`memory`
`?`	`model`
`??`	`mpls-depth`
`access-control-config`	`NAT`
`alarms`	`network`
`arp-tables`	`network-modules`
`audit-log`	`ntp`
`bypass`	`perfstats`
`high-availability`	`portstats`

cpu	power-supply-status
database	process-tree
device-settings	processes
disk	routing-table
disk-manager	serial-number
dns	stacking
expert	summary
fan-status	time
fastpath-rules	traffic-statistics
gui	version
hostname	virtual-routers
hyperthreading	virtual-switches
inline-sets	

Creating Local User Accounts for the FTD CLI

You can create users for CLI access on Firepower Threat Defense devices. These accounts do not allow access to the management application, but to the CLI only. The CLI is useful for troubleshooting and monitoring purposes.

You cannot create local user accounts on more than one device at a time. Each device has its own set of unique local user CLI accounts.

Procedure

Step 1 Log into the device CLI using an account with config privileges.

The admin user account has the required privileges, but any account with config privileges will work. You can use an SSH session or the Console port.

For certain device models, the Console port puts you into the FXOS CLI. Use the **connect ftd** command to get to the Firepower Threat Defense CLI.

Step 2 Create the user account.

configure user add *username* {**basic** | **config**}

You can define the user with the following privilege levels:

- **config**—Gives the user configuration access. This gives the user full administrator rights to all commands.
- **basic**—Gives the user basic access. This does not allow the user to enter configuration commands.

Example:

The following example adds a user account named joecool with config access rights. The password is not shown as you type it.

```
> configure user add joecool config
Enter new password for user joecool: newpassword
Confirm new password for user joecool: newpassword
> show user
Login              UID   Auth Access  Enabled Reset    Exp Warn  Str Lock Max
admin             1000  Local Config  Enabled    No  Never  N/A  Dis   No N/A
joecool           1001  Local Config  Enabled    No  Never  N/A  Dis   No   5
```

Note Tell users they can change their passwords using the **configure password** command.

Step 3 (Optional.) Adjust the characteristics of the account to meet your security requirements.

You can use the following commands to change the default account behavior.

- **configure user aging** *username max_days warn_days*

  Sets an expiration date for the user's password. Specify the maximum number of days for the password to be valid followed by the number of days before expiration the user will be warned about the upcoming expiration. Both values are 1 to 9999, but the warning days must be less than the maximum days. When you create the account, there is no expiration date for the password.

- **configure user forcereset** *username*

  Forces the user to change the password on the next login.

- **configure user maxfailedlogins** *username number*

  Sets the maximum number of consecutive failed logins you will allow before locking the account, from 1 to 9999. Use the **configure user unlock** command to unlock accounts. The default for new accounts is 5 consecutive failed logins.

- **configure user minpasswdlen** *username number*

  Sets a minimum password length, which can be from 1 to 127.

- **configure user strengthcheck** *username* {**enable** | **disable**}

  Enables or disables password strength checking, which requires a user to meet specific password criteria when changing their password. When a user's password expires or if the **configure user forcereset** command is used, this requirement is automatically enabled the next time the user logs in.

Step 4 Manage user accounts as necessary.

Users can get locked out of their accounts, or you might need to remove accounts or fix other issues. Use the following commands to manage the user accounts on the system.

- **configure user access** *username* {**basic** | **config**}

  Changes the privileges for a user account.

- **configure user delete** *username*

  Deletes the specified account.

- **configure user disable** *username*

Disables the specified account without deleting it. The user cannot log in until you enable the account.

- **configure user enable** *username*

  Enables the specified account.

- **configure user password** *username*

  Changes the password for the specified user. Users should normally change their own password using the **configure password** command.

- **configure user unlock** *username*

  Unlocks a user account that was locked due to exceeding the maximum number of consecutive failed login attempts.

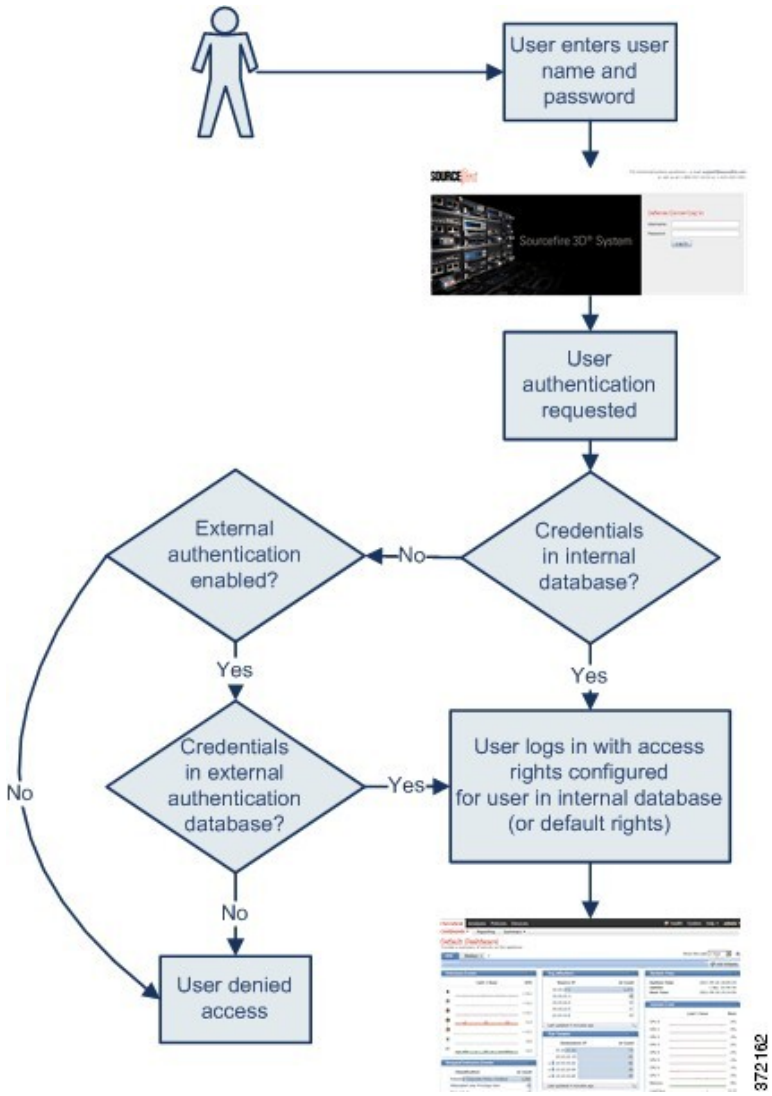
Firepower System User Authentication

When a user logs into the web interface on a Firepower Management Center or a managed device, the appliance looks for a match for the user name and password in the local list of users. This process is called *authentication*.

There are two types of authentication:

- *internal authentication* — The system checks the list in the local database for the user.
- *external authentication* — The system checks the list in the local database for the user and, if the user is not present on that list, queries an external authentication server for its user list.

The authentication process is illustrated below.



When you create a user account, you specify either internal or external authentication for that user.

Internal Authentication

In internal authentication, user credentials are verified against records in the internal Firepower System database. This is the default authentication type.

You set the access rights for internal authentication users when you create the user's account.

Note When an internally authenticated user is converted to external authentication, you cannot revert to internal authentication.

External Authentication

In external authentication, the Firepower Management Center or managed device retrieves user credentials from a repository on an external server. External servers can be either a Lightweight Directory Access Protocol (LDAP) directory server or a Remote Authentication Dial In User Service (RADIUS) authentication server.

You enable external authentication using a platform settings policy and settings in individual user accounts. Note the following guidelines:

- You can use multiple external authentication objects to authenticate users to access the Firepower Management Center web interface. In other words, if you have five external authentication objects, users from any of them can be authenticated to access the web interface.
- You can use only one external authentication object for shell access to the Firepower Management Center. If you have more than one external authentication object set up, users can authenticate using only the first object in the list.

When the user logs into an appliance for the first time, the appliance associates the external credentials with a set of permissions by creating a local user record. The user is assigned permissions based on either:

- the group or access list they belong to
- the default user access role you set in the platform settings policy for the appliance

If permissions are granted through group or list membership, they cannot be modified. However, if they are assigned by default user role, you can modify them in the user account, and the modifications you make override the default settings. For example:

- If the default role for externally authenticated user accounts is set to a specific access role, users can log into the appliance using their external account credentials without any additional configuration by the system administrator.
- If an account is externally authenticated and by default receives no access privileges, users can log in but cannot access any functionality. You (or your system administrator) can then change the permissions to grant the appropriate access to user functionality.

You cannot manage passwords for externally authenticated users or deactivate externally authenticated users through the Firepower System interface. For externally authenticated users, you cannot remove the minimum access rights through the Firepower System user management page for users assigned an access role because of LDAP group or RADIUS list membership or attribute values. On the Edit User page for an externally authenticated user, rights granted because of settings on an external authentication server are marked with a status of **Externally Modified**.

You can, however, assign additional rights. When you modify the access rights for an externally authenticated user, the Authentication Method column on the User Management page provides a status of **External - Locally Modified**.

Related Topics

LDAP Authentication, on page 72
RADIUS Authentication, on page 93

LDAP Authentication

LDAP, or the Lightweight Directory Access Protocol, allows you to set up a directory on your network that organizes objects, such as user credentials, in a centralized location. Multiple applications can then access those credentials and the information used to describe them. If you ever need to change a user's credentials, you can change them in one place.

You must create LDAP authentication objects on a Firepower Management Center, but you can use the external authentication object on any managed devices that have a web interface (that is, on 7000 and 8000 Series devices) by deploying a platform settings policy where the object is enabled to the device. When you deploy the policy, the object is copied to the device.

Note Before enabling external authentication on 7000 and 8000 Series devices, remove any internally-authenticated shell or CLI users that have the same user name as externally-authenticated users included in your shell access filter.

You can use LDAP naming standards for address specification and for filter and attribute syntax in your authentication object. For more information, see the RFCs listed in the Lightweight Directory Access Protocol (v3): Technical Specification, RFC 3377. Examples of syntax are provided throughout this procedure. Note that when you set up an authentication object to connect to a Microsoft Active Directory Server, you can use the address specification syntax documented in the Internet RFC 822 (Standard for the Format of ARPA Internet Text Messages) specification when referencing a user name that contains a domain. For example, to refer to a user object, you might type `JoeSmith@security.example.com` rather than the equivalent user distinguished name of `cn=JoeSmith,ou=security, dc=example,dc=com` when using Microsoft Active Directory Server.

Note Currently, the Firepower System supports LDAP external authentication on LDAP servers running Microsoft Active Directory on Windows Server 2008, Oracle Directory Server Enterprise Edition 7.0 on Windows Server 2008, or OpenLDAP on Linux. However, the Firepower System does not support external authentication for NGIPSv or ASA FirePOWER devices.

Microsoft has announced that Active Directory servers will start enforcing LDAP binding and LDAP signing in 2020. Microsoft is making these a requirement because when using default settings, an elevation of privilege vulnerability exists in Microsoft Windows that could allow a man-in-the-middle attacker to successfully forward an authentication request to a Windows LDAP server. For more information, see 2020 LDAP channel binding and LDAP signing requirement for Windows on the Microsoft support site.

If you have not done so already, we recommend you start using TLS/SSL encryption to authenticate with an Active Directory server.

Required Information for Creating LDAP Authentication Objects

Before you configure a connection to your LDAP server, you should collect the information that you need to create the LDAP authentication object.

Note You must have TCP/IP access from your local appliance to the authentication server where you want to connect.

You need the following, at minimum, to create a basic authentication object:

- the server name or IP address for the server where you plan to connect
- the server type of the server where you plan to connect
- the user name and password for a user account with sufficient privileges to browse the LDAP tree; Cisco recommends that you use a domain admin user account for this purpose
- if there is a firewall between the appliance and the LDAP server, an entry in the firewall to allow outgoing connections
- if possible, the base distinguished name for the server directory where the user names reside

Tip You can use a third-party LDAP client to browse the LDAP tree and see base DN and attribute descriptions. You can also use that client to confirm that your selected user can browse the base DN you select. Ask your LDAP administrator to recommend an approved LDAP client for your LDAP server.

Depending on how you plan to customize your advanced LDAP authentication object configuration, you might also need the information in the following table.

Table 11: Additional LDAP Configuration Information

To...	You need...
connect over a port other than 389	the port number
connect via an encrypted connection	the certificate for the connection
filter the users who can access your appliance based on an attribute value	the attribute-value pair to filter by
use an attribute as a UI access attribute rather than checking the user distinguished name	the name of the attribute
use an attribute as a shell login attribute rather than checking the user distinguished name	the name of the attribute
filter the users who can access your appliance via the shell based on an attribute value	the attribute-value pair to filter by
associate groups with specific user roles	the distinguished name of each group, as well as the group member attribute if the groups are static groups or the group member URL attribute if the groups are dynamic groups

To...	You need...
use CACs for authentication and authorization	your CAC, a server certificate signed by the same CA that issued your CAC, and the certificate chain for both certificates

CAC Authentication

If your organization uses Common Access Cards (CACs), you can configure LDAP authentication to authenticate users logging into the web interface and authorize access to specific functionality based on group membership or default access rights. With CAC authentication and authorization configured, users have the option to log in directly without providing a separate username and password for the appliance.

Note You **must** have a valid user certificate present in your browser (in this case, a certificate passed to your browser via your CAC) to enable user certificates as part of the CAC configuration process. After you configure CAC authentication and authorization, users on your network **must** maintain the CAC connection for the duration of their browsing session. If you remove or replace a CAC during a session, your web browser terminates the session and the system logs you out of the web interface.

CAC-authenticated users are identified in the system by their electronic data interchange personal identifier (EDIPI) numbers. After users log in using their CAC credentials for the first time, you can manually add or remove access privileges for those users on the User Management page. If you did not preconfigure a user's privileges using group-controlled access roles, the user has only the privileges granted by default in the platform settings policy.

Tip The system purges manually configured access privileges when it purges CAC-authenticated users from the User Management page after 24 hours of inactivity. The users are restored to the page after each subsequent login, but you must reconfigure any manual changes to their access privileges.

Configuring CAC Authentication

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	FMC 7000 and 8000 Series	Any	Admin/Network Admin

Before users on your network can log into Firepower Management Centers and 7000 and 8000 Series devices using their CAC credentials, a user with appropriate permissions must complete the multi-step configuration process for CAC authentication and authorization.

Before you begin

- Gather the information described in Required Information for Creating LDAP Authentication Objects, on page 72.

Procedure

Step 1 Insert a CAC as directed by your organization.

Step 2 Direct your browser to the web interface of the FMC or device.

Step 3 If prompted, enter the PIN associated with the CAC you inserted in step 1.

Step 4 If prompted, choose the appropriate certificate from the drop-down list.

Step 5 On the Login page, in the **Username** and **Password** fields, log in as a user with Administrator privileges. User names are case sensitive.

You cannot log in using your CAC credentials until you have fully configured CAC authentication and authorization.

Step 6 Navigate to **System** > **Users** and click the **External Authentication** tab.

Step 7 Create an LDAP authentication object exclusively for CAC authentication and authorization.

See Creating Advanced LDAP Authentication Objects, on page 78. You must configure:

- The **User Name Template** in the advanced options of the **LDAP-Specific Parameters** section.
- The **UI Access Attribute** in the **Attribute Mapping** section.
- The distinguished names for existing LDAP groups in the **Group Controlled Access Roles** section, if you want to preconfigure access rights through LDAP group membership.

You cannot configure both CAC authentication and shell access in the same authentication object. If you also want to authorize users for shell access, create and enable separate authentication objects.

Step 8 Click **Save**.

Step 9 Enable external authentication and CAC authentication.

Step 10 Select **System** > **Configuration** and click **HTTPS Certificate**.

Step 11 Import a HTTPS server certificate, if necessary.

See Importing HTTPS Server Certificates, on page 875. The same certificate authority (CA) must issue the HTTPS server certificate and the user certificates on the CACs you plan to use for authentication and authorization.

Step 12 Under **HTTPS User Certificate Settings**, choose **Enable User Certificates**.

For more information, see Requiring Valid HTTPS Client Certificates, on page 876 .

What to do next

- After the user logs in for the first time, you can manually add or remove the user's access rights. If you do not modify the rights, the user has only the rights granted by default. For more information, see Editing a User Account, on page 63.

Related Topics

LDAP Group Fields, on page 87
LDAP-Specific Fields, on page 83
Logging Into a 7000 or 8000 Series Device with CAC Credentials, on page 26
Logging Into the Firepower Management Center with CAC Credentials, on page 26

Creating Basic LDAP Authentication Objects

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

You can set up an LDAP authentication object where you customize many of the values. However, if you just want to authenticate all the users in a particular directory, you can create a basic authentication object with the base DN for that directory. If you set defaults to those for your server type and supply authentication credentials for the account used to retrieve user data from the server, you can quickly create an authentication object. Follow the procedure below to do so.



Note If you prefer to consider and possibly customize each authentication setting when creating the authentication object (to grant shell access, for example), use the advanced procedure to create the object. You should also use the advanced procedure if you plan to encrypt your connection to the server, set user timeouts, customize the user name template, or assign Firepower user roles based on LDAP group membership.

In a multidomain deployment, external authentication objects are only available in the domain in which they are created.

Before you begin

- Gather the information described in Required Information for Creating LDAP Authentication Objects, on page 72.

Procedure

Step 1 Choose **System** > **Users**.

Step 2 Click the **External Authentication** tab.

Step 3 Click **Add External Authentication Object**.

Step 4 Choose **LDAP** from the **Authentication Method** drop-down list.

Step 5 Provide a **Name**, **Description**, **Server Type**, and **Primary Server Host Name/IP Address** as described in Identifying the LDAP Authentication Server, on page 82.

Tip If you click Set Defaults, the system populates the **User Name Template**, **UI Access Attribute**, **Shell Access Attribute**, **Group Member Attribute**, and **Group Member URL Attribute** fields with default values.

Step 6 Choose **Fetch DNs** to specify a base distinguished name and, optionally, provide a **Base Filter** as described in Configuring LDAP-Specific Parameters, on page 85.

Step 7 Enter a distinguished name as the **User Name** and the **Password** for a user who has sufficient credentials to browse the LDAP server as described in Configuring LDAP-Specific Parameters, on page 85.

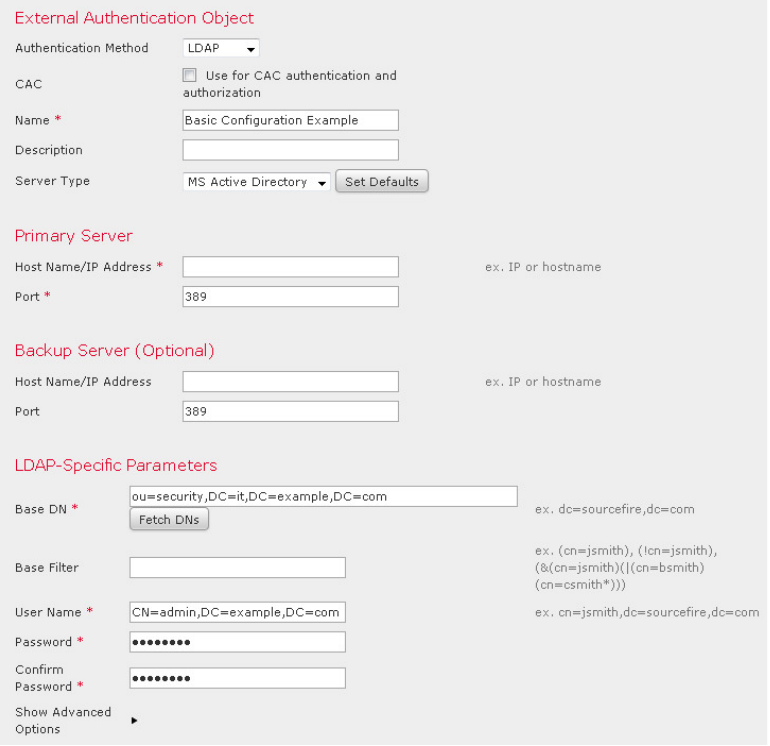
Step 8 Re-enter the password in the **Confirm Password** field.

Step 9 Test the connection as described in Testing LDAP Authentication Connections, on page 91.

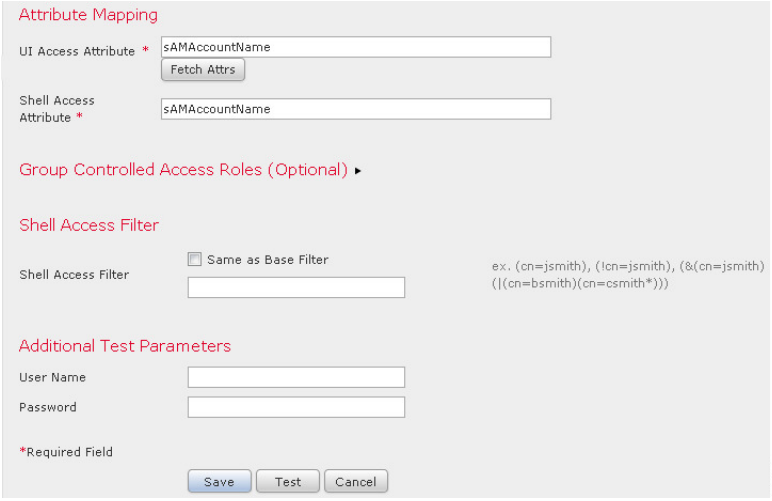
Step 10 Click **Save**.

Example

The following figures illustrate a basic configuration of an LDAP login authentication object for a Microsoft Active Directory Server. The LDAP server in this example has an IP address of 10.11.3.4. The connection uses port 389 for access.



This example shows a connection using a base distinguished name of `OU=security,DC=it,DC=example,DC=com` for the security organization in the information technology domain of the Example company.



However, because this server is a Microsoft Active Directory server, it uses the `sAMAccountName` attribute to store user names rather than the `uid` attribute. Choosing the MS Active Directory server type and clicking **Set Defaults** sets the UI Access Attribute to `sAMAccountName`. As a result, the system checks the `sAMAccountName` attribute for each object for matching user names when a user attempts to log into the system.

In addition, a Shell Access Attribute of `sAMAccountName` causes each `sAMAccountName` attribute to be checked for all objects in the directory for matches when a user logs into a shell or CLI account on the appliance.

Note that because no base filter is applied to this server, the system checks attributes for all objects in the directory indicated by the base distinguished name. Connections to the server time out after the default time period (or the timeout period set on the LDAP server).

What to do next

- If you want to refine the list of users retrieved, see Troubleshooting LDAP Authentication Connections, on page 92 for more information.

Creating Advanced LDAP Authentication Objects

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

When you create a basic authentication object, you define basic settings that let you connect to an authentication server. When you create an advanced authentication object, you define basic settings and you also choose the directory context and search criteria you want to use to retrieve user data from the server. Optionally, you can configure shell access authentication.

Although you can use the default settings for your server type to quickly set up an LDAP configuration, you can also customize advanced settings to control whether the appliance makes an encrypted connection to the LDAP server, the timeout for the connection, and which attributes the server checks for user information.

For the LDAP-specific parameters, you can use LDAP naming standards and filter and attribute syntax. For more information, see the RFCs listed in the Lightweight Directory Access Protocol (v3): Technical Specification, RFC 3377. Examples of syntax are provided throughout this procedure. Note that when you set up an authentication object to connect to a Microsoft Active Directory Server, you can use the address specification syntax documented in the Internet RFC 822 (Standard for the Format of ARPA Internet Text Messages) specification when referencing a user name that contains a domain. For example, to refer to a user object, you might enter `JoeSmith@security.example.com` rather than the equivalent user distinguished name of `cn=JoeSmith,ou=security, dc=example,dc=com` when using Microsoft Active Directory Server.

Note If you are configuring an LDAP authentication object for use with CAC authentication, do **not** remove the CAC inserted in your computer. You **must** have a CAC inserted at all times after enabling user certificates.

In a multidomain deployment, external authentication objects are only available in the domain in which they are created.

Before you begin

- Gather the information described in Required Information for Creating LDAP Authentication Objects, on page 72.
- Remove any internally authenticated shell users that have the same user name as externally authenticated users included in your shell access filter.

Procedure

Step 1 Choose **System** > **Users**.

Step 2 Click **External Authentication**, then **Add External Authentication Object**.

Step 3 Identify the authentication server as described in Identifying the LDAP Authentication Server, on page 82.

Step 4 Configure authentication settings as described in Configuring LDAP-Specific Parameters, on page 85.

Step 5 Optionally, configure LDAP groups to use as the basis for default access role assignments as described in Configuring Access Rights by Group, on page 88.

If you plan to use this object for CAC authentication and authorization, we recommend you configure LDAP groups to manage access role assignments.

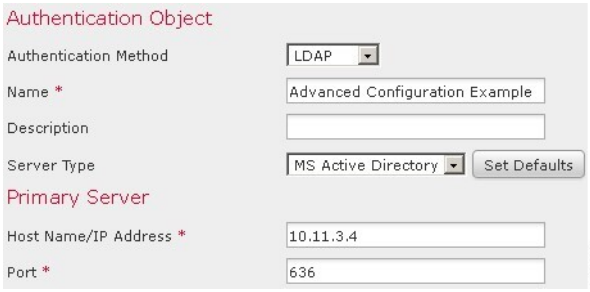
Step 6 Optionally, configure authentication settings for shell access as described in Configuring LDAP Shell Access, on page 90.

Step 7 Test your configuration as described in Testing LDAP Authentication Connections, on page 91.

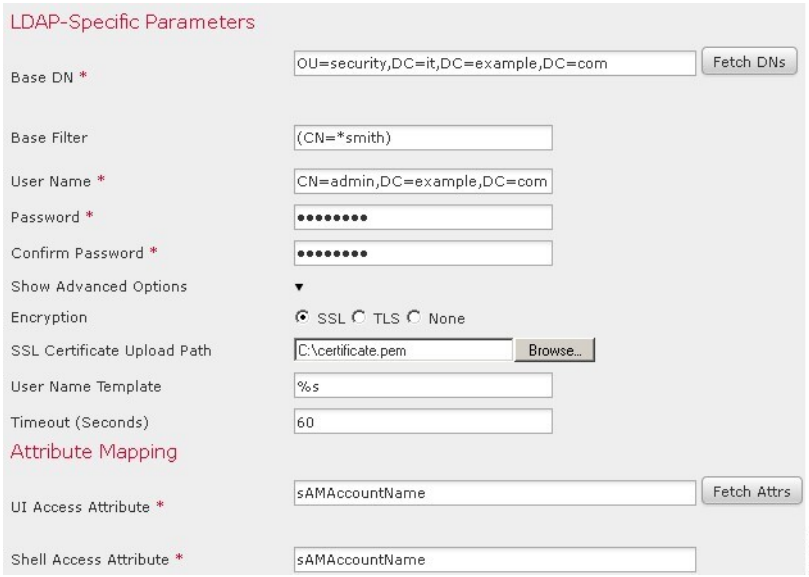
Step 8 Click **Save**.

Example

This example illustrates an advanced configuration of an LDAP login authentication object for a Microsoft Active Directory Server. The LDAP server in this example has an IP address of 10.11.3.4. The connection uses port 636 for access.



This example shows a connection using a base distinguished name of `OU=security,DC=it,DC=example,DC=com` for the security organization in the information technology domain of the Example company. However, note that this server has a base filter of `(cn=*smith)`. The filter restricts the users retrieved from the server to those with a common name ending in `smith`.

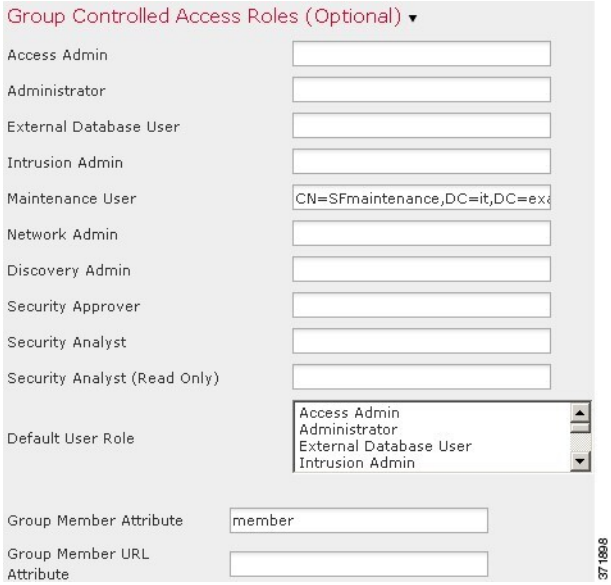


The connection to the server is encrypted using SSL and a certificate named `certificate.pem` is used for the connection. In addition, connections to the server time out after 60 seconds because of the **Timeout** setting.

Because this server is a Microsoft Active Directory server, it uses the `sAMAccountName` attribute to store user names rather than the `uid` attribute. Note that the configuration includes a UI Access Attribute of `sAMAccountName`. As a result, the Firepower System checks the `sAMAccountName` attribute for each object for matching user names when a user attempts to log into the Firepower System.

In addition, a Shell Access Attribute of `sAMAccountName` causes each `sAMAccountName` attribute to be checked for all objects in the directory for matches when a user logs into a shell account on the appliance.

This example also has group settings in place. The Maintenance User role is automatically assigned to all members of the group with a `member` group attribute and the base domain name of `CN=SFmaintenance,DC=it,DC=example,DC=com`.



The shell access filter is set to be the same as the base filter, so the same users can access the appliance through the shell or CLI as through the web interface.



LDAP Authentication Server Fields

CAC

Select this checkbox if you want to use CAC for authentication and authorization.

Name

A name for the authentication server.

Description

A description for the authentication server.

Server Type

The type of LDAP server you plan to connect to. You have the following options when selecting a type:

- If you are connecting to a Microsoft Active Directory server, select **MS Active Directory**.

- If you are connecting to a Sun Java Systems Directory Server or Oracle Directory Server, select **Oracle Directory**.
- If you are connecting to an OpenLDAP server, select **OpenLDAP**.
- If you are connecting to a LDAP server other than those listed above and want to clear default settings, select **Other**.



Tip If you click Set Defaults, the system populates the **User Name Template**, **UI Access Attribute**, **Shell Access Attribute**, **Group Member Attribute**, and **Group Member URL Attribute** fields with default values.

Primary Server Host Name/IP Address

The IP address or host name for the primary server where you want to obtain authentication data.



Note If you are using a certificate to connect via TLS or SSL, the host name in the certificate must match the host name used in this field. In addition, IPv6 addresses are not supported for encrypted connections.

Primary Server Port

The port used by the primary authentication server.

Backup Server Host Name/IP Address

The IP address or host name for the backup server where you want to obtain authentication data.

Backup Server Port

The port used by the backup authentication server.

Identifying the LDAP Authentication Server

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

When you create an authentication object, you first specify the primary and backup server and server port where you want the managed device or Firepower Management Center to connect for authentication.



Note If you are configuring an LDAP authentication object for use with CAC authentication, do **not** remove the CAC inserted in your computer. You **must** have a CAC inserted at all times after enabling user certificates.

In a multidomain deployment, external authentication objects are only available in the domain in which they are created.

Procedure

Step 1 Choose **System** > **Users**.

Step 2 Click the **External Authentication** tab.

Step 3 Click **Add External Authentication Object**.

Step 4 Choose **LDAP** from the **Authentication Method** drop-down list.

Step 5 Optionally, check the check box for **CAC** if you plan to use this authentication object for CAC authentication and authorization.

Note You must follow the procedure in Configuring CAC Authentication, on page 74 to fully configure CAC authentication and authorization.

Step 6 Enter a name and description for the authentication server in the **Name** and **Description** fields.

Step 7 Choose a **Server Type** from the drop-down list as described in LDAP Authentication Server Fields, on page 81. Optionally, click **Set Defaults**.

Step 8 Enter a **Primary Server Host Name/IP Address**.

Note If you are using a certificate to connect via TLS or SSL, the host name in the certificate must match the host name used in this field. In addition, IPv6 addresses are not supported for encrypted connections.

Step 9 Optionally, enter a **Primary Server Port**.

Step 10 Optionally, enter a **Backup Server Host Name/IP Address**.

Step 11 Optionally, enter a **Backup Server Port**.

What to do next

- Continue creating your LDAP authentication object as described in Creating Advanced LDAP Authentication Objects, on page 78.

LDAP-Specific Fields

The following table describes each of the LDAP-specific parameters.

Table 12: LDAP-Specific Parameters

Setting	Description	Examples
Base DN	Supplies the base distinguished name of the directory where the appliance searches for user information on the LDAP server. Typically, the base DN has a basic structure indicating the company domain and operational unit. Note that after you identify a primary server, you can automatically retrieve a list of available base DNs from the server and select the appropriate base DN.	The Security organization of the Example company might have a base DN of `ou=security, dc=example,dc=com`

Setting	Description	Examples
Base Filter	Focuses your search by only retrieving objects in the base DN that have the specific attribute-value pair set in the filter. The base filter is an attribute type, a comparison operator, and the attribute value you want to use as a filter enclosed in parentheses.	To filter for only users with a common name starting with F, use the filter `(cn=F*)`.
User Name/ Password	Allows the local appliance to access the user objects. Supplies user credentials for a user with appropriate rights to the authentication objects you want to retrieve. The distinguished name for the user you specify must be unique to the directory information tree for the LDAP server. Server user names associated with a Microsoft Active Directory Server cannot end with the `$` character.	The user name for the `admin` user in the Security organization of the Example company might have a user name of `cn=admin, ou=security, dc=example,dc=com`
Encryption	Determines whether and how the communications are encrypted. You can choose no encryption, Transport Layer Security (TLS), or Secure Sockets Layer (SSL) encryption. Note that if you are using a certificate to authenticate when connecting via TLS or SSL, the name of the LDAP server in the certificate **must** match the **User Name** you supply. If you change the encryption method after specifying the port, the port resets to the default value for the selected server type.	If you enter `10.10.10.250` in the external authentication settings and `computer1. example.com` in the certificate, the connection fails, even if `computer1. example.com` has an IP address of `10.10.10.250`. Changing the name of the server in the external authentication settings to `computer1. example.com` causes the connection to succeed.
SSL Certificate Upload Path	Indicates the path on your local computer to the certificate to be used for encryption.	`c:/server.crt`
User Name Template	Indicates how user names entered on login should be formatted, by mapping the string conversion character (`%s`) to the value of the **UI Access Attribute** for the user. The user name template is the format for the distinguished name used for authentication. When a user enters a user name into the login page, the appliance substitutes the name for the string conversion character and uses the resulting distinguished name to search for the user credentials. If you want to use this object for CAC authentication and authorization, you **must** enter a **User Name Template**.	`%s@security.example.com`, `%s@mail.com`, `%s@mil`, `%s@smil.mil`,
Timeout	Sets a timeout for the connection attempt to the primary server, so the connection rolls over to the backup server. If the number of seconds indicated in this field (or the timeout on the LDAP server) elapses without a response from the primary authentication server, the appliance then queries the backup server. However, if LDAP is running on the port of the primary LDAP server and for some reason refuses to service the request, the failover to the backup server does not occur.	If the primary server has LDAP disabled, the appliance queries the backup server.

Setting	Description	Examples
UI Access Attribute	Tells the local appliance to match the value of a specific attribute rather than the value of the user distinguished name. You can use any attribute, if the value of the attribute is a valid user name for the Firepower System web interface. If one of the objects has a matching user name and password, the user login request is authenticated. Selecting a server type and setting defaults prepopulates the **UI Access Attribute** with a value typically appropriate for that type of server. If you leave this field blank, the local appliance checks the user distinguished name value for each user record on the LDAP server to see if it matches the user name. If you want to use this object for CAC authentication and authorization, you **must** enter a value that corresponds with your **User Name Template** value.	`sAMAccountName`, `userPrincipalName`, `mail`

Configuring LDAP-Specific Parameters

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

The settings in the LDAP-specific parameters section determine the area of the LDAP directory where the appliance searches for user names, and control details of how the appliance connects to the LDAP server.

Valid user names are unique, and can include underscores (_), periods (.), hyphens (-), and alphanumeric characters.

In addition for most LDAP-specific settings, you can use LDAP naming standards and filter and attribute syntax. For more information, see the RFCs listed in the Lightweight Directory Access Protocol (v3): Technical Specification, RFC 3377. Examples of syntax are provided throughout this procedure. Note that when you set up an authentication object to connect to a Microsoft Active Directory Server, you can use the address specification syntax documented in the Internet RFC 822 (Standard for the Format of ARPA Internet Text Messages) specification when referencing a user name that contains a domain. For example, to refer to a user object, you might enter `JoeSmith@security.example.com` rather than the equivalent user distinguished name of `cn=JoeSmith,ou=security, dc=example,dc=com` when using Microsoft Active Directory Server.



Note If you are configuring an LDAP authentication object for use with CAC authentication, do **not** remove the CAC inserted in your computer. You **must** have a CAC inserted at all times after enabling user certificates.

Procedure

Step 1 In the **LDAP-Specific Parameters** section of the Create External Authentication Object page, you have two options for setting the base DN:

- Click **Fetch DNs**, and choose the appropriate base distinguished name from the drop-down list.

- Enter the base distinguished name for the LDAP directory you want to access in the **Base DN** field. For example, to authenticate names in the Security organization at the Example company, enter `ou=security,dc=example,dc=com`.

Step 2 Optionally, enter a **Base Filter**.

Example:

For example, if the user objects in a directory tree have a `physicalDeliveryOfficeName` attribute and users in the New York office have an attribute value of `NewYork` for that attribute, to retrieve only users in the New York office, enter `(physicalDeliveryOfficeName=NewYork)`.

Step 3 Enter a distinguished name as the **User Name** and the **Password** for a user who has sufficient credentials to browse the LDAP server.

Example:

For example, if you are connecting to an OpenLDAP server where user objects have a `uid` attribute and the object for the administrator in the Security division at our example company has a `uid` value of `NetworkAdmin`, you might enter `uid=NetworkAdmin,ou=security,dc=example,dc=com`.

Caution If you are connecting to a Microsoft Active Directory Server, you cannot provide a server user name that ends with the `$` character.

Step 4 Re-enter the password in the **Confirm Password** field.

Step 5 After you configure the basic LDAP-specific parameters, you have several options:

- To access advanced options, click the arrow next to **Show Advanced Options** and continue with the next step.
- If you want to configure user default roles based on LDAP group membership, continue with Configuring Access Rights by Group, on page 88.
- If you are not using LDAP groups for authentication, continue with Configuring LDAP Shell Access, on page 90.

Step 6 Choose an **Encryption** mode for your LDAP connection.

Note Note that if you change the encryption method after specifying a port, you reset the port to the default value for that method. For none or TLS, the port uses the default value of 389. If you choose SSL encryption, the port uses the default of 636.

Step 7 If you choose TLS or SSL encryption and you want to use a certificate to authenticate, **Browse** to the location of a valid TLS or SSL certificate.

Note If you previously uploaded a certificate and want to replace it, upload the new certificate and redeploy the configuration to your appliances to copy over the new certificate.

Step 8 Optionally, provide a **User Name Template** that corresponds with your **UI Access Attribute**.

Example:

For example, to authenticate all users who work in the Security organization of our example company by connecting to an OpenLDAP server where the UI access attribute is `uid`, you might enter `uid=%s,ou=security,dc=example,dc=com` in the **User Name Template** field. For a Microsoft Active Directory server, you could enter `%s@security.example.com`.

Note If you want to use CAC credentials for authentication and authorization, you **must** enter a value in the **User Name Template** field.

Step 9 Optionally, in the **Timeout** field, enter the number of seconds that should elapse before rolling over to the backup connection.

Step 10 Optionally, to retrieve users based on an attribute instead of the Base DN and Base Filter, you have two options:

- Click **Fetch Attrs** to retrieve a list of available attributes, and choose the appropriate attribute.
- Enter a **UI Access Attribute**. For example, on a Microsoft Active Directory Server, you may want to use the UI Access Attribute to retrieve users, because there may not be a `uid` attribute on Active Directory Server user objects. Instead, you can search the `userPrincipalName` attribute by typing `userPrincipalName` in the **UI Access Attribute** field.

Note If you want to use CAC credentials for authentication and authorization, you **must** enter a value in the **UI Access Attribute** field.

What to do next

- Continue creating your LDAP authentication object as described in Creating Advanced LDAP Authentication Objects, on page 78.

LDAP Group Fields

Any group you reference must exist on the LDAP server. You can reference static LDAP groups or dynamic LDAP groups. Static LDAP groups are groups where membership is determined by group object attributes that point to specific users, and dynamic LDAP groups are groups where membership is determined by creating an LDAP search that retrieves group users based on user object attributes. Group access rights for a role only affect users who are members of the group.

The access rights granted when a user logs into the Firepower System depend on the LDAP configuration:

- If no group access rights are configured for your LDAP server, when a new user logs in, the Firepower System authenticates the user against the LDAP server and then grants user rights based on the default minimum access role set in the platform settings policy.
- If you configure any group settings, new users belonging to specified groups inherit the minimum access setting for the groups where they are members.
- If a new user does not belong to any specified groups, the user is assigned the default minimum access role specified in the Group Controlled Access Roles section of the authentication object.
- If a user belongs to more than one configured group, the user receives the access role for the group with the highest access as a minimum access role.

You cannot use the Firepower System user management page to remove the minimum access rights for users assigned an access role because of LDAP group membership. You can, however, assign additional rights. When you modify the access rights for an externally authenticated user, the Authentication Method column on the User Management page provides a status of **External - Locally Modified**.

Note If you use a dynamic group, the LDAP query is used exactly as it is configured on the LDAP server. For this reason, the Firepower System limits the number of recursions of a search to four to prevent search syntax errors from causing infinite loops. If a user's group membership is not established in those recursions, the default access role defined in the Group Controlled Access Roles section is granted to the user.

Firepower System User Roles

The distinguished names for the LDAP groups that contain users who should be assigned each user role.

Default User Role

The default minimum access role for users that do not belong to any of the specified groups.

Group Member Attribute

The LDAP attribute that contains the LDAP search string in a static group.

Group Member URL Attribute

The LDAP attribute that designates membership in a dynamic group

Configuring Access Rights by Group

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

If you prefer to base default access rights on a user's membership in an LDAP group, you can specify distinguished names for existing groups on your LDAP server for each of the access roles used by your Firepower System. When you do so, you can configure a default access setting for those users detected by LDAP that do not belong to any specified groups. When a user logs in, the Firepower System dynamically checks the LDAP server and assigns default access rights according to the user's current group membership.

If you do not configure a user's privileges using group-controlled access roles, a user has only the privileges granted by default in the platform settings policy.

If you plan to use an object for CAC authentication and authorization, Cisco recommends configuring LDAP groups to manage access role assignments for CAC-authenticated users.

Note If you are configuring an LDAP authentication object for use with CAC authentication, do **not** remove the CAC inserted in your computer. You **must** have a CAC inserted at all times after enabling user certificates.

Before you begin

- Confirm that the group you plan to reference exists on the LDAP server.

Procedure

Step 1 On the Create External Authentication Object page, click the down arrow next to **Group Controlled Access Roles**.

Step 2 Optionally, in the **DN** fields that correspond to Firepower System user roles, enter the distinguished name for the LDAP groups that contain users who should be assigned to those roles.

Example:

For example, you might enter the following in the **Administrator** field to authenticate names in the information technology organization at the `Example` company:

```
cn=itgroup,ou=groups, dc=example,dc=com
```

Step 3 Choose a **Default User Role**.

Step 4 If you use static groups, enter a **Group Member Attribute**.

Example:

For example, if the `member` attribute is used to indicate membership in the static group you reference for default Security Analyst access, enter `member`.

Step 5 If you use dynamic groups, enter a **Group Member URL Attribute**.

Example:

For example, if the `memberURL` attribute contains the LDAP search that retrieves members for the dynamic group you specified for default Admin access, enter `memberURL`.

What to do next

- Continue creating your LDAP authentication object as described in Creating Advanced LDAP Authentication Objects, on page 78.

LDAP Shell Access Fields

With the exception of the admin account, shell access is controlled entirely though the shell access attribute you set. The shell access filter you set determines which set of users on the LDAP server can log into the shell.

Note that a home directory for each shell user is created on login, and when an LDAP shell access user account is disabled (by disabling the LDAP connection), the directory remains, but the user shell is set to `/bin/false` in `/etc/password` to disable the shell. If the user then is re-enabled, the shell is reset, using the same home directory.

Shell users can log in using user names with lowercase, uppercase, or mixed case letters. Login authentication for the shell is case sensitive.

Shell Access Attribute

The access attribute you want to use for filtering. You can use any attribute if the value of the attribute is a valid user name for shell access.

If you leave this field blank, the user distinguished name is used for shell access authentication.

Tip Selecting a server type and setting defaults prepopulates this field with an attribute typically appropriate for that type of server.

Shell Access Filter

The attribute value you want to use to retrieve administrative user entries for shell access. The filter is an attribute name, a comparison operator, and the attribute value.

The **Same as Base Filter** check box allows you to search more efficiently if all users qualified in the base DN are also qualified for shell access privileges. Normally, the LDAP query to retrieve users combines the base filter with the shell access filter. If the shell access filter was the same as the base filter, the same query runs twice, which is unnecessarily time-consuming. You can use the **Same as Base Filter** option to run the query only once for both purposes.

If you leave this field blank, you prevent LDAP authentication of shell access.

Configuring LDAP Shell Access

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

You can use the LDAP server to authenticate accounts for shell access on your managed device or Firepower Management Center. Specify a search filter that retrieves entries for users you want to grant shell access.

You **cannot** configure CAC authentication and authorization and shell access in the same authentication object. Instead, create and enable separate authentication objects.

The authentication object for shell access must be the first authentication object on the Firepower Management Center.

Cisco does not support external authentication for NGIPSv devices or ASA FirePOWER devices. In addition, IPv6 is not supported for shell access authentication.

Caution On all appliances, users with shell access (whether obtained through external authentication or through using the CLI `expert` command) have `sudoers` privileges in the shell, which can present a security risk. If you establish external authentication, make sure that you restrict the list of users with shell access appropriately. Similarly, when granting CLI access privileges, restrict the list of users with **Configuration** level access. Cisco strongly recommends that you do not establish additional shell users on the Firepower Management Center.

You **cannot** configure CAC authentication and authorization and shell access in the same authentication object. Checking the **CAC** check box disables the shell access configuration options on the page. Instead, create and enable separate authentication objects.

Before you begin

- Remove any internally-authenticated CLI or shell users that have the same user name as externally-authenticated users included in your shell access filter.

Procedure

Step 1 On the Create External Authentication Object page, if you want to use a shell access attribute other than the user distinguished type a **Shell Access Attribute**.

Example:

For example, on a Microsoft Active Directory Server, use the `sAMAccountName` shell access attribute to retrieve shell access users by typing `sAMAccountName` in the **Shell Access Attribute** field.

Step 2 Set a shell access account filter. You have multiple options:

- To retrieve administrative user entries based on attribute value, enter the attribute name, a comparison operator, and the attribute value you want to use as a filter, enclosed in parentheses, in the **Shell Access Filter** field. For example, if all network administrators have a `manager` attribute which has an attribute value of `shell`, you can set a base filter of `(manager=shell)`.
- To use the same filter you specified when configuring authentication settings, choose **Same as Base Filter**.
- To prevent LDAP authentication of shell access, leave the field blank.

What to do next

- Continue creating your LDAP authentication object as described in Creating Advanced LDAP Authentication Objects, on page 78.

Testing LDAP Authentication Connections

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

After you configure LDAP server and authentication settings, you can specify user credentials for a user who should be able to authenticate to test those settings.

For the **User Name**, you can enter the value for the `uid` attribute for the user you want to test with. If you are connecting to a Microsoft Active Directory Server and supplied a UI access attribute in place of `uid`, use the value for that attribute as the user name. You can also specify a fully qualified distinguished name for the user.

Use the **Password** for the same user.

The test output lists valid and invalid user names. Valid user names are unique, and can include underscores (_), periods (.), hyphens (-), and alphanumeric characters.

Note that testing the connection to servers with more than 1000 users only returns 1000 users because of web interface page size limitations.

Tip If you mistype the name or password of the test user, the test fails even if the server configuration is correct. Test the server configuration without the additional test parameters first. If that succeeds supply a user name and password to test with the specific user.

Procedure

Step 1 On the Add External Authentication Object page, enter a **User Name** and **Password**.

Example:

For example, to test to see if you can retrieve the `JSmith` user credentials at the Example company, enter `JSmith` and `password`.

Step 2 Click **Test**. You have two options:

- If the test succeeds, the test output appears at the bottom of the page. Click **Save**.
- If the test fails, see Troubleshooting LDAP Authentication Connections, on page 92 for suggestions for troubleshooting the connection.

Troubleshooting LDAP Authentication Connections

If you create an LDAP authentication object and it either does not succeed in connecting to the server you select, or does not retrieve the list of users you want, you can tune the settings in the object.

If the connection fails when you test it, try the following suggestions to troubleshoot your configuration:

- Use the messages displayed at the top of the screen and in the test output to determine which areas of the object are causing the issue.
- Check that the user name and password you used for the object are valid:
  - Check that the user has the rights to browse to the directory indicated in your base distinguished name by connecting to the LDAP server using a third-party LDAP browser.
  - Check that the user name is unique to the directory information tree for the LDAP server.
  - If you see an LDAP bind error 49 in the test output, the user binding for the user failed. Try authenticating to the server through a third-party application to see if the binding fails through that connection as well.
- Check that you have correctly identified the server:
  - Check that the server IP address or host name is correct.
  - Check that you have TCP/IP access from your local appliance to the authentication server where you want to connect.
  - Check that access to the server is not blocked by a firewall and that the port you have configured in the object is open.
  - If you are using a certificate to connect via TLS or SSL, the host name in the certificate must match the host name used for the server.
  - Check that you have not used an IPv6 address for the server connection if you are authenticating shell access.
  - If you used server type defaults, check that you have the correct server type and click **Set Defaults** again to reset the default values.

- If you typed in your base distinguished name, click **Fetch DNs** to retrieve all the available base distinguished names on the server, and select the name from the list.
- If you are using any filters, access attributes, or advanced settings, check that each is valid and typed correctly.
- If you are using any filters, access attributes, or advanced settings, try removing each setting and testing the object without it.
- If you are using a base filter or a shell access filter, make sure that the filter is enclosed in parentheses and that you are using a valid comparison operator.
- To test a more restricted base filter, try setting it to the base distinguished name for the user to retrieve just that user.
- If you are using an encrypted connection:
  - Check that the name of the LDAP server in the certificate matches the host name that you use to connect.
  - Check that you have not used an IPv6 address with an encrypted server connection.
- If you are using a test user, make sure that the user name and password are typed correctly.
- If you are using a test user, remove the user credentials and test the object.
- Test the query you are using by connecting to the LDAP server via the command line on the appliance you want to connect from using this syntax:

  ```
  ldapsearch -x -b 'base_distinguished_name'
  -h LDAPserver_ip_address -p port -v -D
  'user_distinguished_name' -W 'base_filter'
  ```

  For example, if you are trying to connect to the security domain on `myrtle.example.com` using the `domainadmin@myrtle.example.com` user and a base filter of (`cn=*`), you could test the connection using this statement:

  ```
  ldapsearch -x -b 'CN=security,DC=myrtle,DC=example,DC=com'
  -h myrtle.example.com -p 389 -v -D
  'domainadmin@myrtle.example.com' -W '(cn=*)'
  ```

If you can test your connection successfully but authentication does not work after you deploy a platform settings policy, check that authentication and the object you want to use are both enabled in the platform settings policy that is applied to the appliance.

If you connect successfully but want to adjust the list of users retrieved by your connection, you can add or change a base filter or shell access filter or use a more restrictive or less restrictive base DN.

RADIUS Authentication

The Remote Authentication Dial In User Service (RADIUS) is an authentication protocol used to authenticate, authorize, and account for user access to network resources. You can create an authentication object for any RADIUS server that conforms to RFC 2865.

When a user authenticated on a RADIUS server logs in for the first time, the user receives the roles specified for that user in the authentication object. If the user is not listed for any of the user roles, they receive the default access role you selected in the authentication object. If no default access role is selected in the authentication object, they receive the default access role set in the platform settings policy. You can modify a user's roles, if needed, unless the settings are granted through the user lists in the authentication object. Note that when a user authenticated on a RADIUS server using attribute matching attempts to log in for the first time, the login is rejected as the user account is created. The user must log in a second time.

Note Before enabling external authentication on 7000 or 8000 Series devices, remove any internally-authenticated CLI users that have the same user name as externally-authenticated users included in your shell access filter.

The Firepower System implementation of RADIUS supports the use of SecurID® tokens. When you configure authentication by a server using SecurID, users authenticated against that server append the SecurID token to the end of their SecurID PIN and use that as their password when they log into a Cisco system. As long as SecurID is configured correctly to authenticate users outside the Firepower System, those users can log into a Firepower Management Center or 7000 or 8000 Series device using their PIN plus the SecurID token without any additional configuration.

Creating RADIUS Authentication Objects

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

When you create a RADIUS authentication object, you define settings that let you connect to an authentication server. You also grant user roles to specific and default users. If your RADIUS server returns custom attributes for any users you plan to authenticate, you must define those custom attributes. Optionally, you can also configure CLI or shell access authentication.

In a multidomain deployment, external authentication objects are only available in the domain in which they are created.

Before you begin

- Confirm that you have TCP/IP access from your local appliance to the authentication server where you want to connect.

Procedure

Step 1 Choose **System** > **Users**.

Step 2 Click the **External Authentication** tab.

Step 3 Click **Add External Authentication Object**.

Step 4 Choose **RADIUS** from the **Authentication Method** drop-down list.

Step 5 Identify the authentication server as described in Configuring RADIUS Connection Settings, on page 96.

Step 6 Configure user roles as described in Configuring RADIUS User Roles, on page 98.

Step 7 Optionally, configure shell access as described in Configuring RADIUS Shell Access, on page 99.

Step 8 Optionally, define custom attributes as described in Defining Custom RADIUS Attributes, on page 100.

Step 9 Test your configuration as described in Testing RADIUS Authentication Connections, on page 101.

Example

The following figure illustrates a sample RADIUS login authentication object for a server running FreeRADIUS with an IP address of 10.10.10.98. Note that the connection uses port 1812 for access, and note that connections to the server time out after 30 seconds of disuse, then retry three times before attempting to connect to a backup authentication server.

This example illustrates important aspects of RADIUS user role configuration:

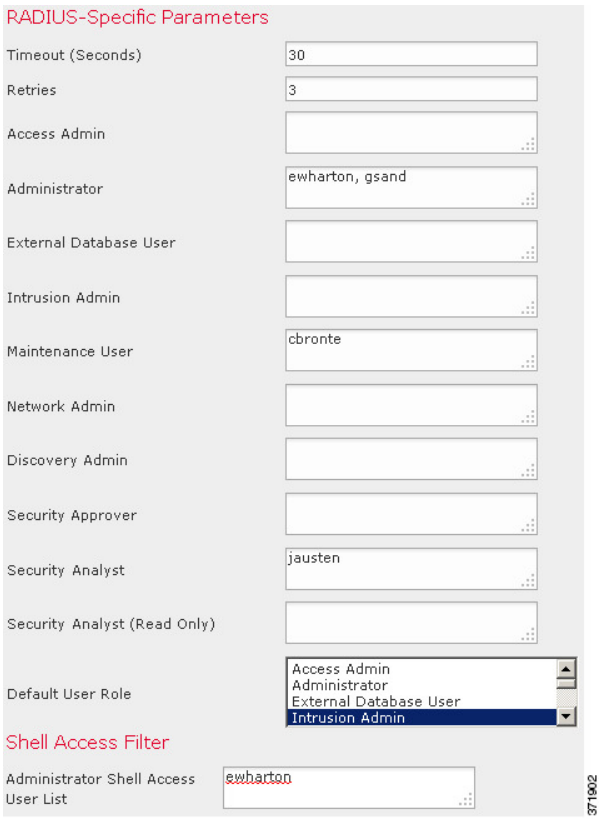
Users `ewharton` and `gsand` are granted administrative access to appliances where this authentication object is enabled.

The user `cbronte` is granted Maintenance User access to appliances where this authentication object is enabled.

The user `jausten` is granted Security Analyst access to appliances where this authentication object is enabled.

The user `ewharton` can log into the appliance using a shell account.

The following graphic depicts the role configuration for the example:

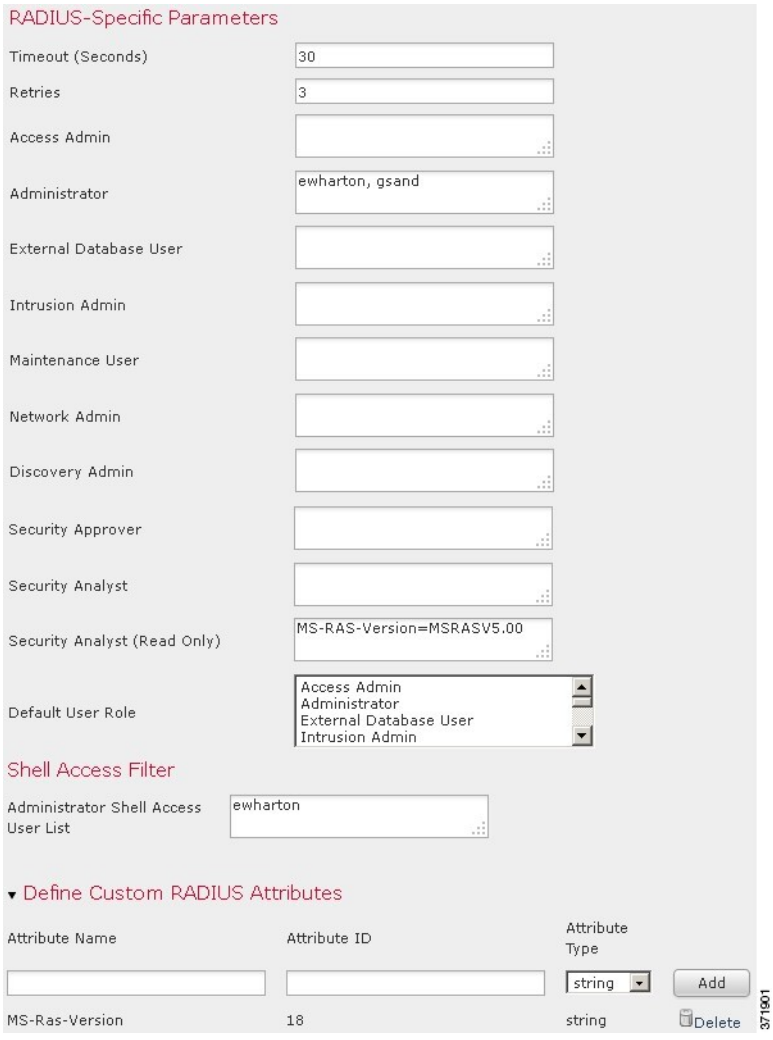


Example

You can use an attribute-value pair to identify users who should receive a particular user role. If the attribute you use is a custom attribute, you must define the custom attribute.

The following figure illustrates the role configuration and custom attribute definition in a sample RADIUS login authentication object for the same FreeRADIUS server as in the previous example.

In this example, however, the `MS-RAS-Version` custom attribute is returned for one or more of the users because a Microsoft remote access server is in use. Note the `MS-RAS-Version` custom attribute is a string. In this example, all users logging in to RADIUS through a Microsoft v. 5.00 remote access server should receive the Security Analyst (Read Only) role, so you enter the attribute-value pair of `MS-RAS-Version=MSRASV5.00` in the **Security Analyst (Read Only)** field.



Configuring RADIUS Connection Settings

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

When you create a RADIUS authentication object, you first specify the primary and backup server and server port where you want the local appliance (managed device or Firepower Management Center) to connect for authentication.

Note For RADIUS to function correctly, you must open its authentication and accounting ports (by default, 1812 and 1813) on your firewall.

If you specify a backup authentication server, you can set a timeout for the connection attempt to the primary server. If the number of seconds indicated in the **Timeout** field (or the timeout on the LDAP server) elapses without a response from the primary authentication server, the appliance then re-queries the primary server.

After the appliance re-queries the primary authentication server the number of times indicated by the **Retries** field and the number of seconds indicated in the **Timeout** field again elapses without a response from the primary authentication server, the appliance then rolls over to the backup server.

If, for example, the primary server has RADIUS disabled, the appliance queries the backup server. If RADIUS is running on the port of the primary RADIUS server and for some reason refuses to service the request (due to misconfiguration or other issues), however, the failover to the backup server does not occur.

Procedure

Step 1 Choose **System** > **Users**.

Step 2 Click the **External Authentication** tab.

Step 3 Click **Create External** > **Authentication Object**.

Step 4 Choose **RADIUS** from the **Authentication Method** drop-down list.

Step 5 Enter a **Name** and **Description** for the authentication server.

Step 6 Enter the IP address or host name for the primary RADIUS server where you want to obtain authentication data in the **Primary Server Host Name/IP Address** field.

Note IPv6 addresses are not supported for shell authentication. To allow shell authentication when using an IPv6 address for your primary RADIUS server, set up an authentication object using an IPv4 address for the server and use that IPv4 object as the first authentication object on the Firepower Management Center.

Step 7 Optionally, modify the port used by the primary RADIUS authentication server in the **Primary Server Port** field.

Note If your authentication port and accounting port numbers are not sequential, leave this field blank. The system then determines RADIUS port numbers from the `radius` and `radacct` data in your appliance's `/etc/services` file.

Step 8 Enter the **RADIUS Secret Key** for the primary RADIUS authentication server.

Step 9 Optionally, enter the IP address or host name for the backup RADIUS authentication server where you want to obtain authentication data in the **Backup Server Host Name/IP Address** field.

Step 10 If you set a backup server, modify the **Backup Server Port**, **RADIUS Secret Key**, and **Timeout** and enter the number of times the primary server connection should be tried before rolling over to the backup connection in the **Retries** field.

Note If your authentication port and accounting port numbers are not sequential, leave this field blank. The system then determines RADIUS port numbers from the `radius` and `radacct` data in your appliance's `/etc/services` file.

What to do next

- Continue creating your RADIUS authentication object as described in Creating RADIUS Authentication Objects, on page 94.

Configuring RADIUS User Roles

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

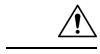
When a user logs in, the Firepower System checks the RADIUS server and grants access rights depending on the RADIUS configuration:

- If specific access rights are not configured for a user and a default access role is not specified, when a new user logs in, the Firepower System authenticates the user against the RADIUS server and then grants user rights based on the default access role (or roles) set in the platform settings policy.
- If a new user is not specified on any lists and default access roles are specified in the **Default User Role** list of the authentication object, the user is assigned those access roles.
- If you add a user to the list for one or more specific role, that user receives all assigned access roles.

You can also use attribute-value pairs, rather than user names, to identify users who should receive a particular user role. For example, if you know all users who should be Security Analysts have the value `Analyst` for their `User-Category` attribute, you can enter `User-Category=Analyst` in the Security Analyst List field to grant that role to those users.

You can assign a default user role (or roles) to be assigned to any users that are authenticated externally but not listed for a specific role. You can specify multiple roles in the **Default User Role** list.

You cannot remove the minimum access rights for users assigned an access role because of RADIUS user list membership through the Firepower System user management page. You can, however, assign additional rights.



Caution If you want to change the minimum access setting for a user, you must not only move the user from one list to another in the RADIUS Specific Parameters section or change the user's attribute on the RADIUS server, you must redeploy the configuration to the managed device and remove the assigned user right on the user management page.

Before you begin

- Define custom attributes if you plan to use them to set user role membership, as described in Defining Custom RADIUS Attributes, on page 100.

Procedure

Step 1 On the Create External Authentication Object page, in the fields that correspond to Firepower System user roles, enter the name of each user or identifying attribute-value pair that should be assigned to those roles.

Separate usernames and attribute-value pairs with commas.

Example:

For example, to grant the Administrator role to the users `jsmith` and `jdoe`, enter `jsmith, jdoe` in the **Administrator** field. As another example, to grant the Maintenance User role to all users with a `User-Category` value of `Maintenance`, enter `User-Category=Maintenance` in the **Maintenance User** field.

Step 2 Choose the default minimum access role for users that do not belong to any of the specified groups from the **Default User Role** list.

What to do next

- Continue creating your RADIUS authentication object as described in Creating RADIUS Authentication Objects, on page 94.

Configuring RADIUS Shell Access

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

You can also use the RADIUS server to authenticate accounts for CLI or shell access on your local appliance (managed device or Firepower Management Center). Specify user names for users you want to grant CLI or shell access.

Note IPv6 addresses are not supported for shell authentication. If you configure a primary RADIUS server with an IPv6 address and also configure administrative shell access, the shell access settings are ignored. To allow shell authentication when using an IPv6 address for your primary RADIUS server, set up another authentication object using an IPv4 address for the server and use that object as the first authentication object on the Firepower Management Center.

With the exception of the admin account, the shell access list you set on the RADIUS authentication object entirely controls CLI or shell access on the appliance. CLI or shell users are configured as local users on the appliance when you deploy the platform settings policy. Note that when a user authenticated on a RADIUS server using attribute matching attempts to log in for the first time, the login is rejected as the user account is created. The user must log in a second time.

Note that a home directory for each CLI or shell user is created on login, and when an RADIUS shell access user account is disabled (by disabling the RADIUS connection), the directory remains, but the user shell is set to `/bin/false` in `/etc/password` to disable the shell. If the user then is re-enabled, the shell is reset, using the same home directory.

CLI or shell users can log in using user names with lowercase, uppercase, or mixed case letters. Login authentication for the CLI or shell is case sensitive.

Caution On all appliances, users with shell access (whether obtained through external authentication or through using the CLI `expert` command) have `sudoers` privileges in the shell, which can present a security risk. If you establish external authentication, make sure that you restrict the list of users with shell access appropriately. Similarly, when granting CLI access privileges, restrict the list of users with **Configuration** level access. Cisco strongly recommends that you do not establish additional shell users on the Firepower Management Center.

Procedure

On the Create External Authentication Object page, enter the user names, separated by commas, in the **Administrator Shell Access User List** field.

Note If you choose not to specify a shell access filter, a warning displays when you save the authentication object to confirm that you meant to leave the filter blank.

What to do next

- Continue creating your RADIUS authentication object as described in Creating RADIUS Authentication Objects, on page 94.

Defining Custom RADIUS Attributes

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

If your RADIUS server returns values for attributes not included in the `dictionary` file in `/etc/radiusclient/` and you plan to use those attributes to set user roles for users with those attributes, you need to define those attributes in the login authentication object. You can locate the attributes returned for a user by looking at the user's profile on your RADIUS server.

When you define an attribute, you provide the name of the attribute, which consists of alphanumeric characters. Note that words in an attribute name should be separated by dashes rather than spaces. You also provide the attribute ID, which should be an integer and should not conflict with any existing attribute IDs in the `etc/radiusclient/dictionary` file. You also specify the type of attribute: string, IP address, integer, or date.

When you create a RADIUS authentication object, a new dictionary file for that object is created on the appliance in the `/var/sf/userauth` directory. Any custom attributes you add to the authentication object are added to the dictionary file.

In a multidomain deployment, external authentication objects are only available in the domain in which they are created.

Procedure

Step 1 On the Add External Authentication Object page, click the arrow to expand the Define Custom RADIUS Attributes section.

Step 2 Enter an attribute name in the **Attribute Name** field.

Step 3 Enter the attribute ID, in integer form, in the **Attribute ID** field.

Step 4 Choose the type of attribute from the **Attribute Type** drop-down list.

Step 5 Click **Add** to add the custom attribute to the authentication object.

Tip You can remove a custom attribute from an authentication object by clicking **Delete** next to the attribute.

Example

If a RADIUS server is used on a network with a Cisco router, you might want to use the `Ascend-Assign-IP-Pool` attribute to grant a specific role to all users logging in from a specific IP address pool. `Ascend-Assign-IP-Pool` is an integer attribute that defines the address pool where the user is allowed to log in, with the integer indicating the number of the assigned IP address pool.

To declare that custom attribute, you create a custom attribute with an attribute name of `Ascend-IP-Pool-Definition`, an attribute ID of `218`, and an attribute type of `integer`.

You could then enter `Ascend-Assign-IP-Pool=2` in the **Security Analyst (Read Only)** field to grant read-only security analyst rights to all users with an `Ascend-IP-Pool-Definition` attribute value of `2`.

What to do next

- Continue creating your RADIUS authentication object as described in Creating RADIUS Authentication Objects, on page 94.

Testing RADIUS Authentication Connections

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin

After you configure RADIUS connection, user role, and custom attribute settings, you can specify user credentials for a user who should be able to authenticate to test those settings.

For the user name, you can enter the user name for the user you want to test with.

Note that testing the connection to servers with more than 1000 users only returns 1000 users because of UI page size limitations.

Tip If you mistype the name or password of the test user, the test fails even if the server configuration is correct. To verify that the server configuration is correct, click **Test** without entering user information in the **Additional Test Parameters** field first. If that succeeds, supply a user name and password to test with the specific user.

Procedure

Step 1 On the Add External Authentication Object page, in the **User Name** and **Password** fields, enter the user name and password for the user whose credentials should be used to validate access to the RADIUS server.

Example:

For example, to test to see if you can retrieve the `jsmith` user credentials at our example company, enter `jsmith`.

Step 2 Choose **Show Details**, and click **Test**.

Step 3 If the test succeeds, click **Save**.

Single Sign-on (SSO)

Single sign-on (SSO) enables integration between Cisco Security Manager (CSM) Version 4.7 or higher and the Firepower Management Center, which allows you to access the Firepower Management Center from CSM without additional authentication to log in. When managing an ASA FirePOWER module, you may want to modify the policies deployed to the module. You can select the managing Firepower Management Center in CSM and launch it in a web browser.

If you have access based on your user role, the system navigates you to the Device tab of the Device Management page for the device you cross-launched from in CSM. Otherwise, the system navigates you to the Summary Dashboard page (**Overview** > **Dashboards**), except for user accounts with no dashboard access, which use the Welcome page.

Note You cannot login with single sign-on if your organization uses CACs for authentication.

Related Topics

Security Certifications Compliance

Configuring SSO

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	ASA FirePOWER	Any	Admin

You must set up a one-way, encrypted authentication path from CSM to the Firepower Management Center before you configure Single sign-on.

In NAT environments, the Firepower Management Center and CSM must reside on the same side of the NAT boundary. You must provide specific criteria to enable communications between CSM and the Firepower Management Center.

Note You cannot login with single sign-on if your organization uses CACs for authentication.

Procedure

Step 1 From CSM, generate an SSO shared encryption key that identifies the connection. See your CSM documentation for more information.

Step 2 From the Firepower Management Center, choose **System** > **Users**.

Step 3 Choose **CSM Single Sign-on**.

Step 4 Enter the **CSM hostname** or **IP** address and the server **Port**.

Step 5 Enter the **Shared key** that you generated from CSM.

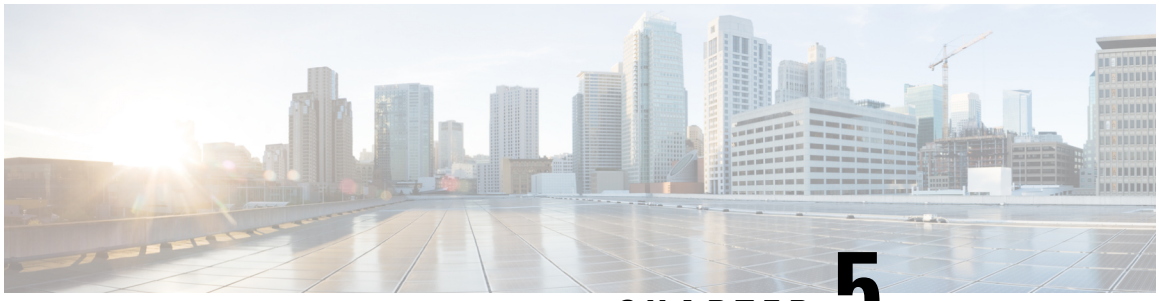
Step 6 Optionally, if you want to use the Firepower Management Center's proxy server to communicate with CSM, choose the **Use Proxy For Connection** check box.

Step 7 Click **Submit**.

Step 8 Click **Confirm Certificate** to save the Certificate.
You can now log in from CSM to the Firepower Management Center without an additional login.

Related Topics

Configure Management Interfaces



CHAPTER 5

Licensing the Firepower System

The Licensing chapter of the Firepower Management Center Configuration Guide provides in-depth information about the different license types, service subscriptions, licensing requirements and more. The chapter also provides procedures and requirements for deploying Smart and Classic licenses and licensing for air-gapped solutions.

The following topics explain how to license Firepower.

- About Firepower Licenses, on page 105
- Requirements and Prerequisites for Licensing, on page 106
- License Requirements for Firepower Management Center, on page 106
- Evaluation License Caveats, on page 106
- Smart vs. Classic Licenses, on page 107
- License Firepower Threat Defense Devices (FTD), on page 107
- License Classic Devices (Firepower 7000/8000 Series, ASA FirePOWER, and NGIPSv), on page 124
- How to Convert a Classic License for Use on an FTD Device, on page 132
- Assign Licenses to Managed Devices from the Device Management Page, on page 134
- Additional Information about Firepower Licensing, on page 135

About Firepower Licenses

Your Firepower products (Firepower Management Center and managed devices) include licenses for basic operation, but some features require separate licensing or service subscriptions, as described in this chapter.

A "right-to-use" license does not expire, but service subscriptions require periodic renewal.

The type of license your products require (Smart or Classic) depends on the software you use, not on the hardware it runs on.

Note "NGFW" means different things to different people, so this documentation does not use this term.

Requirements and Prerequisites for Licensing

Model Support

Any, but the specific licenses requires per model differ as indicated in the procedures.

Supported Domains

Global, except where indicated.

User Roles

- Admin

License Requirements for Firepower Management Center

Firepower Management Center allows you to assign licenses to managed devices and manage licenses for the system.

A single Firepower Management Center can manage both devices that require Classic licenses and devices that require Smart Licenses.

Hardware FMC

A hardware Firepower Management Center does not require purchase of additional licenses or service subscriptions in order to manage devices.

Virtual FMC

Firepower Management Center Virtual has additional licensing requirements. See Firepower Management Center Virtual Licenses, on page 106.

Firepower Management Center Virtual Licenses

If a single FMCv manages Firepower Threat Defense devices that are configured in a high availability pair, you still need one entitlement for each device (*not* one entitlement for the pair of FTDs.)

This entitlement appears in Cisco Smart Software Manager as **Firepower MCv Device License** with different numbers of entitlements.

Evaluation License Caveats

Not all functionality is available with an evaluation license, functionality under an evaluation license may be partial, and transition from evaluation licensing to standard licensing may not be seamless.

For example, if you have Firepower Threat Defense devices configured in a cluster, and you switch from an evaluation license to Smart Licensing, service will be interrupted when you deploy the change.

Review information about evaluation license caveats in information about particular features in this Licensing chapter and in the chapters related to deploying each feature.

Smart vs. Classic Licenses

For managed devices, the licenses you need (Smart or Classic) depend on the software that runs on the device.

Any FMC can simultaneously manage devices with Smart and Classic licenses. You must configure each type of licensing separately.

Software	License Type	More Information
Firepower Management Center (hardware)	None	FMC hardware models themselves require no license.
Firepower Management Center Virtual	Device entitlements	See Firepower Management Center Virtual Licenses, on page 106.
Firepower Threat Defense Firepower Threat Defense Virtual	Smart	See the topics under License Firepower Threat Defense Devices (FTD), on page 107.
NGIPS software: • Firepower 7000/8000 series • ASA FirePOWER • NGIPSv	Classic	See License Classic Devices (Firepower 7000/8000 Series, ASA FirePOWER, and NGIPSv), on page 124.
All other software products, including those that run on Firepower hardware	See licensing information for your software product.	

License Firepower Threat Defense Devices (FTD)

Firepower Threat Defense devices require Smart Licensing.

Cisco Smart Licensing lets you purchase and manage a pool of licenses centrally. Unlike product authorization key (PAK) licenses, Smart Licenses are not tied to a specific serial number or license key. Smart Licensing lets you assess your license usage and needs at a glance.

In addition, Smart Licensing does not prevent you from using product features that you have not yet purchased. You can start using a license immediately, as long as you are registered with the Cisco Smart Software Manager, and purchase the license later. This allows you to deploy and use a feature, and avoid delays due to purchase order approval.

Smart Software Manager (CSSM)

When you purchase one or more Smart Licenses for Firepower features, you manage them in the Cisco Smart Software Manager: http://www.cisco.com/web/ordering/smart-software-manager/index.html. The Smart Software Manager lets you create a primary account for your organization.

By default, your licenses are assigned to the Default Virtual Account under your primary account. As the account administrator, you can create additional virtual accounts; for example, for regions, departments, or subsidiaries. Multiple virtual accounts help you manage large numbers of licenses and appliances.

You manage licenses and appliances by virtual account. Only that virtual account's appliances can use the licenses assigned to the account. If you need additional licenses, you can transfer an unused license from another virtual account. You can also transfer appliances between virtual accounts.

For each virtual account, you can create a Product Instance Registration Token. Enter this token ID when you deploy each Firepower Management Center, or when you register an existing FMC. You can create a new token if an existing token expires. An expired token does not affect a registered FMC that used this token for registration, but you cannot use an expired token to register a FMC. Also, a registered FMC becomes associated with a virtual account based on the token you use.

For more information about the Cisco Smart Software Manager, see *Cisco Smart Software Manager User Guide* or https://www.cisco.com/c/en/us/buy/smart-accounts/software-manager.html or the online help in CSSM, also available from: https://www.cisco.com/web/fw/softwareworkspace/smartlicensing/SSMCompiledHelps/.

Periodic Communication with the License Authority

In order to maintain your product license entitlement, your product must communicate periodically with the Cisco License Authority.

When you use a Product Instance Registration Token to register a Firepower Management Center, the appliance registers with the Cisco License Authority. The License Authority issues an ID certificate for communication between the Firepower Management Center and the License Authority. This certificate is valid for one year, although it will be renewed every six months. If an ID certificate expires (usually in nine months or a year with no communication), the Firepower Management Center reverts to a deregistered state and licensed features usage become suspended.

The Firepower Management Center communicates with the License Authority on a periodic basis. If you make changes in the Smart Software Manager, you can refresh the authorization on the Firepower Management Center so the changes immediately take effect. You also can wait for the appliance to communicate as scheduled.

Your Firepower Management Center must either have direct Internet access to the License Authority through the Cisco Smart Software Manager or access through the Smart Software Satellite Server at scheduled time periods. Normal license communication occurs every 30 days, but with the grace period, your appliance will operate for up to 90 days without calling home. You must contact the License Authority before 90 days have passed.

Service Subscriptions for FTD Features

Some features require a service subscription.

A service subscription enables a specific Firepower feature on a managed device for a set length of time. Service subscriptions can be purchased in one-, three-, or five-year terms. If a subscription expires, Cisco

notifies you that you must renew the subscription. If a subscription expires for a Firepower Threat Defense device, you can continue to use the related features.

Table 13: Service Subscriptions and Corresponding Smart Licenses

Subscription You Purchase	Smart Licenses You Assign in Firepower System
T	Threat
TC	Threat + URL Filtering
TM	Threat + Malware
TMC	Threat + URL Filtering + Malware
URL	URL Filtering (can be added to Threat or used without Threat)
AMP	Malware (the Threat license is also required)

Your purchase of a managed device that uses Smart Licenses automatically includes a Base license. This license is perpetual and enables system updates. All service subscriptions are optional for Firepower Threat Defense devices.

FTD License Types and Restrictions

This section describes the types of Smart Licenses available in a Firepower System deployment. The Firepower Management Center requires Smart Licenses to manage Firepower Threat Defense devices.

The following table summarizes Firepower System Smart Licenses.

Table 14: Firepower System Smart Licenses

License You Assign in Firepower System	Subscription You Purchase	Duration	Granted Capabilities
Base (Base licenses are automatically assigned with all Firepower Threat Defense devices)	No subscription required (license is included with device)	Perpetual	User and application control Switching and routing NAT For details, see Base Licenses, on page 110.
Threat	• T • TC (Threat + URL) • TMC (Threat + Malware + URL)	Term-based	Intrusion detection and prevention File control Security Intelligence filtering For details, see Threat Licenses, on page 111

License You Assign in Firepower System	Subscription You Purchase	Duration	Granted Capabilities
Malware	• TM (Threat + Malware) • TMC (Threat + Malware + URL) • AMP	Term-based	AMP for Networks (network-based Advanced Malware Protection) Cisco Threat Grid File storage For details, see Malware Licenses for Firepower Threat Defense Devices, on page 111 and License Requirements for File and Malware Policies, on page 1403.
URL Filtering	• TC (Threat + URL) • TMC (Threat + Malware + URL) • URL	Term-based	Category and reputation-based URL filtering For details, see URL Filtering Licenses for Firepower Threat Defense Devices, on page 112.
Firepower Management Center Virtual	No subscription required.	Perpetual	The platform license determines the number of devices the virtual appliance can manage. For details, see Firepower Management Center Virtual Licenses, on page 106.
Export-Controlled Features	No subscription required.	Perpetual	Features that are subject to national security, foreign policy, and anti-terrorism laws and regulations; see Licensing for Export-Controlled Functionality, on page 112.

Base Licenses

A base license is automatically included with every purchase of a Firepower Threat Defense or Firepower Threat Defense Virtual device.

The Base license allows you to:

- configure your Firepower Threat Defense devices to perform switching and routing (including DHCP relay and NAT)
- configure Firepower Threat Defense devices as a high availability pair
- configure security modules as a cluster within a Firepower 9300 chassis (intra-chassis clustering)
- configure Firepower 9300 or Firepower 4100 series devices running Firepower Threat Defense as a cluster (inter-chassis clustering)

- implement user and application control by adding user and application conditions to access control rules

Threat and malware detection and URL filtering features require additional, optional licenses.

Base licenses are automatically added to the Firepower Management Center for every Firepower Threat Defense device you register.

Malware Licenses for Firepower Threat Defense Devices

A Malware license for Firepower Threat Defense devices allows you to perform Cisco Advanced Malware Protection (AMP) with AMP for Networks and Cisco Threat Grid. With this feature, you can use Firepower Threat Defense devices to detect and block malware in files transmitted over your network. To support this feature license, you can purchase the Malware (AMP) service subscription as a stand-alone subscription or in combination with Threat (TM) or Threat and URL Filtering (TMC) subscriptions.



Note Firepower Threat Defense managed devices with Malware licenses enabled periodically attempt to connect to the AMP cloud even if you have not configured dynamic analysis. Because of this, the device's Interface Traffic dashboard widget shows transmitted traffic; this is expected behavior.

You configure AMP for Networks as part of a file policy, which you then associate with one or more access control rules. File policies can detect your users uploading or downloading files of specific types over specific application protocols. AMP for Networks allows you to use local malware analysis and file preclassification to inspect a restricted set of those file types for malware. You can also download and submit specific file types to the Cisco Threat Grid cloud for dynamic and Spero analysis to determine whether they contain malware. For these files, you can view the network file trajectory, which details the path the file has taken through your network. The Malware license also allows you to add specific files to a file list and enable the file list within a file policy, allowing those files to be automatically allowed or blocked on detection.

If you disable all your Malware licenses, the system stops querying the AMP cloud, and also stops acknowledging retrospective events sent from the AMP cloud. You cannot re-deploy existing access control policies if they include AMP for Networks configurations. Note that for a very brief time after a Malware license is disabled, the system can use existing cached file dispositions. After the time window expires, the system assigns a disposition of `Unavailable` to those files.

Note that a Malware license is required only if you deploy AMP for Networks and Cisco Threat Grid. Without a Malware license, the Firepower Management Center can receive AMP for Endpoints malware events and indications of compromise (IOC) from the AMP cloud.

See also important information at License Requirements for File and Malware Policies, on page 1403.

Threat Licenses

A Threat license allows you to perform intrusion detection and prevention, file control, and Security Intelligence filtering:

- *Intrusion detection and prevention* allows you to analyze network traffic for intrusions and exploits and, optionally, drop offending packets.
- *File control* allows you to detect and, optionally, block users from uploading (sending) or downloading (receiving) files of specific types over specific application protocols. *AMP for Networks*, which requires a Malware license, allows you to inspect and block a restricted set of those file types based on their dispositions.

- *Security Intelligence filtering* allows you to block —deny traffic to and from—specific IP addresses, URLs, and DNS domain names, before the traffic is subjected to analysis by access control rules. Dynamic feeds allow you to immediately block connections based on the latest intelligence. Optionally, you can use a "monitor-only" setting for Security Intelligence filtering.

You can purchase a Threat license as a stand-alone subscription (T) or in combination with URL Filtering (TC), Malware (TM), or both (TMC).

If you disable Threat on managed devices, the Firepower Management Center stops acknowledging intrusion and file events from the affected devices. As a consequence, correlation rules that use those events as a trigger criteria stop firing. Additionally, the Firepower Management Center will not contact the internet for either Cisco-provided or third-party Security Intelligence information. You cannot re-deploy existing intrusion policies until you re-enable Threat.

URL Filtering Licenses for Firepower Threat Defense Devices

The URL Filtering license allows you to write access control rules that determine the traffic that can traverse your network based on URLs requested by monitored hosts, correlated with information about those URLs. To support this feature license, you can purchase the URL Filtering (URL) service subscription as a stand-alone subscription or in combination with Threat (TC) or Threat and Malware (TMC) subscriptions.



Tip Without a URL Filtering license, you can specify individual URLs or groups of URLs to allow or block. This gives you granular, custom control over web traffic, but does not allow you to use URL category and reputation data to filter network traffic.

Although you can add category and reputation-based URL conditions to access control rules without a URL Filtering license, the Firepower Management Center will not download URL information. You cannot deploy the access control policy until you first add a URL Filtering license to the Firepower Management Center, then enable it on the devices targeted by the policy.

If you disable the URL Filtering license on managed devices, you may lose access to URL filtering. If your license expires or if you disable it, access control rules with URL conditions immediately stop filtering URLs, and your Firepower Management Center can no longer download updates to URL data. You cannot re-deploy existing access control policies if they include rules with category and reputation-based URL conditions.

Licensing for Export-Controlled Functionality

Features that require export-controlled functionality

Certain software features are subject to national security, foreign policy, and anti-terrorism laws and regulations. These export-controlled features include:

- Site to Site VPN with strong encryption
- SSH platform policy with strong encryption
- SSL policy with strong encryption
- Functionality such as SNMPv3 with strong encryption

How to determine whether export-controlled functionality is currently enabled for your system

To determine whether export-controlled functionality is currently enabled for your system: Go to **System** > **Licenses** > **Smart Licenses** and see if **Export-Controlled Features** displays **Enabled**.

About enabling export-controlled functionality

If **Export-Controlled Features** shows **Disabled** and you want to use features that require strong encryption:

- If there is *no* option to enable export-controlled functionality when you generate a new Product Instance Registration Token in Cisco Smart Software Manager (CSSM):

  Your Smart Account is not authorized to use this functionality.

- If the option to use export-controlled functionality appears when you generate a new Product Instance Registration Token in Cisco Smart Software Manager:
  - In order to use export-controlled functionality, your Smart Account must be enabled for this functionality before you license your Firepower Management Center.
  - After export-controlled functionality is enabled for your Smart Account in Cisco Smart Software Manager (CSSM), you must re-register your Firepower Management Center using a new Product Instance Registration Token.
  - When you create the new Product Instance Registration Token, you must select the option "Allow export-controlled functionality on the products registered with this token." This option is enabled by default if this functionality is permitted for your Smart Account.
  - After you install a token with export-controlled functionality on your Firepower Management Center and assign the relevant licenses to managed Firepower Threat Defense devices:
    - Reboot each device to make the newly-enabled features available.
    - In a high availability deployment, the active and standby devices must be rebooted together to avoid an Active-Active condition.

More Information

For general information about export controls, see https://www.cisco.com/c/en/us/about/legal/global-export-trade.html.

Licensing for FTD Clusters

In addition to information in this Licensing chapter, see:

- Licenses for Clustering, on page 713
- FMC: Add a Cluster, on page 727.

Create a Smart Account to Hold Your Licenses

A Smart Account is required for Smart Licenses and can also hold Classic licenses.

You should set up this account before you purchase Smart Licenses.

Before you begin

Your account representative or reseller may have set up a Smart Account on your behalf. If so, obtain the necessary information to access the account from that person instead of using this procedure, then verify that you can access the account.

For general information about Smart Accounts, see http://www.cisco.com/go/smartaccounts.

Procedure

Step 1 Request a Smart Account:

For instructions, see https://community.cisco.com/t5/licensing-enterprise-agreements/request-a-smart-account-for-customers/ta-p/3636515?attachment-id=150577 .

For additional information, see https://communities.cisco.com/docs/DOC-57261.

Step 2 Wait for an email telling you that your Smart Account is ready to set up. When it arrives, click the link it contains, as directed.

Step 3 Set up your Smart Account:

Go here: https://software.cisco.com/software/company/smartaccounts/home?route=module/accountcreation.

For instructions, see https://community.cisco.com/t5/licensing-enterprise-agreements/complete-smart-account-setup-for-customers/ta-p/3636631?attachment-id=132604.

Step 4 Verify that you can access the account in the Cisco Smart Software Manager (CSSM).

Go to https://software.cisco.com/#module/SmartLicensing and sign in.

How to Configure Smart Licensing with Direct Internet Access

Procedure

Step 1 Obtain a token from the Cisco Smart Software Manager licensing portal.

See Obtain a Product License Registration Token for Smart Licensing, on page 115.

Step 2 Register your Firepower Management Center with the Smart licensing portal.

See Register Smart Licenses, on page 116. Be sure to address the prerequisites in this topic.

Step 3 Verify that your FMC registered successfully with the Smart licensing portal.

In the Firepower Management Center web interface, go to **System > Licenses > Smart Licenses**.

Product Registration should show a green checkmark.

Step 4 If you have not yet done so, add devices to your FMC.

See Add a Device to the FMC, on page 229.

Step 5 Assign licenses to the devices that are managed by your FMC.

See Assign Licenses to Multiple Managed Devices, on page 119.

Step 6 Verify that licenses are successfully installed.

See View FTD Licenses and License Status, on page 120.

What to do next

If applicable, set up licensing for high-availability and clustered deployments.

Obtain a Product License Registration Token for Smart Licensing

Before you begin

- Create a Smart Account, if you have not already done so: Visit https://software.cisco.com/smartaccounts/setup#accountcreation-account. For information, see https://www.cisco.com/c/en/us/buy/smart-accounts.html.
- Ensure that you have purchased the type and number of licenses you require.
- Verify that the licenses you need appear in your Smart Account.

  If your licenses do not appear in your Smart Account, ask the person who ordered them (for example, your Cisco sales representative or authorized reseller) to transfer those licenses to your Smart Account.
- Ideally, check the prerequisites for Register Smart Licenses, on page 116 to ensure that your registration process goes smoothly.
- Make sure you have your credentials to sign in to the Cisco Smart Software Manager.

Procedure

Step 1 Go to https://software.cisco.com.

Step 2 Click **Smart Software Licensing** (in the License section.)

Step 3 Sign in to the Cisco Smart Software Manager.

Step 4 Click **Inventory**.

Step 5 Click **General**.

Step 6 Click **New Token**.

Step 7 For **Description**, enter a name that uniquely and clearly identifies the Firepower Management Center for which you will use this token.

Step 8 Enter an expiration time within 365 days.

This determines how much time you have to register the token to a Firepower Management Center. (Your license entitlement term is independent of this setting but may start to count down even if you have not yet registered your token.)

Step 9 If you see an option to enable export-controlled functionality, and you plan to use features that require strong encryption, select this option.

Important If you see this option, you must select it now if you plan to use this functionality. You cannot enable export-controlled functionality later.

If you do not see this option but you expect to, cancel out of this procedure now and contact your Cisco account representative.

Step 10 Click **Create Token**.

Step 11 Locate your new token in the list and click **Actions**, then choose **Copy** or **Download**.

Step 12 If necessary, save your token in a safe place until you are ready to enter it into your Firepower Management Center.

What to do next

Continue with the steps in Register Smart Licenses, on page 116.

Register Smart Licenses

Register the Firepower Management Center with the Cisco Smart Software Manager.

Before you begin

- If your deployment is air-gapped, do not use this procedure. Instead, see Configure the Connection to Smart Software Manager On-Prem, on page 118.
- Ensure that the Firepower Management Center can reach the Cisco Smart Software Manager (CSSM) server at tools.cisco.com:443.
- Make sure the NTP daemon is running on your Firepower Management Center. During registration, a key exchange occurs between the NTP server and the Cisco Smart Software Manager, so time must be in sync for proper registration.

  If you are deploying Firepower Threat Defense on a Firepower 4100/9300 chassis, you must configure NTP on the Firepower chassis using the same NTP server for the chassis as for the Firepower Management Center.
- If your organization has multiple Firepower Management Center appliances, make sure each FMC has a unique name that clearly identifies and distinguishes it from other Firepower Management Center appliances that may be registered to the same virtual account. This name is critical for managing your Smart License entitlements and ambiguous names will lead to problems later.
- Generate the necessary product license registration token from the Cisco Smart Software Manager. See Obtain a Product License Registration Token for Smart Licensing, on page 115, including all prerequisites. Make sure the token is accessible from the machine from which you will access your Firepower Management Center.

Procedure

Step 1 Choose **System** > **Licenses** > **Smart Licenses**.

Step 2 Click **Register**.

Step 3 Paste the token you generated from Cisco Smart Software Manager into the **Product Instance Registration Token** field.

Make sure there are no empty spaces or blank lines at the beginning or end of the text.

Step 4 Click **Apply Changes**.

What to do next

- Add your Firepower Threat Defense devices to the Firepower Management Center; see Add a Device to the FMC, on page 229.
- Assign licenses to your Firepower Threat Defense devices; see Assign Licenses to Multiple Managed Devices, on page 119.

Smart Software Manager On-Prem Overview

As described in Periodic Communication with the License Authority, on page 108, your system must communicate regularly with Cisco to maintain your license entitlement. If you have one of the following situations, you might want to use a Smart Software Manager On-Prem (also known as SSM On-Prem, and formerly known as "Smart Software Satellite Server") as a proxy for connections to the License Authority:

- Your Firepower Management Center is offline or otherwise has limited or no connectivity (in other words, is deployed in an air-gapped network.)
- Your Firepower Management Center has permanent connectivity, but you want to manage your Smart Licenses via a single connection from your network.

Cisco Smart Software Manager On-Prem allows you to schedule synchronization or manually synchronize Smart License authorization with the Smart Software Manager.

For more information about Smart Software Manager On-Prem, see https://www.cisco.com/c/en/us/buy/smart-accounts/software-manager.html#~on-prem

How to Deploy Smart Software Manager On-Prem

Procedure

Step 1 Deploy and set up Smart Software Manager On-Prem.

See the documentation for the Smart Software Manager On-Prem, available from https://www.cisco.com/c/en/us/buy/smart-accounts/software-manager.html#~on-prem.

Step 2 Connect the Firepower Management Center to Smart Software Manager On-Prem, obtain a registration token, and register the FMC to SSM On-Prem.

See Configure the Connection to Smart Software Manager On-Prem, on page 118.

Step 3 Add devices to be managed.

See Add a Device to the FMC, on page 229.

Step 4 Assign licenses to managed devices

See Assign Licenses to Multiple Managed Devices, on page 119

Step 5 Synchronize Smart Software Manager On-Prem to the Cisco Smart Software Management Server (CSSM).

See the Smart Software Manager On-Prem documentation, above.

Step 6 Schedule ongoing synchronization times.

Configure the Connection to Smart Software Manager On-Prem

Before you begin

- Set up Smart Software Manager On-Prem (SSM On-Prem). For information, see How to Deploy Smart Software Manager On-Prem, on page 117.
- Make a note of the CN of the TLS/SSL certificate on your SSM On-Prem.
- Verify that your FMC can reach your SSM On-Prem. For example, verify that the FQDN configured as the SSM On-Prem call-home URL can be resolved by your internal DNS server.
- Go to http://www.cisco.com/security/pki/certs/clrca.cer and copy the entire body of the TLS/SSL certificate (from "-----BEGIN CERTIFICATE-----" to "-----END CERTIFICATE-----") into a place you can access during configuration.

Procedure

Step 1 Choose **System** > **Integration**.

Step 2 Click **Smart Software Satellite**.

Step 3 Select **Connect to Cisco Smart Software Satellite Server**.

Step 4 Enter the **URL** of your Smart Software Manager On-Prem, using the CN value you collected in the prerequisites of this procedure, in the following format:

`https://FQDN_or_hostname_of_your_SSM_On-Prem/Transportgateway/services/DeviceRequestHandler`

The FQDN or hostname must match the CN value of the certificate presented by your SSM On-Prem.

Step 5 Add a new **SSL Certificate** and paste the certificate text that you copied in the prerequisites for this procedure.

Step 6 Click **Apply**.

Step 7 Select **System > Licenses > Smart Licenses** and click **Register**.

Step 8 Create a new token on Smart Software Manager On-Prem.

Step 9 Copy the token.

Step 10 Paste the token into the form on the management center page.

Step 11 Click **Apply Changes**.

The management center is now registered to Smart Software Manager On-Prem.

What to do next

Complete remaining steps in How to Deploy Smart Software Manager On-Prem, on page 117.

Assign Licenses to Multiple Managed Devices

Devices managed by a Firepower Management Center obtain their licenses via the Firepower Management Center, not directly from the Cisco Smart Software Manager.

Use this procedure to enable licensing on multiple Firepower Threat Defense devices at once.



Note For an Firepower Threat Defense cluster, you apply the licenses to the cluster as a whole; note that each unit in the cluster consumes a separate license per feature.

Before you begin

- If you have not yet done so, register your devices with the Firepower Management Center. See Add a Device to the FMC, on page 229.
- Prepare licenses for distribution to managed devices: See Register Smart Licenses, on page 116

Procedure

Step 1 Choose **System** > **Licenses** > **Smart Licenses**.

Step 2 Click **Edit Licenses**.

Step 3 For each type of license you want to add to a device:

a) Click the tab for that type of license.
b) Click a device in the list on the left.
c) Click **Add** to move that device to the list on the right.
d) Repeat for each device to receive that type of license.

For now, don't worry about whether you have licenses for all of the devices you want to add.

e) Repeat this subprocedure for each type of license you want to add.
f) Click **Apply**.

What to do next

- Verify that your licenses are correctly installed. Follow the procedure in View FTD Licenses and License Status, on page 120.
- If export-controlled functionality is newly enabled, reboot each device. If devices are configured in a high-availability pair, reboot both devices at the same time to avoid an Active-Active condition.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

View FTD Licenses and License Status

To view the license status for a Firepower Management Center and its managed Firepower Threat Defense devices, use the Smart Licenses page in FMC.

For each type of license in your deployment, the page lists the total number of licenses consumed, whether the license is in compliance or out of compliance, the device type, and the domain and group where the device is deployed. You can also view the Firepower Management Center's Smart License Status.

Other than the Smart Licenses page, there are a few other ways you can view licenses:

- The Product Licensing dashboard widget provides an at-a-glance overview of your licenses.

  See Adding Widgets to a Dashboard, on page 269 and Dashboard Widget Availability by User Role, on page 257 and The Product Licensing Widget, on page 266.

- The Device Management page (**Devices** > **Device Management**) lists the licenses applied to each of your managed devices.
- The Smart License Monitor health module communicates license status when used in a health policy.

Procedure

Step 1 Choose **System** > **Licenses** > **Smart Licenses**.

Step 2 In the **Smart Licenses** table, click the arrow at the left side of each **License Type** folder to expand that folder.

Step 3 In each folder, verify that each device has a green circle with a **Check Mark** ( ) in the **License Status** column.

Note If you see duplicate Firepower Management Center Virtual licenses, each represents one managed device.

If all devices show a green circle with a **Check Mark** ( ), your devices are properly licensed and ready to use.

If you see any License Status other than a green circle with a **Check Mark** ( ), hover over the status icon to view the message.

What to do next

- If you had any devices that did NOT have a green circle with a **Check Mark** ( ), you may need to purchase more licenses.

FTD License Status

Smart Licensing

The Smart License Status section of the **System > Licenses > Smart Licenses** page provides an overview of license usage on the Firepower Management Center, as described below.

Usage Authorization

Possible status values are:

- **In-compliance** () — All licenses assigned to managed devices are in compliance and the Firepower Management Center is communicating successfully with the Cisco licensing authority.
- **License is in compliance but communication with licensing authority has failed**— Device licenses are in compliance, but the Firepower Management Center is not able to communicate with the Cisco licensing authority.
- **Out-of-compliance icon or unable to communicate with License Authority**— One or more managed devices is using a license that is out of compliance, or the Firepower Management Center has not communicated with the Cisco licensing authority in more than 90 days.

Product Registration

Specifies the last date when the Firepower Management Center contacted the License Authority and registered.

Assigned Virtual Account

Specifies the Virtual Account under the Smart Account that you used to generate the Product Instance Registration Token and register the Firepower Management Center. If this deployment is not associated with a particular virtual account within your Smart Account, this information is not displayed.

Export-Controlled Features

If this option is enabled, you can deploy restricted features. For details, see Licensing for Export-Controlled Functionality, on page 112.

Move or Remove Licenses from FTD Devices

Use this procedure to manage licenses for Firepower Threat Defense devices managed by an Firepower Management Center.

For example, you can move a license from one FTD device to another device registered to the same FMC, or to remove a license from a device.

If you remove (disable) a license for a device, you cannot use the features associated with that license on that device.



Important If you need to move a license to a device managed by a *different* Firepower Management Center, see Transfer FTD Licenses to a Different Firepower Management Center, on page 122.

Procedure

Step 1 Choose **System** > **Licenses** > **Smart Licenses**.

Step 2 Click **Edit Licenses**.

Step 3 Click either the **Malware**, **Threat**, or **URL Filtering**.

Step 4 Choose the devices you want to license, then click **Add**, and/or click each device form which you want to remove a license and click the **Delete** ().

Step 5 Click **Apply**.

What to do next

Deploy the changes to the managed devices.

Transfer FTD Licenses to a Different Firepower Management Center

When you register a Smart License to a Firepower Management Center, your virtual account allocates the license to the FMC. If you need to transfer your Smart Licenses to another Firepower Management Center, you must deregister the currently licensed FMC. This removes it from your virtual account and frees your existing licenses, so you can register the licenses to the new FMC. Otherwise, you cannot reuse these licenses, and you may receive an Out-of-Compliance notification because your virtual account does not have enough free licenses. For instructions, see Deregister a Firepower Management Center from the Cisco Smart Software Manager, on page 122.

Then you can register the licenses to the destination Firepower Management Center.

If FTD License Status is Out of Compliance

If the Usage Authorization status on the Smart Licenses page (**System** > **Licenses** > **Smart Licenses**) shows Out of Compliance, you must take action.

Procedure

Step 1 Look at the Smart Licenses section at the bottom of the page to determine which licenses are needed.

Step 2 Purchase the required licenses through your usual channels.

Step 3 In Cisco Smart Software Manager (https://software.cisco.com/#SmartLicensing-Inventory), verify that the licenses appear in your virtual account.

If the expected licenses are not present, see Troubleshoot FTD Licensing, on page 123.

Step 4 In Firepower Management Center, select **System** > **Licenses** > **Smart Licenses**.

Step 5 Click **Re-Authorize**.

Deregister a Firepower Management Center from the Cisco Smart Software Manager

Deregister (unregister) your Firepower Management Center from the Cisco Smart Software Manager before you reinstall (reimage) the appliance, or if you need to release all of the license entitlements back to your Smart Account for any reason.

Deregistering removes the FMC from your virtual account. All license entitlements associated with the Firepower Management Center release back to your virtual account. After deregistration, the Firepower Management Center enters Enforcement mode where no update or changes on licensed features are allowed.

If you need to remove the licenses from a subset of managed Firepower Threat Defense devices, see Assign Licenses to Multiple Managed Devices, on page 119 or Assign Licenses to Managed Devices from the Device Management Page, on page 134.

Procedure

Step 1 Choose **System** > **Licenses** > **Smart Licenses**.

Step 2 Click **Deregister** ( ).

Synchronize a Firepower Management Center with the Cisco Smart Software Manager

If you make changes in the Cisco Smart Software Manager, you can refresh the authorization on the Firepower Management Center so the changes immediately take effect.

Procedure

Step 1 Choose **System** > **Licenses** > **Smart Licenses**.

Step 2 Click **Refresh** ( ).

Troubleshoot FTD Licensing

Expected Licenses Do Not Appear in My Smart Account

If the licenses you expect to see are not in your Smart Account, try the following:

- Make sure they are not in a different Virtual Account. Your organization's license administrator may need to assist you with this.
- Check with the person who sold you the licenses to be sure that transfer to your account is complete.

Unable to Connect to Smart License Server

Check the obvious causes first. For example, make sure your Firepower system has outside connectivity. See Internet Access Requirements, on page 2426.

Unexpected Out-of-Compliance Notification or Other Error

- If a device is already registered to a different FMC, you need to deregister the original FMC before you can license the device under a new FMC. See Deregister a Firepower Management Center from the Cisco Smart Software Manager, on page 122.
- Try synchronizing: Synchronize a Firepower Management Center with the Cisco Smart Software Manager, on page 123.

Troubleshoot Other Issues

For solutions to other common issues, see https://www.cisco.com/c/en/us/support/docs/security/firepower-management-center/215838-fmc-and-ftd-smart-license-registration-a.html

License Classic Devices (Firepower 7000/8000 Series, ASA FirePOWER, and NGIPSv)

7000 and 8000 Series and NGIPSv devices and ASA FirePOWER modules require Classic licenses. These devices are frequently referred to in this documentation as Classic devices.



Important If you are running Firepower hardware but not Firepower software, see licensing information for the software product you are using. This documentation is not applicable.

Classic licenses require a product authorization key (PAK) to activate and are device-specific. Classic licensing is sometimes also referred to as "traditional licensing."

Product License Registration Portal

When you purchase one or more Classic licenses for Firepower features, you manage them in the Cisco Product License Registration Portal:

https://cisco.com/go/license

For more information on using this portal, see:

https://slexui.cloudapps.cisco.com/SWIFT/LicensingUI/Quickstart

You will need your account credentials in order to access these links.

Service Subscriptions for Firepower Features (Classic Licensing)

Some features require a service subscription.

A service subscription enables a specific Firepower feature on a managed device for a set length of time. Service subscriptions can be purchased in one-, three-, or five-year terms. If a subscription expires, Cisco notifies you that you must renew the subscription. If a subscription expires for a Classic device, you might not be able to use the related features, depending on the feature type.

Table 15: Service Subscriptions and Corresponding Classic Licenses

Subscription You Purchase	Classic Licenses You Assign in Firepower System
TA	Control + Protection (a.k.a. "Threat & Apps," required for system updates)
TAC	Control + Protection + URL Filtering
TAM	Control + Protection + Malware
TAMC	Control + Protection + URL Filtering + Malware
URL	URL Filtering (add-on where TA is already present)
AMP	Malware (add-on where TA is already present)

Your purchase of a managed device that uses Classic licenses automatically includes Control and Protection licenses. These licenses are perpetual, but you must also purchase a TA service subscription to enable system updates. Service subscriptions for additional features are optional.

Classic License Types and Restrictions

This section describes the types of Classic licenses available in a Firepower System deployment. The licenses you can enable on a device depend on its model, version, and the other licenses enabled.

Licenses are model-specific for 7000 and 8000 Series and NGIPSv devices and for ASA FirePOWER modules. You cannot enable a license on a managed device unless the license exactly matches the device's model. For example, you cannot use a Firepower 8250 Malware license (FP8250-TAM-LIC=) to enable Malware capabilities on an 8140 device; you must purchase a Firepower 8140 Malware license (FP8140-TAM-LIC=).



Note For NGIPSv or ASA FirePOWER, the Control license allows you to perform user and application control, but these devices do not support switching, routing, stacking, or 7000 and 8000 Series device high availability.

There are a few ways you may lose access to licensed features in the Firepower System:

- You can remove Classic licenses from the Firepower Management Center, which affects all of its managed devices.
- You can disable licensed capabilities on specific managed devices.

Though there are some exceptions, you cannot use the features associated with an expired or deleted license.

The following table summarizes Classic licenses in the Firepower System.

Table 16: Firepower System Classic Licenses

License You Assign in Firepower System	Service Subscription You Purchase	Platforms	Granted Capabilities	Also Requires	Expire Capable?
Any	TA, TAC, TAM, or TAMC	7000 and 8000 Series ASA FirePOWER NGIPSv	host, application, and user discovery decrypting and inspecting SSL- and TLS-encrypted traffic	none	depends on license
Protection	TA (included with device)	7000 and 8000 Series ASA FirePOWER NGIPSv	intrusion detection and prevention file control Security Intelligence filtering	none	no
Control	none (included with device)	7000 and 8000 Series	user and application control switching and routing 7000 and 8000 Series device high availability 7000 and 8000 Series network address translation (NAT)	Protection	no
Control	none (included with device)	ASA FirePOWER NGIPSv	user and application control	Protection	no
Malware	TAM, TAMC, or AMP	7000 and 8000 Series ASA FirePOWER NGIPSv	AMP for Networks (network-based Advanced Malware Protection) File storage	Protection	yes
URL Filtering	TAC, TAMC, or URL	7000 and 8000 Series ASA FirePOWER NGIPSv	category and reputation-based URL filtering	Protection	yes
VPN	none (contact Sales for more information)	7000 and 8000 Series	deploying virtual private networks	Control	yes

Protection Licenses

A Protection license allows you to perform intrusion detection and prevention, file control, and Security Intelligence filtering:

- *Intrusion detection and prevention* allows you to analyze network traffic for intrusions and exploits and, optionally, drop offending packets.
- *File control* allows you to detect and, optionally, block users from uploading (sending) or downloading (receiving) files of specific types over specific application protocols. *AMP for Networks*, which requires a Malware license, allows you to inspect and block a restricted set of those file types based on their dispositions.
- *Security Intelligence filtering* allows you to block —deny traffic to and from—specific IP addresses, URLs, and DNS domain names, before the traffic is subjected to analysis by access control rules. Dynamic feeds allow you to immediately block connections based on the latest intelligence. Optionally, you can use a "monitor-only" setting for Security Intelligence filtering.

A Protection license (along with a Control license) is automatically included in the purchase of any Classic managed device. This license is perpetual, but you must also purchase a TA subscription to enable system updates.

Although you can configure an access control policy to perform Protection-related inspection without a license, you cannot deploy the policy until you first add a Protection license to the Firepower Management Center, then enable it on the devices targeted by the policy.

If you delete your Protection license from the Firepower Management Center or disable Protection on managed devices, the Firepower Management Center stops acknowledging intrusion and file events from the affected devices. As a consequence, correlation rules that use those events as a trigger criteria stop firing. Additionally, the Firepower Management Center will not contact the internet for either Cisco-provided or third-party Security Intelligence information. You cannot re-deploy existing policies until you re-enable Protection.

Because a Protection license is required for URL Filtering, Malware, and Control licenses, deleting or disabling a Protection license has the same effect as deleting or disabling your URL Filtering, Malware, or Control license.

Control Licenses

A Control license allows you to implement user and application control by adding user and application conditions to access control rules. For 7000 and 8000 Series devices only, this license also allows you to configure switching and routing (including DHCP relay and NAT) and device high-availability pairs. To enable a Control license on a managed device, you must also enable a Protection license. A Control license is automatically included (along with a Protection license) in the purchase of any Classic managed device. This license is perpetual, but you must also purchase a TA subscription to enable system updates.

If you do not enable a Control license for a Classic managed device, you can add user and application conditions to rules in an access control policy, but you cannot deploy the policy to the device. If you do not enable a Control license for 7000 or 8000 Series devices specifically, you also cannot:

- create switched, routed, or hybrid interfaces
- create NAT entries
- configure DHCP relay for virtual routers
- deploy a device configuration that includes switch or routing to the device
- establish high availability between devices

Note Although you can create virtual switches and routers without a Control license, they are not useful without switched and routed interfaces to populate them.

If you delete a Control license from the Firepower Management Center or disable Control on individual devices:

- The affected devices do **not** stop performing switching or routing, nor do device high-availability pairs break.
- You can continue to edit and delete existing configurations, but you cannot deploy those changes to the affected devices.
- You cannot add new switched, routed, or hybrid interfaces, nor can you add new NAT entries, configure DHCP relay, or establish 7000 or 8000 Series device high-availability.
- You cannot re-deploy existing access control policies if they include rules with user or application conditions.

URL Filtering Licenses for Classic Devices

URL filtering allows you to write access control rules that determine the traffic that can traverse your network based on URLs requested by monitored hosts, correlated with information about those URLs. To enable a URL Filtering license, you must also enable a Protection license. You can purchase a URL Filtering license for Classic devices as a services subscription combined with Threat & Apps (TAC) or Threat & Apps and Malware (TAMC) subscriptions, or as an add-on subscription (URL) for a system where Threat & Apps (TA) is already enabled.

Tip Without a URL Filtering license, you can specify individual URLs or groups of URLs to allow or block. This gives you granular, custom control over web traffic, but does not allow you to use URL category and reputation data to filter network traffic.

Although you can add category and reputation-based URL conditions to access control rules without a URL Filtering license, the Firepower Management Center will not download URL information. You cannot deploy the access control policy until you first add a URL Filtering license to the Firepower Management Center, then enable it on the devices targeted by the policy.

You may lose access to URL filtering if you delete the license from the Firepower Management Center or disable URL Filtering on managed devices. Also, URL Filtering licenses may expire. If your license expires or if you delete or disable it, access control rules with URL conditions immediately stop filtering URLs, and your Firepower Management Center can no longer download updates to URL data. You cannot re-deploy existing access control policies if they include rules with category and reputation-based URL conditions.

Malware Licenses for Classic Devices

A Malware license allows you to perform Cisco Advanced Malware Protection (AMP) with AMP for Networks and Cisco Threat Grid. You can use managed devices to detect and block malware in files transmitted over your network. To enable a Malware license, you must also enable Protection. You can purchase a Malware license as a subscription combined with Threat & Apps (TAM) or Threat & Apps and URL Filtering (TAMC) subscriptions, or as an add-on subscription (AMP) for a system where Threat & Apps (TA) is already enabled.

Note 7000 and 8000 Series managed devices with Malware licenses enabled attempt to connect periodically to the AMP cloud even if you have not configured dynamic analysis. Because of this, the device's Interface Traffic dashboard widget shows transmitted traffic; this is expected behavior.

You configure AMP for Networks as part of a file policy, which you then associate with one or more access control rules. File policies can detect your users uploading or downloading files of specific types over specific application protocols. AMP for Networks allows you to use local malware analysis and file preclassification to inspect a restricted set of those file types for malware. You can also download and submit specific file types to the Cisco Threat Grid cloud for dynamic and Spero analysis to determine whether they contain malware. For these files, you can view the network file trajectory, which details the path the file has taken through your network. The Malware license also allows you to add specific files to a file list and enable the file list within a file policy, allowing those files to be automatically allowed or blocked on detection.

Before you can deploy an access control policy that includes AMP for Networks configurations, you **must** add a Malware license, then enable it on the devices targeted by the policy. If you later disable the license on the devices, you cannot re-deploy the existing access control policy to those devices.

If you delete all your Malware licenses or they all expire, the system stops querying the AMP cloud, and also stops acknowledging retrospective events sent from the AMP cloud. You cannot re-deploy existing access control policies if they include AMP for Networks configurations. Note that for a very brief time after a Malware license expires or is deleted, the system can use existing cached file dispositions. After the time window expires, the system assigns a disposition of `Unavailable` to those files.

A Malware license is required only if you deploy AMP for Networks and Cisco Threat Grid. Without a Malware license, the Firepower Management Center can receive AMP for Endpoints malware events and indications of compromise (IOC) from the AMP cloud.

See also important information at License Requirements for File and Malware Policies, on page 1403.

VPN Licenses for 7000 and 8000 Series Devices

VPN allows you to establish secure tunnels between endpoints via a public source, such as the Internet or other network. You can configure the Firepower System to build secure VPN tunnels between the virtual routers of 7000 and 8000 Series devices. To enable VPN, you must also enable Protection and Control licenses. To purchase a VPN license, contact Sales.

Without a VPN license, you cannot configure a VPN deployment with your 7000 and 8000 Series devices. Although you can create deployments, they are not useful without at least one VPN-enabled routed interface to populate them.

If you delete your VPN license from the Firepower Management Center or disable VPN on individual devices, the affected devices do **not** break the current VPN deployments. Although you can edit and delete existing deployments, you cannot deploy your changes to the affected devices.

Classic Licenses in Device Stacks and High-Availability Pairs

Individual devices must have equivalent licenses before they can be stacked or configured into 7000 or 8000 Series device high-availability pairs. After you stack devices, you can change the licenses for the entire stack. However, you cannot change the enabled licenses on a 7000 or 8000 Series device high-availability pair.

See also About Device Stacks, on page 533 and Device High Availability Requirements, on page 518.

View Your Classic Licenses

Procedure

Do one of the following, depending on your needs:

To View	Do This
The Classic licenses that you have added to the Firepower Management Center and details including their type, status, usage, expiration dates, and the managed devices to which they are applied.	Choose **System** > **Licenses** > **Classic Licenses**. The summary shows the number of licenses you have purchased, followed by the number of licenses that are in used in parentheses.
The licenses applied to each of your managed devices	Choose **Devices** > **Device Management**.
License status in the Health Monitor	Use the Classic License Monitor health module in a health policy. For information, see Health Monitoring, on page 275, including #unique_230 and Creating Health Policies, on page 283.
An overview of your licenses in the Dashboard	Add the Product Licensing widget to the dashboard of your choice. For instructions, see The Product Licensing Widget, on page 266 and Adding Widgets to a Dashboard, on page 269 and Dashboard Widget Availability by User Role, on page 257.

Identify the License Key

The license key uniquely identifies the Firepower Management Center in the Cisco License Registration Portal. It is composed of a product code (for example, 66) and the MAC address of the management port (eth0) of the Firepower Management Center; for example, 66:00:00:77:FF:CC:88.

You will use the license key in the Cisco License Registration Portal to obtain the license text required to add licenses to the Firepower Management Center.

Procedure

Step 1 Choose **System** > **Licenses** > **Classic Licenses**.

Step 2 Click **Add New License**.

Step 3 Note the value in the **License Key** field at the top of the **Add Feature License** dialog.

What to do next

- Add a license to the Firepower Management Center; see Generate a Classic License and Add It to the Firepower Management Center, on page 131.

  This procedure includes the process of generating the actual license text using the license key.

Generate a Classic License and Add It to the Firepower Management Center

Note If you add licenses after a backup has completed, these licenses will not be removed or overwritten if this backup is restored. To prevent a conflict on restore, remove those licenses before restoring the backup, noting where the licenses were used, and add and reconfigure them after restoring the backup. If a conflict occurs, contact Support.

Tip You can also request licenses on the **Licenses** tab after you log into the Support Site.

Before you begin

- Make sure you have the product activation key (PAK) from the Software Claim Certificate that Cisco provided when you purchased the license. If you have a legacy, pre-Cisco license, contact Support.
- Identify the license key for the Firepower Management Center; see Identify the License Key, on page 130.
- You will need your account credentials to complete this procedure.

Procedure

Step 1 Choose **System** > **Licenses** > **Classic Licenses**.

Step 2 Click **Add New License**.

Step 3 Continue as appropriate:

- If you have already obtained the license text, skip to Step 8.
- If you still need to obtain the license text, go to the next step.

Step 4 Click **Get License** to open the Cisco License Registration Portal.

Note If you cannot access the Internet using your current computer, switch to a computer that can, and browse to http://cisco.com/go/license.

Step 5 Generate a license from the PAK in the License Registration Portal.

This step requires the PAK you received during the purchase process, as well as the license key for the Firepower Management Center.

For information, see Product License Registration Portal, on page 124.

Step 6 Copy the license text from either the License Registration Portal display, or the email the License Registration Portal sends you.

Important The licensing text block in the portal or email message may include more than one license. Each license is bounded by a BEGIN LICENSE line and an END LICENSE line. Make sure that you copy and paste only one license at a time.

Step 7 Return to the **Add Feature License** page in the Firepower Management Center's web interface.

Step 8 Paste the license text into the **License** field.

Step 9 Click **Verify License**.

If the license is invalid, make sure that you correctly copied the license text.

Step 10 Click **Submit License**.

What to do next

- Assign the license to a managed device; see Assign Licenses to Managed Devices from the Device Management Page, on page 134. You **must** assign licenses to your managed devices before you can use licensed features on those devices.

How to Convert a Classic License for Use on an FTD Device

You can convert licenses using either the License Registration Portal (LRP) or the Cisco Smart Software Manager (CSSM), and you can convert an unused Product Authorization Key (PAK) or a Classic license that has already been assigned to a device.



Important You cannot undo this process. You cannot convert a Smart License to a Classic license, even if the license was originally a Classic license.

In documentation on Cisco.com, Classic licenses may also be referred to as "traditional" licenses.

Before you begin

- It is easiest to convert a Classic license to a Smart License when it is still an unused PAK that has not yet been assigned to a product instance.
- Your hardware must be able to run Firepower Threat Defense. See the *Cisco Firepower Compatibility Guide* at https://www.cisco.com/c/en/us/support/security/defense-center/products-device-support-tables-list.html.
- You must have a Smart Account. If you do not have one, create one. See Create a Smart Account to Hold Your Licenses, on page 113.
- The PAKs or licenses that you want to convert must appear in your Smart Account.
- If you convert using the License Registration Portal instead of the Cisco Smart Software Manager, you must have your Smart Account credentials in order to initiate the conversion process.

Procedure

Step 1 The conversion process you will follow depends on whether or not the license has been consumed:

- If the PAK that you want to convert has never been used, follow instructions for converting a PAK.
- If the PAK you want to convert has already been assigned to a device, follow instructions for converting a Classic license.

  Make sure your existing classic license is still registered to your device.

Step 2 See instructions for your type of conversion (PAK or installed Classic license) in the following documentation:

- To convert PAKs or licenses using the LRP:
  - To view a video that steps you through the License Registration Portal part of the conversion process, click https://salesconnect.cisco.com/#/content-detail/7da52358-0fc1-4d85-8920-14a1b7721780.
  - Search for "Convert" in the following document: https://cisco.app.box.com/s/mds3ab3fctk6pzonq5meukvcpjizt7wu.

    There are three conversion procedures. Choose the conversion procedure applicable to your situation.
  - Sign in to the License Registration Portal (LRP) at https://tools.cisco.com/SWIFT/LicensingUI/Home and follow the instructions in the documentation above.
- To convert PAKs or licenses using CSSM:
  - *Converting Hybrid Licenses to Smart Software Licenses QRG*:

    https://community.cisco.com/t5/licensing-enterprise-agreements/converting-hybrid-licenses-to-smart-software-licenses-qrg/ta-p/3628609?attachment-id=134907
  - Sign in to CSSM at https://software.cisco.com/#SmartLicensing-LicenseConversion and follow the instructions for your type of conversion (PAK or installed Classic license) in the documentation above.

Step 3 Freshly install Firepower Threat Defense on your hardware.

See the instructions for your hardware at https://www.cisco.com/c/en/us/support/security/firepower-ngfw/products-installation-guides-list.html.

Step 4 If you will use Firepower Device Manager to manage this device as a standalone device:

See information about licensing the device in the *Cisco Firepower Threat Defense Configuration Guide for Firepower Device Manager* at https://www.cisco.com/c/en/us/support/security/firepower-ngfw/products-installation-and-configuration-guides-list.html.

Skip the rest of this procedure.

Step 5 If you have already deployed Smart Licensing on your Firepower Management Center:

a) Set up Smart Licensing on your new Firepower Threat Defense device.

   See Assign Licenses to Multiple Managed Devices, on page 119.

b) Verify that the new Smart License has been successfully applied to the device.

See View FTD Licenses and License Status, on page 120.

Step 6 If you have not yet deployed Smart Licensing on your Firepower Management Center:

See Register Smart Licenses, on page 116.

Assign Licenses to Managed Devices from the Device Management Page

Although there are some exceptions, you cannot use the features associated with a license if you disable it on a managed device.

Note For an Firepower Threat Defense cluster, you apply the licenses to the cluster as a whole; note that each unit in the cluster consumes a separate license per feature.

Before you begin

- Add your devices to the Firepower Management Center. See Add a Device to the FMC, on page 229.
- You must have Admin or Network Admin privileges to perform this task. When operating with multiple domains, you must do this task in leaf domains.
- If you will assign Smart Licenses:
  - If you need to apply Smart Licenses to many devices at one time, use the Smart Licenses page instead of following this procedure. See Assign Licenses to Multiple Managed Devices, on page 119
  - Prepare Smart Licenses for distribution to managed devices: See Register Smart Licenses, on page 116

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to assign or disable a license, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Device**.

Step 4 Next to the License section, click **Edit** ( ).

Step 5 Check or clear the appropriate check boxes to assign or disable licenses for the device.

Step 6 Click **Save**.

What to do next

- If you assigned Smart Licenses, verify license status:

  Go to **System** > **Licenses** > **Smart Licenses**, enter the hostname or IP address of the device into the filter at the top of the Smart Licenses table, and verify that only a green circle with a **Check Mark** ( ) appears for each device, for each license type. If you see any other icon, hover over the icon for more information.

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

- If you are licensing Firepower Threat Defense devices and you applied a Base license with export-controlled functionality enabled, reboot each device. For devices configured in a high-availability pair, reboot both devices at the same time to avoid an Active-Active condition.

Additional Information about Firepower Licensing

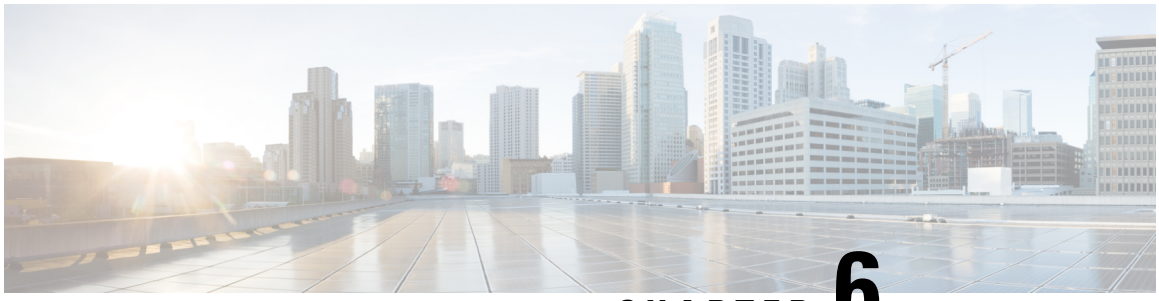
For additional information to help resolve common licensing questions, see the following documents:

- The *Frequently Asked Questions (FAQ) about Firepower Licensing* document at:

  https://www.cisco.com/c/en/us/td/docs/security/firepower/licensing/faq/firepower-license-FAQ.html

- The *Cisco Firepower System Feature Licenses* document at:

  https://www.cisco.com/c/en/us/td/docs/security/firepower/roadmap/firepower-licenseroadmap.html



CHAPTER 6

System Updates

The following topics explain how to update Firepower deployments:

- About System Updates, on page 137
- Requirements and Prerequisites for System Updates, on page 138
- Guidelines and Limitations for System Updates, on page 139
- Upgrade System Software, on page 139
- Update the Vulnerability Database (VDB), on page 139
- Update the Geolocation Database, on page 141
- Update Intrusion Rules, on page 143
- Maintain Your Air-Gapped Deployment, on page 153

About System Updates

You can use the FMC to upgrade the system software for itself and the devices it manages. You can also update various databases and feeds that provide advanced services.

For FMCs with internet access, the system can often obtain updates directly from Cisco. We recommend you schedule or enable automatic updates whenever possible. Some updates are auto-enabled when you enable the related feature. Other updates you must schedule yourself. After initial setup, we recommend you review all auto-updates and adjust them if necessary.

Table 17: Upgrades and Updates in FMC Deployments

Component	Description	Details
Firepower software	*Major* software releases contain new features, functionality, and enhancements. They may include infrastructure or architectural changes. *Patches* are on-demand updates limited to critical fixes with time urgency. *Hotfixes* can address specific customer issues.	**Direct Download:** Select releases only, usually some time after the release is available for manual download. The length of the delay depends on release type, release adoption, and other factors. **Schedule:** Patches only, on **System** > **Tools** > **Scheduling**. **Uninstall:** Patches only. **Reimage:** Major releases only. **See:** Upgrade System Software, on page 139

Component	Description	Details
Vulnerability database (VDB)	The Cisco vulnerability database (VDB) is a database of known vulnerabilities to which hosts may be susceptible, as well as fingerprints for operating systems, clients, and applications. The system uses the VDB to help determine whether a particular host increases your risk of compromise.	**Direct Download:** Yes. **Schedule:** Yes, on **System** > **Tools** > **Scheduling**. **Uninstall:** No. **See:** Update the Vulnerability Database (VDB), on page 139
Geolocation database (GeoDB)	The Cisco geolocation database (GeoDB) is a database of geographical and connection-related data associated with routable IP addresses.	**Direct Download:** Yes. **Schedule:** Yes, on **System** > **Updates**. **Uninstall:** No. **See:** Update the Geolocation Database, on page 141
Intrusion rules (SRU)	Intrusion rule updates provide new and updated intrusion rules and preprocessor rules, modified states for existing rules, and modified default intrusion policy settings. Rule updates may also delete rules, provide new rule categories and default variables, and modify default variable values.	**Direct Download:** Yes. **Schedule:** Yes, on **System** > **Updates**. **Uninstall:** No. **See:** Update Intrusion Rules, on page 143
Security Intelligence feeds	Security Intelligence feeds are collections of IP addresses, domain names, and URLs that you can use to quickly filter traffic that matches an entry.	**Direct Download:** Yes. **Schedule:** Yes, on **Objects** > **Object Management**. **Uninstall:** No. **See:** List and Feed Updates for Security Intelligence, on page 432
URL categories and reputations	URL filtering allows you to control access to websites based on the URL's general classification (category) and risk level (reputation).	**Direct Download:** Yes. **Schedule:** Yes, on **System** > **Integration** > **Cloud Services** *or* **System** > **Tools** > **Scheduling**, depending on your requirements. **Uninstall:** No. **See:** Enable URL Filtering Using Category and Reputation, on page 1235

Requirements and Prerequisites for System Updates

Model Support

Any

Supported Domains

Global unless indicated otherwise.

User Roles

Admin

Guidelines and Limitations for System Updates

Before You Update

Before you update any component of your Firepower deployment (including intrusion rules, VDB, or GeoDB) read the release notes or advisory text that accompanies the update. These provide critical and release-specific information, including compatibility, prerequisites, new capabilities, behavior changes, and warnings.

Scheduled Updates

The system schedules tasks — including updates — in UTC. This means that when they occur locally depends on the date and your specific location. Also, because updates are scheduled in UTC, they do not adjust for Daylight Saving Time, summer time, or any such seasonal adjustments that you may observe in your location. If you are affected, scheduled updates occur one hour "later" in the summer than in the winter, according to local time.



Important We *strongly* recommend you review scheduled updates to be sure they occur when you intend.

Bandwidth Guidelines

To upgrade a Firepower appliance (or perform a readiness check), the upgrade package must be on the appliance. Firepower upgrade package sizes vary. Make sure you have the bandwidth to perform a large data transfer to your managed devices. See Guidelines for Downloading Data from the Firepower Management Center to Managed Devices (Troubleshooting TechNote).

Upgrade System Software

This guide does not contain detailed upgrade instructions for either system software or companion operating systems. Instead, see the Cisco Firepower Management Center Upgrade Guide, Version 6.0–7.0.

For information on scheduling downloads and installations for system software patches, see Software Update Automation, on page 188.

Update the Vulnerability Database (VDB)

The Cisco vulnerability database (VDB) is a database of known vulnerabilities to which hosts may be susceptible, as well as fingerprints for operating systems, clients, and applications. The system uses the VDB to help determine whether a particular host increases your risk of compromise.

Cisco issues periodic updates to the VDB. The time it takes to update the VDB and its associated mappings on the FMC depends on the number of hosts in your network map. As a rule of thumb, divide the number of hosts by 1000 to determine the approximate number of minutes to perform the update.

If the FMC has internet access, we recommend you schedule tasks to perform automatic recurring VDB update downloads and installations.

Caution Updating the VDB *immediately* restarts the Snort process on all managed devices. Additionally, the first deploy after installing the VDB *might* cause a Snort restart depending on the VDB content. Snort restarts cause an interruption in traffic inspection and, depending on how the managed device handles traffic, possibly interrupts traffic flow. For more information, see Snort® Restart Traffic Behavior, on page 350.

Manually Update the VDB

To update the VDB, the VDB update package must be on the FMC.

Before you begin

- Download the update from https://www.cisco.com/go/firepower-software.

Note Beginning with VDB Release 343, all application detector information is available through Cisco Secure Firewall Application Detectors. This site includes a searchable database of application detectors. The Release Notes provide information on changes for a particular VDB release.

- Consider the update's effect on traffic flow and inspection due to Snort restarts. We recommend performing updates in a maintenance window.

Procedure

Step 1 Choose **System** > **Updates**, then click **Product Updates**.

Step 2 Choose how you want to upload the VDB update to the FMC.

- Download directly from Cisco.com: Click **Download Updates**. If it can access the Cisco Support & Download site, the Firepower Management Center downloads the latest VDB. Note that the Firepower Management Center also downloads a package for each patch and hotfix (but not major release) associated with the version your appliances are currently running.
- Upload manually: Click **Upload Update**, then **Choose File**. Browse to the update you downloaded earlier, and click **Upload**.

VDB updates appear on the same page as Firepower software upgrade and uninstaller packages.

Step 3 Install the update.

a) Click Install next to the Vulnerability and Fingerprint Database update.
b) Choose the Firepower Management Center.
c) Click **Install**.

Step 4 (Optional) Monitor update progress in the Message Center.

Do not perform tasks related to mapped vulnerabilities until the update completes. Even if the Message Center shows no progress for several minutes or indicates that the update has failed, do not restart the update. Instead, contact Cisco TAC.

After the update completes and Snort restarts, the system uses the new vulnerability information. However, you must deploy before updated application detectors and operating system fingerprints can take effect.

Step 5 Verify update success.

Choose **Help** > **About** to view the current VDB version.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Schedule VDB Updates

If your FMC has internet access, we recommend you schedule regular VDB updates. See Vulnerability Database Update Automation, on page 191.

Update the Geolocation Database

The geolocation database (GeoDB) is a database that you can leverage to view and filter traffic based on geographical location.

The system comes with an initial GeoDB that maps IP addresses to countries/continents, so that information should always be available. If you update the GeoDB, the system also downloads contextual data. This contextual data includes additional location details, as well as connection information such as ISP, connection type, proxy type, domain name, and so on. We issue periodic updates to the GeoDB. You must regularly update the GeoDB to have accurate geolocation information.

The time needed to update the GeoDB depends on your appliance, but can take up to 45 minutes depending on the size of the update—for example, if this is the first time you are downloading the full GeoDB. Although a GeoDB update does not interrupt any other system functions (including the ongoing collection of geolocation information), the update does consume system resources while it completes. Consider this when planning your updates.

The GeoDB update overrides any previous versions of the GeoDB and is effective immediately. When you update the GeoDB, the FMC automatically updates the related data on its managed devices. It may take a few minutes for a GeoDB update to take effect throughout your deployment. You do not need to re-deploy after you update.

The **System** > **Updates** > **Geolocation Updates** page and the **Help** > **About** page both list the current version.

Note In May 2022 we split the GeoDB into two packages: a country code package that maps IP addresses to countries/continents, and an IP package that contains contextual data. The new country code package has the same file name as the old all-in-one package. This allows FMCs running Version 7.1 and earlier to continue to obtain GeoDB updates. However, because this package now contains only country code mappings, the contextual data is no longer updated and will grow stale. To obtain fresh data, upgrade or reimage to Version 7.2+ and update the GeoDB. Note that this split does not affect geolocation rules or traffic handling in any way—those rules rely only on the data in the country code package.

Manually Update the GeoDB (Internet Connection)

You can import a new GeoDB update by automatically connecting to the Support Site only if the appliance has Internet access.

Procedure

Step 1 Choose **System** > **Updates**.

Step 2 Click **Geolocation Updates**.

Step 3 Choose **Download and install geolocation update from the Support Site**.

Step 4 Click **Import**.
The system queues a Geolocation Update task, which checks for the latest updates on the Cisco Support Site (http://www.cisco.com/cisco/web/support/index.html).

Step 5 Optionally, monitor the task status; see Viewing Task Messages, on page 316.

Step 6 After the update finishes, return to the Geolocation Updates page or choose **Help** > **About** to confirm that the GeoDB build number matches the update you installed.

Manually Update the GeoDB (No Internet Connection)

Use this procedure to perform an on-demand update of the GeoDB if the FMC does not have internet access.

Procedure

Step 1 Download the GeoDB from the Cisco Support & Download site: https://www.cisco.com/go/firepower-software.

Select or search for your model (or choose any model—you use the same GeoDB for all FMCs), then browse to the *Coverage and Content Updates* page.

Make sure you download the country code package: Cisco_GEODB_Update-*date-build*. The IP package is for Version 7.2+.

Step 2 Choose **System** > **Updates** > **Geolocation Updates**.

Step 3 Under One-Time Geolocation Update, choose **Upload and install geolocation update**.

Step 4 Click **Choose File**, then browse to the country code package you downloaded earlier.

Step 5 Click **Import**.

You can monitor update progress in the Message Center.

Step 6 Verify update success.

The Geolocation Updates page and the **Help** > **About** page both list the current version.

Schedule GeoDB Updates

If the Firepower Management Center has internet access, we recommend you schedule weekly GeoDB updates.

Before you begin

Make sure the FMC can access the internet.

Procedure

Step 1 Choose **System** > **Updates**, then click **Geolocation Updates**.

Step 2 Under **Recurring Geolocation Updates**, check **Enable Recurring Weekly Updates...**

Step 3 Specify the **Update Start Time**.

Step 4 Click **Save**.

Update Intrusion Rules

As new vulnerabilities become known, the Cisco Talos Intelligence Group (Talos) releases intrusion rule updates that you can import onto your Firepower Management Center, and then implement by deploying the changed configuration to your managed devices. These updates affect intrusion rules, preprocessor rules, and the policies that use the rules.

Intrusion rule updates are cumulative, and Cisco recommends you always import the latest update. You cannot import an intrusion rule update that either matches or predates the version of the currently installed rules.

An intrusion rule update may provide the following:

- **New and modified rules and rule states**—Rule updates provide new and updated intrusion and preprocessor rules. For new rules, the rule state may be different in each system-provided intrusion policy. For example, a new rule may be enabled in the Security over Connectivity intrusion policy and disabled in the Connectivity over Security intrusion policy. Rule updates may also change the default state of existing rules, or delete existing rules entirely.
- **New rule categories**—Rule updates may include new rule categories, which are always added.
- **Modified preprocessor and advanced settings**—Rule updates may change the advanced settings in the system-provided intrusion policies and the preprocessor settings in system-provided network analysis policies. They can also update default values for the advanced preprocessing and performance options in your access control policies.

- **New and modified variables**—Rule updates may modify default values for existing default variables, but do not override your changes. New variables are always added.

In a multidomain deployment, you can import local intrusion rules in any domain, but you can import intrusion rule updates from Talos in the Global domain only.

Caution The first deploy after importing an intrusion rule update restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during the interruption or passes without further inspection depends on how the target device handles traffic. For more information, see Snort® Restart Traffic Behavior, on page 350.

Understanding When Intrusion Rule Updates Modify Policies

Intrusion rule updates can affect both system-provided and custom network analysis policies, as well as all access control policies:

- **system provided**—Changes to system-provided network analysis and intrusion policies, as well as any changes to advanced access control settings, automatically take effect when you re-deploy the policies after the update.
- **custom**—Because every custom network analysis and intrusion policy uses a system-provided policy as its base, or as the eventual base in a policy chain, rule updates can affect custom network analysis and intrusion policies. However, you can prevent rule updates from automatically making those changes. This allows you to update system-provided base policies manually, on a schedule independent of rule update imports. Regardless of your choice (implemented on a per-custom-policy basis), updates to system-provided policies do **not** override any settings you customized.

Note that importing a rule update discards all cached changes to network analysis and intrusion policies. For your convenience, the Rule Updates page lists policies with cached changes and the users who made those changes.

Deploying Intrusion Rule Updates

For changes made by an intrusion rule update to take effect, you must redeploy configurations. When importing a rule update, you can configure the system to automatically redeploy to affected devices. This approach is especially useful if you allow the intrusion rule update to modify system-provided base intrusion policies.

Recurring Intrusion Rule Updates

You can import rule updates on a daily, weekly, or monthly basis, using the Rule Updates page.

If your deployment includes a high availability pair of Firepower Management Centers, import the update on the primary only. The secondary Firepower Management Center receives the rule update as part of the regular synchronization process.

Applicable subtasks in the intrusion rule update import occur in the following order: download, install, base policy update, and configuration deploy. When one subtask completes, the next subtask begins.

At the scheduled time, the system installs the rule update and deploys the changed configuration as you specified in the previous step. You can log off or use the web interface to perform other tasks before or during the import. When accessed during an import, the Rule Update Log displays a **Red Status** (), and you can view messages as they occur in the Rule Update Log detailed view. Depending on the rule update size and content, several minutes may pass before status messages appear.

As a part of initial configuration the FMC configures a daily automatic intrusion rule update from the Cisco support site. (The FMC deploys automatic intrusion rule updates to affected managed devices when it next deploys affected policies.) You can observe the status of this update using the web interface Message Center. If configuring the update fails and your FMC has internet access, we recommend you configure regular intrusion rule updates as described in Schedule Intrusion Rule Updates, on page 146.

Importing Local Intrusion Rules

A local intrusion rule is a custom standard text rule that you import from a local machine as a plain text file with ASCII or UTF-8 encoding. You can create local rules using the instructions in the Snort users manual, which is available at http://www.snort.org.

In a multidomain deployment, you can import local intrusion rules in any domain. You can view local intrusion rules imported in the current domain and ancestor domains.

Update Intrusion Rules One-Time Manually



Caution The first deploy after importing an intrusion rule update restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during the interruption or passes without further inspection depends on how the target device handles traffic. For more information, see Snort® Restart Traffic Behavior, on page 350.

Procedure

Step 1 Manually download the update from the Cisco Support Site (http://www.cisco.com/cisco/web/support/index.html).

Step 2 Choose **System** > **Updates**, then click **Rule Updates**.

Step 3 If you want to move all user-defined rules that you have created or imported to the deleted folder, you must click **Delete All Local Rules** in the toolbar, then click **OK**.

Step 4 Choose **Rule Update or text rule file to upload and install** and click **Browse** to navigate to and choose the rule update file.

Step 5 If you want to automatically re-deploy policies to your managed devices after the update completes, choose **Reapply all policies after the rule update import completes**.

Step 6 Click **Import**. The system installs the rule update and displays the Rule Update Log detailed view.

Note Contact Support if you receive an error message while installing the rule update.

Update Intrusion Rules One-Time Automatically

To import a new intrusion rule update automatically, your appliance must have Internet access to connect to the Support Site.

Caution The first deploy after importing an intrusion rule update restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during the interruption or passes without further inspection depends on how the target device handles traffic. For more information, see Snort® Restart Traffic Behavior, on page 350.

Before you begin

- Ensure the Firepower Management Center has internet access; see Security, Internet Access, and Communication Ports, on page 2425.

Procedure

Step 1 Choose **System** > **Updates**.

Tip You can also click **Import Rules** on the Rule Editor page, which you access by choosing **Policies** > **Intrusion** > **Intrusion Rules**.

Step 2 Click **Rule Updates**.

Step 3 If you want to move all user-defined rules that you have created or imported to the deleted folder, click **Delete All Local Rules** in the toolbar, then click **OK**.

Step 4 Choose **Download new Rule Update from the Support Site**.

Step 5 If you want to automatically re-deploy the changed configuration to managed devices after the update completes, check the **Reapply all policies after the rule update import completes** check box.

Step 6 Click **Import**.
The system installs the rule update and displays the Rule Update Log detailed view.

Caution Contact Support if you receive an error message while installing the rule update.

Schedule Intrusion Rule Updates

Caution The first deploy after importing an intrusion rule update restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during the interruption or passes without further inspection depends on how the target device handles traffic. For more information, see Snort® Restart Traffic Behavior, on page 350.

Procedure

Step 1 Choose **System** > **Updates**.

Tip You can also click **Import Rules** on the Rule Editor page, which you access by choosing **Policies** > **Intrusion** > **Intrusion Rules**.

Step 2 Click **Rule Updates**.

Step 3 If you want to move all user-defined rules that you have created or imported to the deleted folder, click **Delete All Local Rules** in the toolbar, then click **OK**.

Step 4 Check **Enable Recurring Rule Update Imports from the Support Site** check box.

Import status messages appear beneath the **Recurring Rule Update Imports** section heading.

Step 5 In the **Import Frequency** field, specify:

- The frequency of the update (**Daily**, **Weekly**, or **Monthly**)
- The day of the week or month you want the update to occur
- The time you want the update to start

Step 6 If you want to automatically re-deploy the changed configuration to your managed devices after the update completes, check the **Deploy updated policies to targeted devices after rule update completes** check box.

Step 7 Click **Save**.

Caution Contact Support if you receive an error message while installing the intrusion rule update.

The status message under the Recurring Rule Update Imports section heading changes to indicate that the rule update has not yet run.

Best Practices for Importing Local Intrusion Rules

Observe the following guidelines when importing a local rule file:

- The rules importer requires that all custom rules are imported in a plain text file encoded in ASCII or UTF-8.
- The text file name can include alphanumeric characters, spaces, and no special characters other than underscore (_), period (.), and dash (-).
- The system imports local rules preceded with a single pound character (#), but they are flagged as deleted.
- The system imports local rules preceded with a single pound character (#), and does not import local rules preceded with two pound characters (##).
- Rules cannot contain any escape characters.
- You do not have to specify a Generator ID (GID) when importing a local rule. If you do, specify only GID 1 for a standard text rule.
- When importing a rule for the first time, do *not* specify a Snort ID (SID) or revision number. This avoids collisions with SIDs of other rules, including deleted rules. The system will automatically assign the rule the next available custom rule SID of 1000000 or greater, and a revision number of 1.

  If you must import rules with SIDs, the SIDs must be unique numbers between 1,000,000 and 9,999,999.

  In a multidomain deployment, the system assigns SIDs to imported rules from a shared pool used by all domains on the Firepower Management Center. If multiple administrators are importing local rules at the same time, SIDs within an individual domain might appear to be non-sequential, because the system assigned the intervening numbers in the sequence to another domain.

- When importing an updated version of a local rule you have previously imported, or when reinstating a local rule you have deleted, you *must* include the SID assigned by the system and a revision number greater than the current revision number. You can determine the revision number for a current or deleted rule by editing the rule.



Note The system automatically increments the revision number when you delete a local rule; this is a device that allows you to reinstate local rules. All deleted local rules are moved from the local rule category to the deleted rule category.

- Import local rules on the primary Firepower Management Center in a high availability pair to avoid SID numbering issues.
- The import fails if a rule contains any of the following: .
  - A SID greater than 2147483647.
  - A list of source or destination ports that is longer than 64 characters.
- Policy validation fails if you enable an imported local rule that uses the deprecated `threshold` keyword in combination with the intrusion event thresholding feature in an intrusion policy.
- All imported local rules are automatically saved in the local rule category.
- The system always sets local rules that you import to the disabled rule state. You must manually set the state of local rules before you can use them in your intrusion policy.

Import Local Intrusion Rules

- Make sure your local rule file follows the guidelines described in Best Practices for Importing Local Intrusion Rules, on page 147.
- Make sure your process for importing local intrusion rules complies with your security policies.
- Consider the import's effect on traffic flow and inspection due to bandwidth constraints and Snort restarts. We recommend scheduling rule updates during maintenance windows.
- You can perform this task in any domain.

Use this procedure to import local intrusion rules. Imported intrusion rules appear in the local rule category in a disabled state.

Procedure

Step 1 Choose **System** > **Updates**, then click **Rule Updates**.

Step 2 (Optional) Delete existing local rules.

Click **Delete All Local Rules**, then confirm that you want to move all created and imported intrusion rules to the deleted folder.

Step 3 Under **One-Time Rule Update/Rules Import**, choose **Rule update or text rule file to upload and install**, then click **Choose File** and browse to your local rule file.

Step 4 Click **Import**.

Step 5 Monitor import progress in the Message Center.

To display the Message Center, click System Status on the menu bar. Even if the Message Center shows no progress for several minutes or indicates that the import has failed, do not restart the import. Instead, contact Cisco TAC.

What to do next

- Edit intrusion policies and enable the rules you imported.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Rule Update Log

The Firepower Management Center generates a record for each rule update and local rule file that you import.

Each record includes a time stamp, the name of the user who imported the file, and a status icon indicating whether the import succeeded or failed. You can maintain a list of all rule updates and local rule files that you import, delete any record from the list, and access detailed records for all imported rules and rule update components.

The Rule Update Import Log detailed view lists a detailed record for each object imported in a rule update or local rule file. You can also create a custom workflow or report from the records listed that includes only the information that matches your specific needs.

Intrusion Rule Update Log Table

Table 18: Intrusion Rule Update Log Fields

Field	Description
Summary	The name of the import file. If the import fails, a brief statement of the reason for the failure appears under the file name.
Time	The time and date that the import started.
User ID	The user name of the user that triggered the import.
Status	Whether the import: • **Succeeded** ( ) • failed or is currently in progress **Red Status** ( ) The red status icon indicating an unsuccessful or incomplete import appears on the Rule Update Log page during the import and is replaced by the green icon only when the import has successfully completed.

Tip You can view import details as they appear while an intrusion rule update import is in progress.

Viewing the Intrusion Rule Update Log

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **System** > **Updates**.

Tip You can also click **Import Rules** on the intrusion rules editor page (**Objects** > **Intrusion Rules**).

Step 2 Click **Rule Updates**.

Step 3 Click **Rule Update Log**.

Step 4 You have two options:

- View — To view details for each object imported in a rule update or local rule file, click **View** ( ) next to the file you want to view; see Viewing Details of the Intrusion Rule Update Import Log, on page 152.
- Delete — To delete an import file record from the import log, including detailed records for all objects included with the file, click **Delete** ( ) next to the import file name.

  Note Deleting the file from the log does not delete any object imported in the import file, but only deletes the import log records.

Fields in an Intrusion Rule Update Log

Tip You search the entire Rule Update Import Log database even when you initiate a search by clicking **Search** on the toolbar from the Rule Update Import Log detailed view with only the records for a single import file displayed. Make sure you set your time constraints to include all objects you want to include in the search.

Table 19: Rule Update Import Log Detailed View Fields

Field	Description
Action	An indication that one of the following has occurred for the object type: • `new` (for a rule, this is the first time the rule has been stored on this appliance) • `changed` (for a rule update component or rule, the rule update component has been modified, or the rule has a higher revision number and the same GID and SID) • `collision` (for a rule update component or rule, import was skipped because its revision conflicts with an existing component or rule on the appliance) • `deleted` (for rules, the rule has been deleted from the rule update) • `enabled` (for a rule update edit, a preprocessor, rule, or other feature has been enabled in a default policy provided with the system) • `disabled` (for rules, the rule has been disabled in a default policy provided with the system) • `drop` (for rules, the rule has been set to Drop and Generate Events in a default policy provided with the system) • `error` (for a rule update or local rule file, the import failed) • `apply` (the **Reapply all policies after the rule update import completes** option was enabled for the import)
Default Action	The default action defined by the rule update. When the imported object type is `rule`, the default action is `Pass`, `Alert`, or `Drop`. For all other imported object types, there is no default action.
Details	A string unique to the component or rule. For rules, the GID, SID, and previous revision number for a changed rule, displayed as `previously (GID:SID:Rev)`. This field is blank for a rule that has not changed.
Domain	The domain whose intrusion policies can use the updated rule. Intrusion policies in descendant domains can also use the rule. This field is only present in a multidomain deployment.
GID	The generator ID for a rule. For example, `1` (standard text rule) or `3` (shared object rule).
Name	The name of the imported object, which for rules corresponds to the rule Message field, and for rule update components is the component name.
Policy	For imported rules, this field displays `All`. This means that the rule was imported successfully, and can be enabled in all appropriate default intrusion policies. For other types of imported objects, this field is blank.
Rev	The revision number for a rule.
Rule Update	The rule update file name.
SID	The SID for a rule.
Time	The time and date the import began.

Field	Description
Type	The type of imported object, which can be one of the following: • `rule update component` (an imported component such as a rule pack or policy pack) • `rule` (for rules, a new or updated rule; note that in Version 5.0.1 this value replaced the `update` value, which is deprecated) • `policy apply` (the **Reapply all policies after the rule update import completes** option was enabled for the import)
Count	The count (`1`) for each record. The Count field appears in a table view when the table is constrained, and the Rule Update Log detailed view is constrained by default to rule update records. This field is not searchable.

Viewing Details of the Intrusion Rule Update Import Log

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **System** > **Updates**.

Tip You can also click **Import Rules** on the Rule Editor page, which you access by choosing **Policies** > **Intrusion** > **Intrusion Rules**.

Step 2 Click **Rule Updates**.

Step 3 Click **Rule Update Log**.

Step 4 Click **View** ( ) next to the file whose detailed records you want to view.

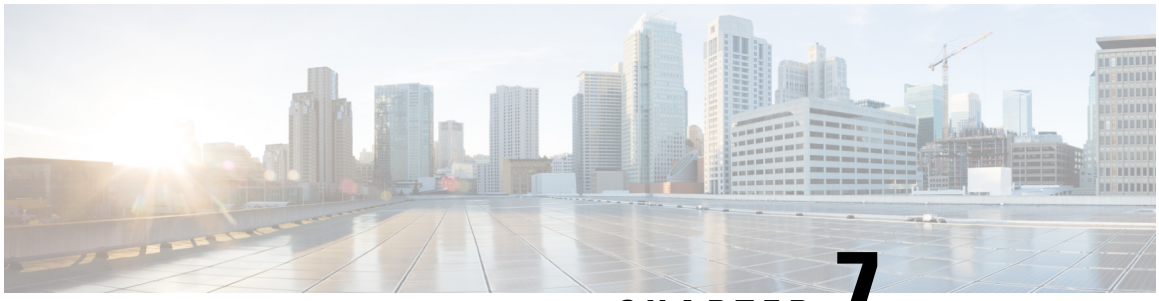
Step 5 You can take any of the following actions:

- Bookmark—To bookmark the current page, click **Bookmark This Page**.
- Edit Search—To open a search page prepopulated with the current single constraint, choose **Edit Search** or **Save Search** next to Search Constraints.
- Manage bookmarks—To navigate to the bookmark management page, click **Report Designer**.
- Report—To generate a report based on the data in the current view, click **Report Designer**.
- Search—To search the entire Rule Update Import Log database for rule update import records, click **Search**.
- Sort—To sort and constain records on the current workflow page, see Using Drill-Down Pages, on page 2158 for more information.
- Switch workflows—To temporarily use a different workflow, click **(switch workflows)**.

Maintain Your Air-Gapped Deployment

If your Firepower system is not connected to the internet, essential updates will not occur automatically. You must manually obtain and install these updates. See the following information:

- Manually Update the VDB, on page 140
- Update Intrusion Rules One-Time Manually, on page 145
- Manually Update the GeoDB (No Internet Connection), on page 142
- The *Firepower Management Center Software Upgrade Guide* at https://www.cisco.com/c/en/us/td/docs/security/firepower/upgrade/fpmc-upgrade-guide.html



CHAPTER 7

Backup and Restore

- About Backup and Restore, on page 155
- Requirements for Backup and Restore, on page 156
- Guidelines and Limitations for Backup and Restore, on page 157
- Best Practices for Backup and Restore, on page 158
- Backing Up FMCs or Managed Devices, on page 161
- Restoring FMCs and Managed Devices, on page 166
- Manage Backups and Remote Storage, on page 168

About Backup and Restore

The ability to recover from a disaster is an essential part of any system maintenance plan. As part of your disaster recovery plan, we recommend that you perform periodic backups to a secure remote location.

On-Demand Backups

You can perform on-demand backups for the FMC and 7000/8000 series devices from the FMC.

You can also use the local web interface on a 7000/8000 series device to perform on-demand backups. Local backup management on 7000/8000 series devices is slightly different and has fewer options than backup management on the FMC, but in general works in the same way. Note that you can use the FMC to back up these devices remotely.

For more information, see Backing Up FMCs or Managed Devices, on page 161.

Scheduled Backups

You can use the scheduler on an FMC or 7000/8000 series device to automate backups. You cannot schedule remote device backups from the FMC.

For more information, see Scheduled Backups, on page 181.

Storing Backup Files

You can store backups locally. However, we recommend you back up FMCs and managed devices to a secure remote location by mounting an NFS, SMB, or SSHFS network volume as remote storage. After you do this, all subsequent backups are copied to that volume, but you can still use the FMC to manage them.

For more information, see Remote Storage Management, on page 890 and Manage Backups and Remote Storage, on page 168.

Restoring the FMC and Managed Devices

You restore the FMC and 7000/8000 series devices from the local Backup Management page.

For more information, see Restoring FMCs and Managed Devices, on page 166.

What Is Backed Up?

FMC backups can include:

- Configurations.

  All configurations you can set on the FMC web interface are included in a configuration backup, with the exception of remote storage and audit log server certificate settings. In a multidomain deployment, you must back up configurations. You cannot back up events or only.

- Events.

  Event backups include all events in the FMC database. However, FMC event backups do not include intrusion event review status. Restored intrusion events do not appear on Reviewed Events pages.

7000/8000 series device backups are always configuration-only.

What Is Restored?

Restoring configurations overwrites *all* backed-up configurations, with very few exceptions. On the FMC, restoring events overwrites *all* existing events, with the exception of intrusion events.

Make sure you understand and plan for the following:

- You cannot restore what is not backed up.

  FMC configuration backups do not include remote storage and audit log server certificate settings, so you must reconfigure these after restore. Also, because FMC event backups do not include intrusion event review status, restored intrusion events do not appear on Reviewed Events pages.

- Restoring to a configured FMC — instead of factory-fresh or reimaged — merges intrusion events and file lists.

  The FMC event restore process does not overwrite intrusion events. Instead, the intrusion events in the backup are added to the database. To avoid duplicates, delete existing intrusion events before you restore.

  The FMC configuration restore process does not overwrite clean and custom detection file lists used by AMP for Networks. Instead, it merges existing file lists with the file lists in the backup. To replace file lists, delete existing file lists before you restore.

Requirements for Backup and Restore

Backup and restore has the following requirements.

Model Requirements: Backup

You can back up:

- FMCs
- 7000/8000 series devices

Backup is *not* supported for:

- Firepower Threat Defense
- NGIPSv
- ASA FirePOWER

If you need to replace a device where backup and restore is not supported, you must manually recreate device-specific configurations. However, backing up the FMC does back up policies and other configurations that you deploy to managed devices, as well as events already transmitted from the devices to the FMC.

Model Requirements: Restore

A replacement appliance must be the same model as the one you are replacing. Replacement managed devices should have the same number of network modules and same type and number of physical interfaces.

Version Requirements

As the first step in any backup, note the patch level. To restore a backup, the old and the new appliance must be running the same Firepower version, including patches.

For FMC backups, you must also have the same VDB. You are *not* required to have the same SRU.

License Requirements

Address licensing or orphan entitlements concerns as described in the best practices and procedures. If you notice licensing conflicts, contact Cisco TAC.

Domain Requirements

To:

- Back up or restore the FMC: Global only.
- Back up a device from the FMC: Global only.
- Restore a device: None. Restore devices locally.

In a multidomain deployment you cannot back up only events. You must also back up configurations.

Guidelines and Limitations for Backup and Restore

Backup and restore has the following guidelines and limitations.

Backup and Restore is for Disaster Recovery/RMA

Backup and restore is primarily intended for RMA scenarios. Before you begin the restore process of a faulty or failed physical appliance, contact Cisco TAC for replacement hardware.

Backup and Restore is not Configuration Import/Export

A backup file contains information that uniquely identifies an appliance, and cannot be shared. Do not use the backup and restore process to copy configurations between appliances or devices, or as a way to save configurations while testing new ones. Instead, use the import/export feature.

Restore is Individual and Local

You restore to cloud-delivered Firewall Management Centers and threat defensemanaged devices individually and locally. This means:

- You cannot batch-restore to high availability (HA) FMCs or devices. The restore procedures in this guide explain how to restore in an HA environment.
- You cannot use the cloud-delivered Firewall Management Center to restore a device. For the cloud-delivered Firewall Management Center and 7000/8000 series devices, you can use the local web interface to restore.
- You cannot use an cloud-delivered Firewall Management Center user account to log into and restore one of its managed devices. cloud-delivered Firewall Management Centers and devices maintain their own user accounts.

Best Practices for Backup and Restore

Backup and restore has the following best practices.

When to Back Up

We recommend backing up during a maintenance window or other time of low use.

While the system collects backup data, there may be a temporary pause in data correlation (cloud-delivered Firewall Management Center only), and you may be prevented from changing configurations related to the backup. If you include event data, event-related features such as eStreamer are not available.

You should back up in the following situations:

- Regular scheduled backups.

  As part of your disaster recovery plan, we recommend that you perform periodic backups. To automate this process, see Scheduled Backups, on page 181.

- Before upgrade or reimage.

  If an upgrade fails catastrophically, you may have to reimage and restore. Reimaging returns most settings to factory defaults, including the system password. If you have a recent backup, you can return to normal operations more quickly.

- After upgrade.

  Back up after you upgrade, so you have a snapshot of your freshly upgraded deployment. We recommend you back up the cloud-delivered Firewall Management Center *after* you upgrade its managed devices, so your new cloud-delivered Firewall Management Center backup file 'knows' that its devices have been upgraded.

Maintaining Backup File Security

Backups are stored as unencrypted archive (.tar) files.

Private keys in PKI objects—which represent the public key certificates and paired private keys required to support your deployment—are decrypted before they are backed up. The keys are reencrypted with a randomly generated key when you restore the backup.



Caution We recommend you back up cloud-delivered Firewall Management Centers and devices to a secure remote location and verify transfer success. Backups left locally may be deleted, either manually or by the upgrade process, which purges locally stored backups.

Especially because backup files are unencrypted, do *not* allow unauthorized access. If backup files are modified, the restore process will fail. Keep in mind that anyone with the Admin/Maint role can access the Backup Management page, where they can move and delete files from remote storage.

In the cloud-delivered Firewall Management Center's system configuration, you can mount an NFS, SMB, or SSHFS network volume as remote storage. After you do this, all subsequent backups are copied to that volume, but you can still use the cloud-delivered Firewall Management Center to manage them. For more information, see Remote Storage Management, on page 890 and Manage Backups and Remote Storage, on page 168.

Note that only the cloud-delivered Firewall Management Center mounts the network volume. Managed device backup files are routed through the cloud-delivered Firewall Management Center. Make sure you have the bandwidth to perform a large data transfer between the cloud-delivered Firewall Management Center and its devices. For more information, see Guidelines for Downloading Data from the Firepower Management Center to Managed Devices (Troubleshooting TechNote).

Backup and Restore in FMC High Availability Deployments

In an cloud-delivered Firewall Management Center high availability deployment, backing up one cloud-delivered Firewall Management Center does not back up the other. You should regularly back up both peers. Do not restore one HA peer with the backup file from the other. A backup file contains information that uniquely identifies an appliance, and cannot be shared.

Note that you can replace an HA cloud-delivered Firewall Management Center without a successful backup. For more information on replacing HAcloud-delivered Firewall Management Centers, both with and without successful backups, see Replacing FMCs in a High Availability Pair, on page 214.

Before Backup

Before you back up, you should:

- Update the VDB and SRU on the cloud-delivered Firewall Management Center.

  We always recommend you use the latest vulnerability database (VDB) and intrusion rules (SRU). Before you back up an cloud-delivered Firewall Management Center, check the Cisco Support & Download site for newer versions.

  This is especially important for the VDB, because the VDB versions must match to restore a backup. Because you cannot downgrade the VDB, you do not want a situation where your replacement cloud-delivered Firewall Management Center has a newer VDB than the backed up cloud-delivered Firewall Management Center.

- Check Disk Space.

Before you begin a backup, make sure you have enough disk space on the appliance or on your remote storage server. The space available is displayed on the Backup Management page.

Backups can fail if there is not enough space. Especially if you schedule backups, make sure you regularly prune backup files or allocate more disk space to the remote storage location.

Before Restore

Before restore, you should:

- Revert licensing changes.

  Revert any licensing changes made since you took the backup.

  Otherwise, you may have license conflicts or orphan entitlements after the restore. However, do *not* unregister from Cisco Smart Software Manager (CSSM). If you unregister from CSSM, you must unregister again after you restore, then re-register.

  After the restore completes, reconfigure licensing. If you notice licensing conflicts or orphan entitlements, contact Cisco TAC.

- Disconnect faulty appliances.

  Disconnect the management interface, and for devices, the data interfaces.

  Note that restoring an cloud-delivered Firewall Management Center or 7000/8000 series device does *not* change the management IP address. You must set that manually on the replacement — just make sure you disconnect the old appliance from the network before you do.

- Do *not* unregister managed devices.

  Whether you are restoring an FMC or managed device, do not unregister devices from the cloud-delivered Firewall Management Center, even if you physically disconnect an appliance from the network.

  If you unregister, you will need to redo some device configurations, such as security zone to interface mappings. After you restore, the cloud-delivered Firewall Management Center and devices should begin communicating normally.

- Reimage.

  In an RMA scenario, the replacement appliance will arrive configured with factory defaults. However, if the replacement appliance is already configured, we recommend you reimage. Reimaging returns most settings to factory defaults, including the system password. You can only reimage to major versions, so you may need to patch after you reimage.

  If you do not reimage, keep in mind that cloud-delivered Firewall Management Center intrusion events and file lists are merged rather than overwritten.

After Restore

After restore, you should:

- Reconfigure anything that was not restored.

  This can include reconfiguring licensing, remote storage, and audit log server certificate settings.

- Update the VDB and SRU on the cloud-delivered Firewall Management Center.

  We always recommend you use the latest vulnerability database (VDB) and intrusion rules (SRU).

- Deploy.

  After you restore an cloud-delivered Firewall Management Center, deploy to all managed devices. After you restore a device, deploy to that device. You *must* deploy. If the a device or devices are not marked out of date, force deploy from the Device Management page: Redeploy Existing Configurations to a Device, on page 348.

Backing Up FMCs or Managed Devices

You can perform on-demand or scheduled backups for supported appliances.

You do not need a backup profile to back up 7000/8000 series devices from the FMC. However, FMC backups require backup profiles, as do local backups on 7000/8000 series devices.. The on-demand backup process allows you to create a new backup profile.

For more information, see:

- Back up the FMC, on page 161
- Back up a Device from the FMC, on page 163
- Back up a 7000/8000 Series Device Locally, on page 163
- Create a Backup Profile, on page 165
- Scheduled Backups, on page 181

Back up the FMC

Use this procedure to perform an on-demand FMC backup. To back up a 7000/8000 series device from its local web interface, see Back up a 7000/8000 Series Device Locally, on page 163.

Before you begin

You must read and understand the requirements, guidelines, limitations, and best practices. You do not want to skip any steps or ignore security concerns. Careful planning and preparation can help you avoid missteps.

- Requirements for Backup and Restore, on page 156
- Guidelines and Limitations for Backup and Restore, on page 157
- Best Practices for Backup and Restore, on page 158

Procedure

Step 1 Select **System** > **Tools** > **Backup/Restore**.

The Backup Management page lists all locally and remotely stored backups. It also lists how much disk space you have available to store backups. Backups can fail if there is not enough space.

Step 2 Choose whether to use an existing backup profile or start fresh.

FMC backups require that you use or create a backup profile.

- Click **Backup Profiles** to use an existing backup profile.

  Next to the profile you want to use, click the edit icon. You can then click **Start Backup** to begin the backup right now. Or, if you want to edit the profile, go on to the next step.

- Click **Firepower Management Backup** to start fresh and create a new backup profile.

  Enter a **Name** for the backup profile.

Step 3 Choose what to back up:

- **Back Up Configuration**
- **Back Up Events**

In a multidomain deployment, you must back up configurations. You cannot back up events or only. For details on what is and what is not backed up for each of these choices, see About Backup and Restore, on page 155.

Step 4 Note the **Storage Location** for FMC backup files.

This will either be local storage in `/var/sf/backup/`, or a remote network volume. For more information, see Manage Backups and Remote Storage, on page 168.

Step 5 (Optional) Enable **Copy when complete** to copy completed FMC backups to a remote server.

Provide a hostname or IP address, the path to the remote directory, and a username and password. To use an SSH public key instead of a password, copy the contents of the **SSH Public Key** field to the specified user's `authorized_keys` file on the remote server.

Note This option is useful if you want to store backups locally and also SCP them to a remote location. If you configured SSH remote storage, do *not* copy backup files to the same directory using **Copy when complete**.

Step 6 (Optional) Enable **Email** and enter an email address to be notified when the backup completes.

To receive email notifications, you must configure the FMC to connect to a mail server: Configuring a Mail Relay Host and Notification Address, on page 906.

Step 7 Click **Start Backup** to start the on-demand backup.

If you are not using an existing backup profile, the system automatically creates one and uses it. If you decide not to run the backup now, you can click **Save** or **Save As New** to save the profile. In either case, you can use the newly created profile to configure scheduled backups.

Step 8 Monitor progress in the Message Center.

While the system collects backup data, there may be a temporary pause in data correlation, and you may be prevented from changing configurations related to the backup. If you configured remote storage or enabled **Copy when complete**, the FMC may write temporary files to the remote server. These files are cleaned up at the end of the backup process.

What to do next

If you configured remote storage or enabled **Copy when complete**, verify transfer success of the backup file.

Back up a Device from the FMC

Use this procedure to perform an on-demand backup of a 7000/8000 series device from the FMC. At this time, backup and restore is not supported for Firepower Threat Defense.

Before you begin

You must read and understand the requirements, guidelines, limitations, and best practices. You do not want to skip any steps or ignore security concerns. Careful planning and preparation can help you avoid missteps.

- Requirements for Backup and Restore, on page 156
- Guidelines and Limitations for Backup and Restore, on page 157
- Best Practices for Backup and Restore, on page 158

Procedure

Step 1 Select **System** > **Tools** > **Backup/Restore**, then click **Managed Device Backup**.

Step 2 Select one or more **Managed Devices**.

Step 3 To back up event data that has not yet been sent to the FMC, select **Include All Unified Files**.

Step 4 Note the **Storage Location** for device backup files.

This will either be local storage in `/var/sf/remote-backup/`, or a remote network volume. For more information, see Manage Backups and Remote Storage, on page 168.

Step 5 If you did not configure remote storage, choose whether you want to **Retrieve to Management Center**.

- Enabled: Saves the backup to the FMC in `/var/sf/remote-backup/`.
- Disabled (default): Saves the backup to the device in `/var/sf/backup`.

Step 6 Click **Start Backup** to start the on-demand backup.

Step 7 Monitor progress in the Message Center.

What to do next

If you configured remote storage, verify transfer success of the backup file.

Back up a 7000/8000 Series Device Locally

Use this procedure to perform a local, on-demand backup for a 7000/8000 series device. Device backups are always configuration-only.

Note that local backup management on 7000/8000 series devices is slightly different and has fewer options than backup management on the FMC, but in general works in the same way. Unless you have a specific need (such as scheduling backups), we recommend you use the FMC to back up these devices remotely.

Before you begin

You must read and understand the requirements, guidelines, limitations, and best practices. You do not want to skip any steps or ignore security concerns. Careful planning and preparation can help you avoid missteps.

- Requirements for Backup and Restore, on page 156
- Guidelines and Limitations for Backup and Restore, on page 157
- Best Practices for Backup and Restore, on page 158

Procedure

Step 1 On the device's local web interface, select **System** > **Tools** > **Backup/Restore**.

The Backup Management page lists all locally stored backups. It also lists how much disk space you have available to store backups. Backups can fail if there is not enough space.

Step 2 Choose whether to use an existing backup profile or start fresh.

7000/8000 series local backups require that you use or create a backup profile. When you perform an on-demand backup, if you do not pick an existing backup profile, the system automatically creates one and uses it. You can then use the newly created profile to configure scheduled backups.

- Click **Backup Profiles** to use an existing backup profile.

  Next to the profile you want to use, click the edit icon. You can then click **Start Backup** to begin the backup right now. Or, if you want to edit the profile, go on to the next step.

- Click **Device Backup** to start fresh and create a new backup profile.

  Enter a **Name** for the backup profile.

Step 3 (Optional) Enable **Copy when complete** to copy completed backups to a remote server.

This is your only option for remote storage for 7000/8000 series local backups.

Provide a hostname or IP address, the path to the remote directory, and a username and password. To use an SSH public key instead of a password, copy the contents of the **SSH Public Key** field to the specified user's `authorized_keys` file on the remote server.

Step 4 (Optional) Enable **Email** and enter an email address to be notified when the backup completes.

To receive email notifications, you must configure the device to connect to a mail server: Configuring a Mail Relay Host and Notification Address, on page 906.

Step 5 Click **Start Backup** to start the on-demand backup.

If you are not using an existing backup profile, the system automatically creates one and uses it. If you decide not to run the backup now, you can click **Save** or **Save As New** to save the profile. In either case, you can use the newly created profile to configure scheduled backups.

Step 6 Monitor progress in the Message Center.

While the system collects backup data, you may be prevented from changing configurations related to the backup. If you enabled **Copy when complete**, the device may write temporary files to the remote server. These files are cleaned up at the end of the backup process.

What to do next

If you enabled **Copy when complete**, verify transfer success of the backup file.

Create a Backup Profile

A backup profile is a saved set of preferences—what to back up, where to store the backup file, and so on.

FMC backups and 7000/8000 series local backups require backup profiles. Backup profiles are not required to back up a device from the FMC.

When you perform an on-demand FMC or 7000/8000 series local backup, if you do not pick an existing backup profile, the system automatically creates one and uses it. You can then use the newly created profile to configure scheduled backups. Note that you cannot schedule 7000/8000 series device backups from the FMC.

The following procedure explains how to create a backup profile without performing an on-demand backup.

Procedure

Step 1 Select **System** > **Tools** > **Backup/Restore**, then click **Backup Profiles**.

Step 2 Click **Create Profile** and enter a **Name**.

Step 3 (FMC only) Choose what to back up.

7000/8000 series backups are always configuration-only.

- **Back Up Configuration**
- **Back Up Events**

In a multidomain deployment, you must back up configurations. You cannot back up events or only. For details on what is and what is not backed up for each of these choices, see About Backup and Restore, on page 155.

Step 4 Note the **Storage Location** for backup files.

For FMC backup profiles, this will either be local storage in `/var/sf/backup/`, or a remote network volume. For 7000/8000 local backup profiles, this is always `/var/sf/backup/`. For more information, see Manage Backups and Remote Storage, on page 168.

Step 5 (Optional) Enable **Copy when complete** to copy completed FMC backups to a remote server.

Provide a hostname or IP address, the path to the remote directory, and a username and password. To use an SSH public key instead of a password, copy the contents of the **SSH Public Key** field to the specified user's `authorized_keys` file on the remote server.

Note This option is useful if you want to store backups locally and also SCP them to a remote location. If you configured SSHFS remote storage, do *not* copy backup files to the same directory using **Copy when complete**.

Step 6 (Optional) Enable **Email** and enter an email address to be notified when the backup completes.

To receive email notifications, you must configure the FMC to connect to a mail server: Configuring a Mail Relay Host and Notification Address, on page 906.

Step 7 Click **Save**.

Restoring FMCs and Managed Devices

For the FMC and 7000/8000 series devices, you use the local web interface to restore from backup. You cannot use the FMC to restore a device.

The following sections explain how to restore FMCs and managed devices.

- Restore an FMC from Backup, on page 166
- Replacing FMCs in a High Availability Pair, on page 214
- Restore a 7000/8000 Series Device from Backup, on page 167

Restore an FMC from Backup

When you restore an FMC backup, you can choose to restore any or all of the components included in the backup file (events, configurations).

Note Restoring configurations overwrites *all* configurations, with very few exceptions. It also reboots the FMC. Restoring events overwrites *all* existing events, with the exception of intrusion events. Make sure you are ready.

Use this procedure to restore an FMC from backup. For more information on backup and restore in an FMC HA deployment, see Replacing FMCs in a High Availability Pair, on page 214. To restore a 7000/8000 series device, see Restore a 7000/8000 Series Device from Backup, on page 167.

Before you begin

You must read and understand the requirements, guidelines, limitations, and best practices. You do not want to skip any steps or ignore security concerns. Careful planning and preparation can help you avoid missteps.

- Requirements for Backup and Restore, on page 156
- Guidelines and Limitations for Backup and Restore, on page 157
- Best Practices for Backup and Restore, on page 158

Procedure

Step 1 Log into the FMC you want to restore.

Step 2 Select **System** > **Tools** > **Backup/Restore**.

The Backup Management page lists all locally and remotely stored backup files. You can click a backup file to view its contents.

If the backup file is not in the list and you have it saved on your local computer, click **Upload Backup**; see Manage Backups and Remote Storage, on page 168.

Step 3 Select the backup file you want to restore and click **Restore**.

Step 4 Select from the available components to restore, then click **Restore** again to begin.

Step 5 Monitor progress in the Message Center.

If you are restoring configurations, you can log back in after the FMC reboots.

What to do next

- If necessary, reconfigure any licensing settings that you reverted before the restore. If you notice licensing conflicts or orphan entitlements, contact Cisco TAC.
- If necessary, reconfigure remote storage and audit log server certificate settings. These settings are not included in backups.
- (Optional) Update the SRU and VDB. If the SRU or the VDB available on the Cisco Support & Download site is newer than the version currently running, we recommend you install the newer version.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Restore a 7000/8000 Series Device from Backup

This procedure explains how to use the 7000/8000 series local web interface to restore from backup. Restoring overwrites *all* configurations, with very few exceptions. It also reboots the device.

Before you begin

You must read and understand the requirements, guidelines, limitations, and best practices. You do not want to skip any steps or ignore security concerns. Careful planning and preparation can help you avoid missteps.

- Requirements for Backup and Restore, on page 156
- Guidelines and Limitations for Backup and Restore, on page 157
- Best Practices for Backup and Restore, on page 158

Procedure

Step 1 Log into the device you want to restore.

Step 2 Select **System** > **Tools** > **Backup/Restore**.

The Backup Management page lists all locally stored backup files. You can click a backup file to view its contents.

If the backup file is not in the list and you have it saved on your local computer, click **Upload Backup**; see Manage Backups and Remote Storage, on page 168.

Step 3 Select the backup file you want to restore and click **Restore**.

Step 4 Make sure **Replace Configuration Data** is enabled, then click **Restore** again to begin.

Device backups are always configuration-only.

Step 5 Monitor progress in the Message Center until the device reboots.

What to do next

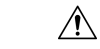
Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Manage Backups and Remote Storage

Backups are stored as unencrypted archive (.tar) files. The file name includes identifying information that can include:

- The name of the backup profile or scheduled task associated with the backup.
- The display name or IP address of the backed-up appliance.
- The appliance's role, such as a member of an HA pair.

We recommend you back up appliances to a secure remote location and verify transfer success. Backups left on an appliance may be deleted, either manually or by the upgrade process; upgrades purge locally stored backups. For more information on your options, see Backup Storage Locations, on page 169.



Caution Especially because backup files are unencrypted, do *not* allow unauthorized access. If backup files are modified, the restore process will fail. Keep in mind that anyone with the Admin/Maint role can access the Backup Management page, where they can move and delete files from remote storage.

The following procedure describes how to manage backup files.

Procedure

Step 1 Select **System** > **Tools** > **Backup/Restore**.

The Backup Management page lists available backups. It also lists how much disk space you have available to store backups. Backups can fail if there is not enough space.

Step 2 Do one of the following:

Table 20: Remote Storage and Backup File Management

To	Do This
Enable or disable remote storage for backups without having to edit the FMC system configuration.	Click **Enable Remote Storage for Backups**. This option appears only after you configure remote storage. Toggling it here also toggles it in the system configuration (**System > Configuration > Remote Storage Device**). **Tip** To quickly access your remote storage configuration, click **Remote Storage** at the upper right of the Backup Management page. **Note** To store backup on the remote storage location, you must also enable the **Retrieve to Management Center** option (see Back up a Device from the FMC, on page 163).
Move a file between the FMC and the remote storage location.	Click **Move**. You can move a file back and forth as many times as you want. This will delete—not copy—the file from the current location. When you move a backup file from remote storage to the FMC, where it is stored on the FMC depends on the kind of backup: • FMC backups: `/var/sf/backup` • Device backups: `/var/sf/remote-backup`
View the contents of the backup.	Click the backup file.
Delete a backup file.	Choose a backup file and click **Delete**. You can delete both locally and remotely stored backup files.
Upload a backup file from your computer.	Click **Upload Backup**, choose a backup file, and click **Upload Backup** again.
Download a backup to your computer.	Choose a backup file and click **Download**. Unlike moving a backup file, this does not delete the backup from the FMC.

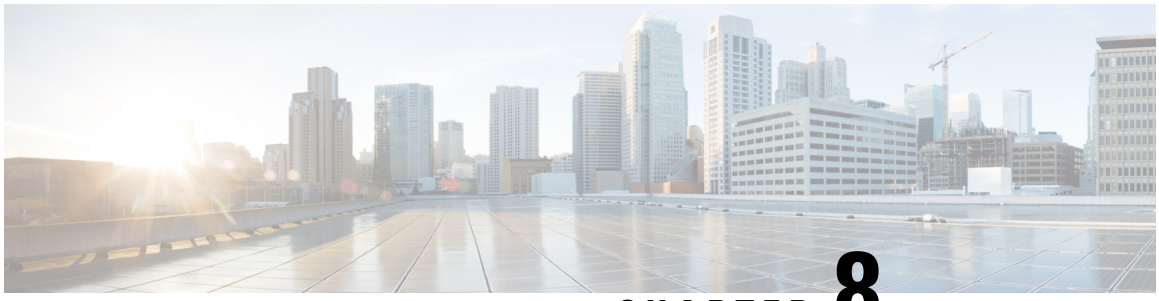
Backup Storage Locations

The following table describes backup storage options for FMCs and managed devices.

Table 21: Backup Storage Locations

Location	Details
Remote, by mounting a network volume (NFS, SMB, SSHFS).	**Note** Backup is stored on a remote storage location only when you have configured remote storage and enabled the **Retrieve to Management Center** option (see Back up a Device from the FMC, on page 163). In the FMC's system configuration, you can mount an NFS, SMB, or SSHFS network volume as remote storage for FMC and device backups; see Remote Storage Management, on page 890.) After you do this, all subsequent FMC backups *and FMC-initiated device backups* are copied to that volume, but you can still use the FMC to manage them (restore, download, upload, delete, move). Note that only the FMC mounts the network volume. Managed device backup files are routed through the FMC. Make sure you have the bandwidth to perform a large data transfer between the FMC and its devices. For more information, see Guidelines for Downloading Data from the Firepower Management Center to Managed Devices (Troubleshooting TechNote).
Remote, by copying (SCP).	**Note** Backup is stored on a remote storage location only when you have configured remote storage and enabled the **Retrieve to Management Center** option (see Back up a Device from the FMC, on page 163). For the FMC and for 7000/8000 series *local* backups, you can use a **Copy when complete** option to securely copy (SCP) completed backups to a remote server. Compared with remote storage by mounting a network volume, **Copy when complete** cannot copy to NFS or SMB volumes. You cannot provide CLI options or set a disk space threshold, and it does not affect remote storage of reports. You also cannot manage backup files after they are copied out. This option is useful if you want to store backups locally *and* SCP them to a remote location. It is also your only option for remote storage for 7000/8000 series local backups. **Note** If you configure SSHFS remote storage in the FMC system configuration, do *not* copy backup files to the same directory using **Copy when complete**.
Local, on the FMC.	If you do not configure remote storage by mounting a network volume, you can save backup files on the FMC: • FMC backups are saved to `/var/sf/backup`. • Device backups are saved to `/var/sf/remote-backup` on the FMC if you enable the **Retrieve to Management Center** option when you perform the backup. Note that you cannot save 7000/8000 series local backups to the FMC.

Location	Details
Local, on the device internal flash memory.	Device backup files are saved to `/var/sf/backup` on the device if you: • Do not configure remote storage by mounting a network volume. • Do not enable **Retrieve to Management Center**.



CHAPTER 8

Configuration Import and Export

The following topics explain how to use the Import/Export feature:

- About Configuration Import/Export, on page 173
- Requirements and Prerequisites for Configuration Import/Export, on page 175
- Exporting Configurations, on page 175
- Importing Configurations, on page 176

About Configuration Import/Export

You can use the Import/Export feature to copy configurations between appliances. Import/Export is not a backup tool, but can simplify the process of adding new appliances to your deployment.

You can export a single configuration, or you can export a set of configurations (of the same type or of different types) with a single action. When you later import the package onto another appliance, you can choose which configurations in the package to import.

An exported package contains revision information for that configuration, which determines whether you can import that configuration onto another appliance. When the appliances are compatible but the package includes a duplicate configuration, the system offers resolution options.

Note The importing and exporting appliances must be running the same version of the Firepower System. For access control and its subpolicies (including intrusion policies), the intrusion rule update version must also match. If the versions do not match, the import fails. You cannot use the Import/Export feature to update intrusion rules. Instead, download and apply the latest rule update version.

Configurations that Support Import/Export

Import/Export is supported for the following configurations:

- Access control policies and the policies they invoke: prefilter, network analysis, intrusion, SSL, file
- Intrusion policies, independently of access control
- NAT policies (Firepower Threat Defense only)
- FlexConfig policies.

- Platform settings
- Health policies
- Alert responses
- Application detectors (both user-defined and those provided by Cisco Professional Services)
- Dashboards
- Custom tables
- Custom workflows
- Saved searches
- Custom user roles
- Report templates
- Third-party product and vulnerability mappings

Special Considerations for Configuration Import/Export

When you export a configuration, the system also exports other required configurations. For example, exporting an access control policy also exports any subpolicies it invokes, objects and object groups it uses, ancestor policies (in a multidomain deployment), and so on. As another example, if you export a platform settings policy with external authentication enabled, the authentication object is exported as well. There are some exceptions, however:

- System-provided databases and feeds—The system does not export URL filtering category and reputation data, Cisco Intelligence Feed data, or the geolocation database (GeoDB). Make sure all the appliances in your deployment obtain up-to-date information from Cisco.
- Global Security Intelligence lists—The system exports Global Security Intelligence Block and Do Not Block lists associated with exported configurations. (In a multidomain deployment, this occurs regardless of your current domain. The system does **not** export descendant domain lists.) The import process converts these lists to user-created lists, then uses those new lists in the imported configurations. This ensures that imported lists do not conflict with existing Global Block and Do Not Block lists. To use Global lists on the importing Firepower Management Center in your imported configurations, add them manually.
- Intrusion policy shared layers—The export process breaks intrusion policy shared layers. The previously shared layer is included in the package, and imported intrusion policies do not contain shared layers.
- Intrusion policy default variable set—The export package includes a default variable set with custom variables and system-provided variables with user-defined values. The import process updates the default variable set on the importing Firepower Management Center with the imported values. However, the import process does **not** delete custom variables not present in the export package. The import process also does not revert user-defined values on the importing Firepower Management Center, for values not set in the export package. Therefore, an imported intrusion policy may behave differently than expected if the importing Firepower Management Center has differently configured default variables.
- Custom user objects—If you have created custom user groups or objects in your Firepower Management Center and if such a custom user object is a part of any rule in your access control policy, note that the export file (.sfo) does not carry the user object information and therefore while importing such a policy, any reference to such custom user objects will be removed and will not be imported to the destination

Firepower Management Center. To avoid detection issues due to the missing user group, add the customized user objects manually to the new Firepower Management Center and re-configure the access control policy after import.

When you import objects and object groups:

- The import process imports objects and groups as new. You cannot replace existing objects and groups.
- If the names of imported objects match existing objects on the importing Firepower Management Center, the system appends autogenerated numbers to the imported object and group names to make them unique.
- You must map any security zones and interface groups used in the imported configurations to matching-type zones and groups managed by the importing Firepower Management Center.
- If you export a configuration that uses PKI objects containing private keys, the system decrypts the private keys before export. On import, the system encrypts the keys with a randomly generated key.

Requirements and Prerequisites for Configuration Import/Export

Model Support

Any

Supported Domains

Any

User Roles

- Admin

Exporting Configurations

Depending on the number of configurations being exported and the number of objects those configurations reference, the export process may take several minutes.



Tip Many list pages in the Firepower System include an **YouTube EDU** () next to list items. Where this icon is present, you can use it as a quick alternative to the export procedure that follows.

Before you begin

- Confirm that the importing and exporting appliances are running the same version of the Firepower System. For access control and its subpolicies (including intrusion policies), the intrusion rule update version must also match.

Procedure

Step 1 Choose **System** > **Tools** > **Import/Export**.

Step 2 Click **Collapse** () and **Expand** () to collapse and expand the list of available configurations.

Step 3 Check the configurations you want to export and click **Export**.

Step 4 Follow your web browser's prompts to save the exported package to your computer.

Importing Configurations

Depending on the number of configurations being imported and the number of objects those configurations reference, the import process may take several minutes.

Before you begin

- Confirm that the importing and exporting appliances are running the same version of the Firepower System. For access control and its subpolicies (including intrusion policies), the intrusion rule update version must also match.

Procedure

Step 1 On the importing appliance, choose **System** > **Tools** > **Import/Export**.

Step 2 Click **Upload Package**.

Step 3 Enter the path to the exported package or browse to its location, then click **Upload**.

Step 4 If there are no version mismatches or other issues, choose the configurations you want to import, then click **Import**.
If you do not need to perform any conflict resolution or interface object mapping, the import completes and a success message appears. Skip the rest of this procedure.

Step 5 If prompted, on the Import Conflict Resolution page, map interface objects used in the imported configurations to zones and groups with matching interface types managed by the importing Firepower Management Center.

Interface object type (security zone or interface group) and interface type (passive, inline, routed, and so on) of source and destination objects must match. For information, see Interface Objects: Interface Groups and Security Zones, on page 410.

If the configurations you are importing reference security zones or interface groups that do not already exist, you can map them to existing interface objects, or create new ones.

Note For individual access control policies, you have the option of replacing an existing policy with the imported ones. However, for nested access control policies, you can only import them as new policies.

Step 6 Click **Import**.

Step 7 If prompted, on the Import Resolution page, expand each configuration and choose the appropriate option as described in Import Conflict Resolution, on page 177.

Step 8 Click **Import**.

Step 9 Update all feeds.

For example, go to **Objects > Object Management > Security Intelligence** and click the **Update Feed** button on the URL, Network, and DNS Lists and Feeds pages.

Imported policies do not include feed contents.

Step 10 Wait for all feed updates to complete before deploying the policies to devices.

What to do next

- Optionally, view a report summarizing the imported configurations; see Viewing Task Messages, on page 316.

Import Conflict Resolution

When you attempt to import a configuration, the system determines whether a configuration of the same name and type already exists on the appliance. In a multidomain deployment, the system also determines whether a configuration is a duplicate of a configuration defined in the current domain or any of its ancestor or descendant domains. (You cannot view configurations in descendant domains, but if a configuration with a duplicate name exists in a descendant domain, the system notifies you of the conflict.) When an import includes a duplicate configuration, the system offers resolution options suitable to your deployment from among the following:

- **Keep existing**

  The system does not import that configuration.

- **Replace existing**

  The system overwrites the current configuration with the configuration selected for import.

- **Keep newest**

  The system imports the selected configuration only if its timestamp is more recent than the timestamp on the current configuration on the appliance.

- **Import as new**

  The system imports the selected duplicate configuration, appending a system-generated number to the name to make it unique. (You can change this name before completing the import process.) The original configuration on the appliance remains unchanged.

The resolution options the system offers depends on whether your deployment uses domains, and whether the imported configuration is a duplicate of a configuration defined in the current domain, or a configuration defined in an ancestor or descendant of the current domain. The following table lists when the system does or does not present a resolution option.

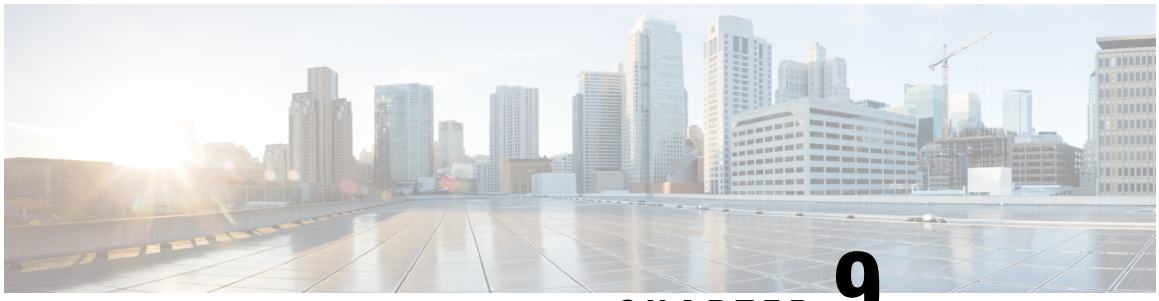
Resolution Option	Firepower Management Center		Managed Device
	Duplicate in current domain	Duplicate in ancestor or descendant domain	
Keep existing	Yes	Yes	Yes

Resolution Option	Firepower Management Center		Managed Device
	Duplicate in current domain	Duplicate in ancestor or descendant domain	
Replace existing	Yes	No	Yes
Keep newest	Yes	No	Yes
Import as new	Yes	Yes	Yes

When you import an access control policy with a file policy that uses clean or custom detection file lists and a file list presents a duplicate name conflict, the system offers conflict resolution options as described in the table above, but the action the system performs on the policies and file lists varies as described in the table below:

Resolution Option	System Action	
	Access control policy and its associated file policy are imported as new and the file lists are merged	Existing access control policy and its associated file policy and file lists remain unchanged
Keep existing	No	Yes
Replace existing	Yes	No
Import as new	Yes	No
Keep newest and access control policy being imported is the newest	Yes	No
Keep newest and existing access control policy is the newest	No	Yes

If you modify an imported configuration on an appliance, and later re-import that configuration to the same appliance, you must choose which version of the configuration to keep.



CHAPTER 9

Task Scheduling

The following topics explain how to schedule tasks:

- About Task Scheduling, on page 179
- Requirements and Prerequisites for Task Scheduling, on page 179
- Configuring a Recurring Task, on page 180
- Scheduled Task Review, on page 194

About Task Scheduling

You can schedule many different types of administrative tasks to run at designated times, either once or on a recurring basis.

Tasks configured using this feature are scheduled in UTC, which means when they occur locally depends on the date and your specific location. Also, because tasks are scheduled in UTC, they do not adjust for Daylight Saving Time, summer time, or any such seasonal adjustments that you may observe in your location. If you are affected, scheduled tasks occur one hour "later" in the summer than in the winter, according to local time.

Important We *strongly* recommend you review scheduled tasks to be sure they occur when you intend.

Note Some tasks (such as those involving automated software updates or that require pushing updates to managed devices) may place a significant load on networks with low bandwidths. You should schedule tasks like these to run during periods of low network use.

Requirements and Prerequisites for Task Scheduling

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Maintenance User

Configuring a Recurring Task

You set the frequency for a recurring task using the same process for all types of tasks.

Note that the time displayed on most pages on the web interface is the local time, which is determined by using the time zone you specify in your local configuration. Further, the Firepower Management Center automatically adjusts its local time display for daylight saving time (DST), where appropriate. However, recurring tasks that span the transition dates from DST to standard time and back do not adjust for the transition. That is, if you create a task scheduled for 2:00 AM during standard time, it will run at 3:00 AM during DST. Similarly, if you create a task scheduled for 2:00 AM during DST, it will run at 1:00 AM during standard time.

Procedure

Step 1 Select **System** > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From the **Job Type** drop-down list, select the type of task that you want to schedule.

Step 4 Click **Recurring** next to the **Schedule task to run** option.

Step 5 In the **Start On** field, specify the date when you want to start your recurring task.

Step 6 In the **Repeat Every** field, specify how often you want the task to recur.

You can either type a number or click **Up** (▲) and **Down** (▼) to specify the interval. For example, type `2` and click **Days** to run the task every two days.

Step 7 In the **Run At** field, specify the time when you want to start your recurring task.

Step 8 For a task to be run on a weekly or monthly basis, select the days when you want to run the task in the **Repeat On** field.

Step 9 Select the remaining options for the type of task you are creating:

- Backup - Schedule backup jobs as described in Schedule FMC Backups, on page 181.
- Download CRL - Schedule certificate revocation list downloads as described in Configuring Certificate Revocation List Downloads, on page 182.
- Deploy Policies - Schedule policy deployment as described in Automating Policy Deployment, on page 183.
- Nmap Scan - Schedule Nmap scans as described in Scheduling an Nmap Scan, on page 184.
- Report - Schedule report generation as described in Automating Report Generation, on page 185

- Firepower Recommended Rules - Schedule automatic update of Firepower recommended rules as described in Automating Firepower Recommendations, on page 187
- Download Latest Update - Schedule software or VDB update downloads as described in Automating Software Downloads, on page 189 or Automating VDB Update Downloads, on page 192.
- Install Latest Update - Schedule installation of software or VDB updates on a Firepower Management Center or managed device as described in Automating Software Installs, on page 190 or Automating VDB Update Installs, on page 192
- Push Latest Update - Schedule push of software updates to managed devices as described in Automating Software Pushes, on page 190.
- Update URL Filtering Database - Scheduling automatic update of URL filtering data as described in Automating URL Filtering Updates Using a Scheduled Task, on page 193

Step 10 Click **Save**

Scheduled Backups

You can use the scheduler on a Firepower Management Center or a 7000/8000 series device to automate its own backups. For more information on backups, see Backup and Restore, on page 155.

You cannot schedule device backups from the FMC.

Schedule FMC Backups

You can use the scheduler on the Firepower Management Center to automate its own backups. You cannot schedule a device backup from the FMC.

Before you begin

Create a backup profile that specifies your backup preferences: Create a Backup Profile, on page 165.

You must be in the global domain to perform this task.

Procedure

Step 1 Choose **System** > **Tools** > **Scheduling**.

Step 2 From the **Job Type** list, select **Backup**.

Step 3 Specify whether you want to back up **Once** or **Recurring**.

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 180.

Step 4 Enter a **Job Name**.

Step 5 Choose a **Backup Profile**.

Step 6 (Optional) Enter a **Comment**.

Keep comments brief. They will appear in the Task Details section of the schedule calendar page.

Step 7 (Optional) Enter an email address, or a comma-separated list of email addresses, in the **Email Status To:** field.

For information on setting up an email relay server to send task status messages, see Configuring a Mail Relay Host and Notification Address, on page 906.

Step 8 Click **Save**.

Schedule Local 7000 & 8000 Series Device Backups

You can use the scheduler on a 7000 or 8000 Series device to automate its own backups. You *cannot* schedule any kind of device backup from the FMC.

Before you begin

Create a backup profile that specifies your backup preferences: Create a Backup Profile, on page 165.

Procedure

Step 1 On the device's web interface, choose **System** > **Tools** > **Scheduling**.

Step 2 From the **Job Type** list, select **Backup**.

Step 3 Specify whether you want to back up **Once** or **Recurring**.

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 180.

Step 4 Enter a **Job Name**.

Step 5 Choose a **Backup Profile**.

Step 6 (Optional) Enter a **Comment**.

Keep comments brief. They will appear in the Task Details section of the schedule calendar page.

Step 7 (Optional) Enter an email address, or a comma-separated list of email addresses, in the **Email Status To:** field.

For information on setting up an email relay server to send task status messages, see Configuring a Mail Relay Host and Notification Address, on page 906.

Step 8 Click **Save**.

Configuring Certificate Revocation List Downloads

You must perform this procedure using the local web interface for the Firepower Management Center or the 7000 or 8000 Series device. In a multidomain deployment, this task is only supported in the Global domain for the Firepower Management Center.

The system automatically creates the Download CRL task when you enable downloading a certificate revocation list (CRL) in the local configuration on an appliance where you enable user certificates or audit log certificates for the appliance. You can use the scheduler to edit the task to set the frequency of the update.

Before you begin

- Enable and configure user certificates or audit log certificates and set one or more CRL download URLs. See Requiring Valid HTTPS Client Certificates, on page 876 and Require Valid Audit Log Server Certificates, on page 903 for more information.

Procedure

Step 1 Select **System** > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From **Job Type**, select **Download CRL**.

Step 4 Specify how you want to schedule the CRL download, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 180 for details.

Step 5 Type a name in the **Job Name** field.

Step 6 If you want to comment on the task, type a comment in the **Comment** field.

The comment field appears in the Task Details section of the schedule calendar page; keep comments brief.

Step 7 If you want to email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field. You must have a valid email relay server configured on the Firepower Management Center to send status messages.

Step 8 Click **Save**.

Related Topics

Configuring a Mail Relay Host and Notification Address, on page 906

Automating Policy Deployment

After modifying configuration settings in the FMC, you must deploy those changes to the affected devices.

In a multidomain deployment, you can schedule policy deployments only for your current domain.



Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 and Configurations that Restart the Snort Process When Deployed or Activated, on page 352.

Procedure

Step 1 Select **System** > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From **Job Type**, select **Deploy Policies**.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 180 for details.

Step 5 Type a name in the **Job Name** field.

Step 6 In the **Device** field, select a device where you want to deploy policies.

Step 7 If you want to comment on the task, type a comment in the **Comment** field.

The comment field displays in the Tasks Details section of the schedule calendar page; keep comments brief.

Step 8 If you want to email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field. You must have a valid email relay server configured to send status messages.

Step 9 Click **Save**.

Related Topics

Configuring a Mail Relay Host and Notification Address, on page 906
Out-of-Date Policies, on page 357

Nmap Scan Automation

You can schedule regular Nmap scans of targets on your network. Automated scans allow you to refresh information previously supplied by an Nmap scan. Because the Firepower System cannot update Nmap-supplied data, you need to rescan periodically to keep that data up to date. You can also schedule scans to automatically test for unidentified applications or servers on hosts in your network.

Note that a Discovery Administrator can also use an Nmap scan as a remediation. For example, when an operating system conflict occurs on a host, that conflict may trigger an Nmap scan. Running the scan obtains updated operating system information for the host, which resolves the conflict.

If you have not used the Nmap scanning capability before, you configure Nmap scanning before defining a scheduled scan.

Related Topics

Nmap Scanning, on page 1855

Scheduling an Nmap Scan

After Nmap replaces a host's operating system, applications, or servers detected by the system with the results from an Nmap scan, the system no longer updates the information replaced by Nmap for the host. Nmap-supplied service and operating system data remains static until you run another Nmap scan. If you plan to scan a host using Nmap, you may want to set up regularly scheduled scans to keep Nmap-supplied operating systems, applications, or servers up to date. If the host is deleted from the network map and re-added, any

Nmap scan results are discarded and the system resumes monitoring of all operating system and service data for the host.

In a multidomain deployment:

- You can schedule scans only for your current domain
- The remediation and Nmap targets you select must exist at your current domain or an ancestor domain.
- Choosing to perform an Nmap scan on a non-leaf domain scans the same targets in each descendant of that domain.

Procedure

Step 1 Select **System** > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From **Job Type**, select **Nmap Scan**.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 180 for details.

Step 5 Type a name in the **Job Name** field.

Step 6 In the **Nmap Remediation** field, select an Nmap remediation.

Step 7 In the **Nmap Target** field, select the scan target.

Step 8 In the **Domain** field, select the domain whose network map you want to augment.

Step 9 If you want to comment on the task, type a comment in the **Comment** field.

Tip The comment field appears in the Task Details section of the calendar schedule page; keep comments brief.

Step 10 If you want to email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field. You must have a valid email relay server configured to send status messages.

Step 11 Click **Save**.

Related Topics

Configuring a Mail Relay Host and Notification Address, on page 906

Automating Report Generation

You can automate reports so that they run at regular intervals.

In a multidomain deployment, you can schedule reports only for your current domain.

Before you begin

- For reports other than risk reports: Create a report template. See Report Templates, on page 2051 for more information.

- If you want to distribute email reports using the scheduler, configure a mail relay host and specify report recipients and message information. See Configuring a Mail Relay Host and Notification Address, on page 906 and (for reports other than risk reports) Distributing Reports by Email at Generation Time, on page 2071 or (for risk reports) Generating, Viewing, and Printing Risk Reports, on page 2050.
- (Optional) Set or change the file name, output format, time window, or email distribution settings of the scheduled report. See Specify Report Generation Settings for a Scheduled Report, on page 186.

Procedure

Step 1 Select **System** > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From the **Job Type** list, select **Report**.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 180 for details.

Step 5 Type a name in the **Job Name** field.

Step 6 In the **Report Template** field, select a risk report or report template.

Step 7 If you want to comment on the task, type a comment in the **Comment** field.

The comment field appears in the Tasks Details section of the schedule calendar page; keep comments brief.

Step 8 If you want to email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field. You must have a valid email relay server configured to send status messages.

Note Configuring this option does **not** distribute the reports.

Step 9 If you do not want to receive report email attachments when reports have no data (for example, when no events of a certain type occurred during the report period), select the **If report is empty, still attach to email** check box.

Step 10 Click **Save**.

Specify Report Generation Settings for a Scheduled Report

You must have Admin or Security Analyst privileges to perform this task.

To specify or change the file name, output format, time window, or email distribution settings of a scheduled report:

Procedure

Step 1 Select **Overview** > **Reporting** > **Report Templates**.

Step 2 Click **Edit** for the report template to change.

Step 3 Click **Generate**.

Note If you want to change report generation settings without generating the report now, you must click **Generate** from the template configuration page. Changes will not be saved if you click **Generate** from the template list view unless you generate the report.

Step 4 Modify settings.

Step 5 To save the new settings without generating the report, click **Cancel**.

To save the new settings and generate the report, click **Generate** and skip the rest of the steps in this procedure.

Step 6 Click **Save**.

Step 7 If you see a prompt to save even though you haven't made changes, click **OK**.

Automating Firepower Recommendations

You can automatically generate rule state recommendations based on network discovery data for your network using the most recently saved configuration settings in a custom intrusion policy.



Note If the system automatically generates scheduled recommendations for an intrusion policy with unsaved changes, you must discard your changes in that policy and commit the policy if you want the policy to reflect the automatically generated recommendations.

When the task runs, the system automatically generates recommended rule states, and modifies the states of intrusion rules based on the configuration of your policy. Modified rule states take effect the next time you deploy your intrusion policy.

In a multidomain deployment, you can automate recommendations for intrusion policies at the current domain level. The system builds a separate network map for each leaf domain. In a multidomain deployment, if you enable this feature in an intrusion policy in an ancestor domain, the system generates recommendations using data from all descendant leaf domains. This can enable intrusion rules tailored to assets that may not exist in all leaf domains, which can affect performance.

Before you begin

- Configure Firepower recommended rules in an intrusion policy as described in Generating and Applying Firepower Recommendations, on page 1520
- If you want to email task status messages, configure a valid email relay server.
- You must have the Threat Smart License or Protection Classic License to generate recommendations.

Procedure

Step 1 Choose **System** > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From **Job Type**, choose **Firepower Recommended Rules**.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 180 for details.

Step 5 Enter a name in the **Job Name** field.

Step 6 Next to **Policies**, choose one or more intrusion policies where you want to generate recommendations. Check **All Policies** check box to choose all intrusion policies.

Step 7 (Optional) Enter a comment in the **Comment** field.

Keep comments brief. Comments appear in the Task Details section of the schedule calendar page.

Step 8 (Optional) To email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field.

Step 9 Click **Save**.

Related Topics

Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461
About Firepower Recommended Rules, on page 1517
Configuring a Mail Relay Host and Notification Address, on page 906

Software Update Automation

You can automatically download and apply most patches and feature releases to the Firepower System.

The tasks you must schedule to install software updates vary depending on whether you are updating the FMC or are using a FMC to update managed devices.



Note Cisco **strongly** recommends that you use your FMCs to update the devices they manage.

- To update the FMC, schedule the software installation using the Install Latest Update task.
- To use a FMC to automate software updates for its managed devices, you must schedule two tasks:
  - Push (copy) the update to managed devices using the Push Latest Update task.
  - Install the update on managed devices using the Install Latest Update task.

When scheduling updates to managed devices, schedule the push and install tasks to happen in succession; you must first push the update to the device before you can install it. To automate software updates on a device group, you must select all the devices within the group. Allow enough time between tasks for the process to complete; schedule tasks at least 30 minutes apart. If you schedule a task to install an update and the update has not finished copying from the FMC to the device, the installation task will not succeed. However, if the scheduled installation task repeats daily, it will install the pushed update when it runs the next day.

Note You must manually upload and install updates in two situations. First, you cannot schedule major updates to the Firepower System. Second, you cannot schedule updates for or pushes from FMC that cannot access the Support Site. If your FMC is not directly connected to the Internet, you should use management interfaces configuration to set up a proxy to allow it to download updates from the Support Site.

Note that a task scheduled to install an update on a device group will install the pushed update to each device within the device group simultaneously. Allow enough time for the scheduled task to complete for each device within the device group.

If you want to have more control over this process, you can use the **Once** option to download and install updates during off-peak hours after you learn that an update has been released.

Related Topics

Management Interfaces, on page 881
System Updates, on page 137

Automating Software Downloads

You can create a scheduled task that automatically downloads the latest software updates from Cisco. You can use this task to schedule download of updates you plan to install manually.

You must be in the global domain to perform this task.

Procedure

Step 1 Select **System** > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From the **Job Type** list, select **Download Latest Update**.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 180 for details.

Step 5 Type a name in the **Job Name** field.

Step 6 Next to **Update Items**, check **Software** check box.

Step 7 If you want to comment on the task, type a comment in the **Comment** field.

The comment field appears in the Task Details section of the schedule calendar page; keep comments brief.

Step 8 If you want to email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field. You must have a valid email relay server configured to send status messages.

Step 9 Click **Save**.

Related Topics

Configuring a Mail Relay Host and Notification Address, on page 906

Automating Software Pushes

If you want to automate the installation of software updates on managed devices, you must push the updates to the devices before installing.

When you create the task to push software updates to managed devices, make sure you allow enough time between the push task and a scheduled install task for the updates to be copied to the device.

You must be in the global domain to perform this task.

Procedure

Step 1 Select **System** > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From the **Job Type** list, select **Push Latest Update**.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 180 for details.

Step 5 Type a name in the **Job Name** field.

Step 6 From the **Device** drop-down list, select the device that you want to update.

Step 7 If you want to comment on the task, type a comment in the **Comment** field.

The comment field appears in the Task Details section of the schedule calendar page; keep comments brief.

Step 8 If you want to email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field. You must have a valid email relay server configured to send status messages.

Step 9 Click **Save**.

Related Topics

Configuring a Mail Relay Host and Notification Address, on page 906

Automating Software Installs

Make sure you allow enough time between the task that pushes the update to a managed device and the task that installs the update.

You must be in the global domain to perform this task.

Caution Depending on the update being installed, the appliance may reboot after the software is installed.

Procedure

Step 1 Select **System** > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From the **Job Type** list, select **Install Latest Update**.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 180 for details.

Step 5 Type a name in the **Job Name** field.

Step 6 From the **Device** drop-down list, select the appliance (including the Firepower Management Center) where you want to install the update.

Step 7 Next to **Update Items**, check the **Software** check box.

Step 8 If you want to comment on the task, type a comment in the **Comment** field.

The comment field appears in the Task Details section of the schedule calendar page; keep comments brief.

Step 9 If you want to email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field. You must have a valid email relay server configured to send status messages.

Step 10 Click **Save**.

Related Topics

Configuring a Mail Relay Host and Notification Address, on page 906

Vulnerability Database Update Automation

Cisco uses vulnerability database (VDB) updates to expand the list of network assets, traffic, and vulnerabilities that the Firepower System recognizes. You can use the scheduling feature to update the VDB, thereby ensuring that you are using the most up-to-date information to evaluate the hosts on your network.

When automating VDB updates, you must automate two separate steps:

- Downloading the VDB update.
- Installing the VDB update.



Caution Installing a vulnerability database (VDB) update *immediately* restarts the Snort process on all managed devices. Additionally, the first deploy after installing the VDB *might* cause a Snort restart depending on the VDB content. In either scenario, the restart interrupts traffic inspection. Whether traffic drops during the interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Allow enough time between tasks for the process to complete. For example, if you schedule a task to install an update and the update has not fully downloaded, the installation task will not succeed. However, if the scheduled installation task repeats daily, it will install the downloaded VDB update when the task runs the next day.

Note:

- You cannot schedule updates for appliances that cannot access the Support Site. If your FMC is not directly connected to the Internet, you should use management interfaces configuration to set up a proxy to allow it to download updates from the Support Site.
- If you want to have more control over this process, you can use the **Once** option to download and install VDB updates during off-peak hours after you learn that an update has been released.
- In multidomain deployments, you can only schedule VDB updates for the Global domain. The changes take effect when you redeploy policies.

Related Topics

Management Interfaces, on page 881

Automating VDB Update Downloads

You must be in the global domain to perform this task.

Procedure

Step 1 Select **System** > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From the **Job Type** list, select **Download Latest Update**.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 180 for details.

Step 5 Type a name in the **Job Name** field.

Step 6 Next to **Update Items**, check the **Vulnerability Database** check box.

Step 7 If you want to comment on the task, type a comment in the **Comment** field.

The comment field appears in the Task Details section of the calendar schedule page; keep comments brief.

Step 8 If you want to email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field. You must have a valid email relay server configured to send status messages.

Step 9 Click **Save**.

Related Topics

Configuring a Mail Relay Host and Notification Address, on page 906

Automating VDB Update Installs

Allow enough time between the task that downloads the VDB update and the task that installs the update.

You must be in the global domain to perform this task.

Caution Installing a vulnerability database (VDB) update *immediately* restarts the Snort process on all managed devices. Additionally, the first deploy after installing the VDB *might* cause a Snort restart depending on the VDB content. In either scenario, the restart interrupts traffic inspection. Whether traffic drops during the interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Procedure

Step 1 Select **System** > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From the **Job Type** list, select **Install Latest Update**.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 180 for details.

Step 5 Type a name in the **Job Name** field.

Step 6 From the **Device** drop-down list, select the FMC.

Step 7 Next to **Update Items**, check the **Vulnerability Database** check box.

Step 8 If you want to comment on the task, type a comment in the **Comment** field.

Tip The comment field appears in the View Tasks section of the page, so you should try to keep it relatively short.

Step 9 If you want to email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field. You must have a valid email relay server configured to send status messages.

Step 10 Click **Save**.

Related Topics

Configuring a Mail Relay Host and Notification Address, on page 906

Automating URL Filtering Updates Using a Scheduled Task

In order to ensure that threat data for URL filtering is current, the system must obtain data updates from the Cisco Collective Security Intelligence (CSI) cloud.

By default, when you enable URL filtering, automatic updates are enabled. However, if you need to control when these updates occur, use the procedure described in this topic instead of the default update mechanism.

Although daily updates tend to be small, if it has been more than five days since your last update, new URL filtering data may take up to 20 minutes to download, depending on your bandwidth. Then, it may take up to 30 minutes to perform the update itself.

Before you begin

- Ensure the Firepower Management Center has internet access; see Security, Internet Access, and Communication Ports, on page 2425.
- Ensure that URL filtering is enabled. See Enable URL Filtering Using Category and Reputation, on page 1235 for more information.
- Verify that **Enable Automatic Updates** is not selected on the **Cisco CSI** under the **System** > **Integration** menu.
- You must be in the global domain to perform this task. You must also have the URL Filtering license.

Procedure

Step 1 Select **System** > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From the **Job Type** list, select **Update URL Filtering Database**.

Step 4 Specify how you want to schedule the update, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 180 for details.

Step 5 Type a name in the **Job Name** field.

Step 6 If you want to comment on the task, type a comment in the **Comment** field.

The comment field appears in the Task Details section of the schedule calendar page; keep comments brief.

Step 7 If you want to email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field. You must have a valid email relay server configured to send status messages.

Step 8 Click **Save**.

Related Topics

Configuring a Mail Relay Host and Notification Address, on page 906

Scheduled Task Review

After adding scheduled tasks, you can view them and evaluate their status. The View Options section of the page allows you to view scheduled tasks using a calendar and a list of scheduled tasks.

The Calendar view option allows you to view which scheduled tasks occur on which day.

The Task List shows a list of tasks along with their status. The task list appears below the calendar when you open the calendar. In addition, you can view it by selecting a date or task from the calendar.

You can edit a scheduled task that you previously created. This feature is especially useful if you want to test a scheduled task once to make sure that the parameters are correct. Later, after the task completes successfully, you can change it to a recurring task.

There are two types of deletions you can perform from the Schedule View page. You can delete a specific one-time task that has not yet run or you can delete every instance of a recurring task. If you delete an instance of a recurring task, all instances of the task are deleted. If you delete a task that is scheduled to run once, only that task is deleted.

Task List Details

Table 22: Task List Columns

Column	Description
Name	Displays the name of the scheduled task and the comment associated with it.
Type	Displays the type of scheduled task.
Start Time	Displays the scheduled start date and time.
Frequency	Displays how often the task is run.
Last Run Time	Displays the actual start date and time. For a recurring task, this applies to the most recent execution.
Last Run Status	Describes the current status for a scheduled task: • A **Check Mark** () indicates that the task ran successfully. • A question mark icon (**Question Mark** ()) indicates that the task is in an unknown state. • An exclamation mark icon () indicates that the task failed. For a recurring task, this applies to the most recent execution.
Next Run Time	Displays the next execution time for a recurring task. Displays N/A for a one-time task.
Creator	Displays the name of the user that created the scheduled task.
Edit	Edits the scheduled task.
Delete	Deletes the scheduled task.

Viewing Scheduled Tasks on the Calendar

In a multidomain deployment, you can view scheduled tasks only for your current domain.

Procedure

Step 1 Select **System** > **Tools** > **Scheduling**.

Step 2 You can perform the following tasks using the calendar view:

- Click **Double Left Arrow** (<<) to move back one year.
- Click **Single Left Arrow** (<) to move back one month.
- Click **Single Right Arrow** (>) to move forward one month.
- Click **Double Right Arrow** (>>) to move forward one year.
- Click **Today** to return to the current month and year.
- Click **Add Task** to schedule a new task.
- Click a date to view all scheduled tasks for the specific date in a task list table below the calendar.
- Click a specific task on a date to view the task in a task list table below the calendar.

Editing Scheduled Tasks

In a multidomain deployment, you can edit scheduled tasks only for your current domain.

Procedure

Step 1 Select **System** > **Tools** > **Scheduling**.

Step 2 On the calendar, click either the task that you want to edit or the day on which the task appears.

Step 3 In the **Task Details** table, click **Edit** ( ) next to the task you want to edit.

Step 4 Edit the task.

Step 5 Click **Save**.

Deleting Scheduled Tasks

In a multidomain deployment, you can delete scheduled tasks only for your current domain.

Procedure

Step 1 Select **System** > **Tools** > **Scheduling**.

Step 2 In the calendar, click the task you want to delete. For a recurring task, click an instance of the task.

Step 3 In the **Task Details** table, click **Delete** ( ), then confirm your choice.



CHAPTER 10

Data Storage

- Data Stored on the FMC, on page 197
- External Data Storage, on page 199

Data Stored on the FMC

For	See
General information about data storage on the FMC	The Disk Usage Widget, on page 264
Purging old data	Purging Data from the FMC Database, on page 198
Allowing external access to the data on the FMC (this is an advanced feature)	External Database Access Settings, on page 877
Backups	Manage Backups and Remote Storage, on page 168 and subtopics
Reports	Configuring Local Storage, on page 890
Events	Connection Logging, on page 2207 Database Event Limits, on page 878 and subtopics
Network discovery data	Network Discovery Data Storage Settings, on page 1934 and subsequent topics
Files	Information about storing files in File Policies and Malware Protection, on page 1401, including best practices. File and Malware Inspection Performance and Storage Tuning, on page 1441
Packet data	Edit General Settings, on page 240
Users and user activity	The Users Database, on page 1831 The User Activity Database, on page 1830

Purging Data from the FMC Database

You can use the database purge page to purge discovery, identity, connection, and Security Intelligence data files from the FMC databases. Note that when you purge a database, the appropriate process is restarted.

Caution Purging a database removes the data you specify from the Firepower Management Center. After the data is deleted, it *cannot* be recovered.

Before you begin

You must have Admin or Security Analyst privileges to purge data. You can be in the global domain only.

Procedure

Step 1 Choose **System** > **Tools** > **Data Purge**.

Step 2 Under **Discovery and Identity**, perform any or all of the following:

- Check the **Network Discovery Events** check box to remove all network discovery events from the database.
- Check the **Hosts** check box to remove all hosts and Host Indications of Compromise flags from the database.
- Check the **User Activity** check box to remove all user activity events from the database.
- Check the **User Identities** check box to remove all user login and user history data from the database, as well as User Indications of Compromise flags.

Step 3 Under **Connections**, perform any or all of the following:

- Check the **Connection Events** check box to remove all connection data from the database.
- Check the **Connection Summary Events** check box to remove all connection summary data from the database.
- Check the **Security Intelligence Events** check box to remove all Security Intelligence data from the database.

Note Checking the **Connection Events** check box does not remove Security Intelligence events. Connections with Security Intelligence data will still appear in the Security Intelligence event page (available under the Analysis > Connections menu). Correspondingly, checking the **Security Intelligence Events** check box does not remove connection events with associated Security Intelligence data.

Step 4 Click **Purge Selected Events**.
The items are purged and the appropriate processes are restarted.

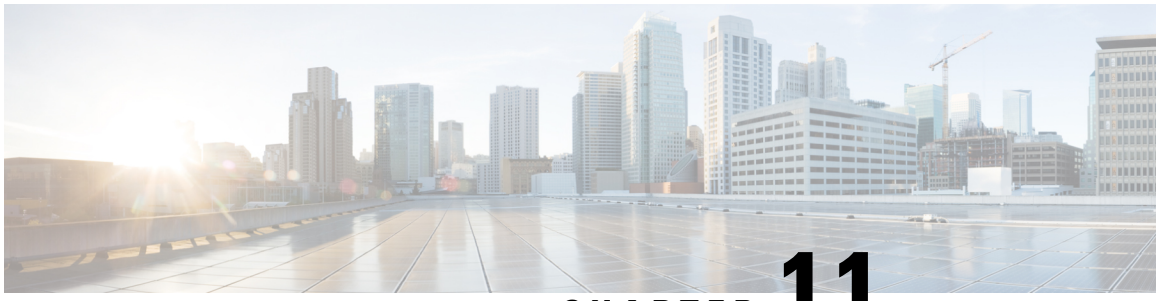
External Data Storage

You can optionally use remote data storage for store certain types of data.

For	See
Backups	Manage Backups and Remote Storage, on page 168 and subtopics Remote Storage Management, on page 890 and subtopics
Reports	Remote Storage Management, on page 890 and subtopics Moving Reports to Remote Storage, on page 2073



Important If you will use syslog or store events externally, avoid special characters in object names such as policy and rule names. Object names should not contain special characters, such as commas, that the receiving application may use as separators.



CHAPTER 11

Firepower Management Center High Availability

The following topics describe how to configure Active/Standby high availability of Cisco Firepower Management Centers:

- About Firepower Management Center High Availability, on page 201
- Requirements for Firepower Management Center High Availability, on page 205
- Prerequisites for Firepower Management Center High Availability, on page 207
- Establishing Firepower Management Center High Availability, on page 208
- Viewing Firepower Management Center High Availability Status, on page 209
- Configurations Synced on Firepower Management Center High Availability Pairs, on page 210
- Configuring External Access to the FMC Database in a High Availability Pair, on page 211
- Using CLI to Resolve Device Registration in Firepower Management Center High Availability, on page 211
- Switching Peers in a Firepower Management Center High Availability Pair, on page 212
- Pausing Communication Between Paired Firepower Management Centers, on page 212
- Restarting Communication Between Paired Firepower Management Centers, on page 212
- Changing the IP address of a Firepower Management Center in a High Availability Pair, on page 213
- Disabling Firepower Management Center High Availability, on page 213
- Replacing FMCs in a High Availability Pair, on page 214

About Firepower Management Center High Availability

To ensure the continuity of operations, the high availability feature allows you to designate redundant Firepower Management Centers to manage devices. Firepower Management Centers support Active/Standby high availability where one appliance is the active unit and manages devices. The standby unit does not actively manage devices. The active unit writes configuration data into a data store and replicates data for both units, using synchronization where necessary to share some information with the standby unit.

Active/Standby high availability lets you configure a secondary Firepower Management Center to take over the functionality of a primary Firepower Management Center if the primary fails. When the primary Firepower Management Center fails, you must promote the secondary Firepower Management Center to become the active unit.

Event data streams from managed devices to both Firepower Management Centers in the high availability pair. If one Firepower Management Center fails, you can monitor your network without interruption using the other Firepower Management Center.

Note that Firepower Management Centers configured as a high availability pair do not need to be on the same trusted management network, nor do they have to be in the same geographic location.

Caution Because the system restricts some functionality to the active Firepower Management Center, if that appliance fails, you must promote the standby Firepower Management Center to active.

Roles v. Status in Firepower Management Center High Availability

Primary/Secondary Roles

When setting up Firepower Management Centers in a high availability pair, you configure one Firepower Management Center to be primary and the other as secondary. During configuration, the primary unit's policies are synchronized to the secondary unit. After this synchronization, the primary Firepower Management Center becomes the active peer, while the secondary Firepower Management Center becomes the standby peer, and the two units act as a single appliance for managed device and policy configuration.

Active/Standby Status

The main differences between the two Firepower Management Centers in a high availability pair are related to which peer is active and which peer is standby. The active Firepower Management Center remains fully functional, where you can manage devices and policies. On the standby Firepower Management Center, functionality is hidden; you cannot make any configuration changes.

Event Processing on Firepower Management Center High Availability Pairs

Since both Firepower Management Centers in a high availability pair receive events from managed devices, the management IP addresses for the appliances are not shared. This means that you do not need to intervene to ensure continuous processing of events if a Firepower Management Center fails.

AMP Cloud Connections and Malware Information

Although they share file policies and related configurations, Firepower Management Centers in a high availability pair share neither Cisco AMP cloud connections nor malware dispositions. To ensure continuity of operations, and to ensure that detected files' malware dispositions are the same on both Firepower Management Centers, both primary and secondary Firepower Management Centers must have access to the AMP cloud.

URL Filtering and Security Intelligence

URL filtering and Security Intelligence configurations and information are synchronized between Firepower Management Centers in a high availability deployment. However, only the primary Firepower Management Center downloads URL category and reputation data for updates to Security Intelligence feeds.

If the primary Firepower Management Center fails, not only must you make sure that the secondary Firepower Management Center can access the internet to update threat intelligence data, but you must also use the web interface on the secondary Firepower Management Center to promote it to active.

User Data Processing During Firepower Management Center Failover

If the primary Firepower Management Center fails, the Secondary Firepower Management Center propagates to managed devices user-to-IP mappings from the and user agent identity source; and propagates SGT mappings from the ISE identity source. Users not yet seen by identity sources are identified as Unknown.

After the downtime, the Unknown users are re identified and processed according to the rules in your identity policy.

Configuration Management on Firepower Management Center High Availability Pairs

In a high availability deployment, only the active Firepower Management Center can manage devices and apply policies. Both Firepower Management Centers remain in a state of continuous synchronization.

If the active Firepower Management Center fails, the high availability pair enters a degraded state until you manually promote the standby appliance to the active state. Once the promotion is complete, the appliances leave maintenance mode.

Firepower Management Center High Availability Behavior During a Backup

When you perform a Backup on a Firepower Management Center high availability pair, the Backup operation pauses synchronization between the peers. During this operation, you may continue using the active Firepower Management Center, but not the standby peer.

After Backup is completed, synchronization resumes, which briefly disables processes on the active peer. During this pause, the High Availability page briefly displays a holding page until all processes resume.

Firepower Management Center High Availability Split-Brain

If the active Firepower Management Center in a high-availability pair goes down (due to power issues, network/connectivity issues), you can promote the standby Firepower Management Center to an active state. When the original active peer comes up, both peers can assume they are active. This state is defined as 'split-brain'. When this situation occurs, the system prompts you to choose an active appliance, which demotes the other appliance to standby.

If the active Firepower Management Center goes down (or disconnects due to a network failure), you may either break high availability or switch roles. The standby Firepower Management Center enters a degraded state.



Note Whichever appliance you use as the secondary loses all of its device registrations and policy configurations when you resolve split-brain. For example, you would lose modifications to any policies that existed on the secondary but not on the primary. If the Firepower Management Center is in a high availability split-brain scenario where both appliances are active, and you register managed devices and deploy policies before you resolve split-brain, you must export any policies and unregister any managed devices from the intended standby Firepower Management Center before re-establishing high availability. You may then register the managed devices and import the policies to the intended active Firepower Management Center.

Upgrading Firepower Management Centers in a High Availability Pair

Cisco electronically distributes several different types of updates periodically. These include major and minor upgrades to the system software. You may need to install these updates on Firepower Management Centers in a high availability setup.

Warning Make sure that there is at least one operational Firepower Management Center during an upgrade.

Before you begin

Read the release notes or advisory text that accompanies the upgrade. The release notes provide important information, including supported platforms, compatibility, prerequisites, warnings, and specific installation and uninstallation instructions.

Procedure

Step 1 Access the web interface of the active Firepower Management Center and pause data synchronization; see Pausing Communication Between Paired Firepower Management Centers, on page 212.

Step 2 Upgrade the standby Firepower Management Center; see the upgrade guide.
When the upgrade completes, the standby unit becomes active. When both peers are active, the high availability pair is in a degraded state (split-brain).

Step 3 Upgrade the other Firepower Management Center.

Step 4 Decide which Firepower Management Center you want to use as the standby. Any additional devices or policies added to the standby after pausing synchronization are not synced to the active Firepower Management Center. Unregister only those additional devices and export any configurations you want to preserve.

When you choose a new active Firepower Management Center, the Firepower Management Center you designate as secondary will lose device registrations and deployed policy configurations, which are not synced.

Step 5 Resolve split-brain by choosing the new active Firepower Management Center which has all the latest required configurations for policies and devices.

Troubleshooting Firepower Management Center High Availability

This section lists troubleshooting information for some common Firepower Management Center high availability operation errors.

Error	Description	Solution
500 Internal	May appear when attempting to access the web interface while performing critical Firepower Management Center high availability operations, including switching peer roles or pausing and resuming synchronization.	Wait until the operation completes before using the web interface.

Error	Description	Solution
System processes are starting, please wait Also, the web interface does not respond.	May appear when the Firepower Management Center reboots (manually or while recovering from a power down) during a high availability or data synchronization operation.	1. Access the Firepower Management Center shell and use the `manage_hadc.pl` command to access the Firepower Management Center high availability configuration utility. **Note** Run the utility as a root user, using `sudo`. 2. Pause mirroring operations by using option 5. Reload the Firepower Management Center web interface. 3. Use the web interface to resume synchronization. Choose **System > Integration**, then click the **High Availability** tab and choose **Resume Synchronization**.

Requirements for Firepower Management Center High Availability

Model Support

See Hardware Requirements, on page 205.

Supported Domains

Global

User Roles

Admin

Hardware Requirements

- Supported hardware models:

  MC1000, MC1500, MC2000, MC2500, MC3500, MC4000, MC4500

- The two Firepower Management Centers in a high availability configuration must be the same model.
- The primary Firepower Management Center backup must not be restored to the secondary Firepower Management Center.

- Bandwidth requirement for high availability configuration depends on various factors such as the size of the network, the number of managed devices, the volume of events and logs, and the size and frequency of configuration updates. For a typical Firepower Management Center high availability deployment, a minimum of 5 Mbps network bandwidth between the peers is recommended.
- The two Firepower Management Centers in a high availability configuration may be physically and geographically separated from each other in different data centers.
- See also License Requirements for FMC High Availability Configurations, on page 206.

Software Requirements

Access the **Appliance Information** widget to verify the software version, the intrusion rule update version and the vulnerability database update. By default, the widget appears on the **Status** tab of the **Detailed Dashboard** and the**Summary Dashboard**. For more information, see The Appliance Information Widget, on page 258

- The two Firepower Management Centers in a high availability configuration must have the same major (first number), minor (second number), and maintenance (third number) software version.
- The two Firepower Management Centers in a high availability configuration must have the same version of the intrusion rule update installed.
- The two Firepower Management Centers in a high availability configuration must have the same version of the vulnerability database update installed.
- The two Firepower Management Centers in a high availability configuration must have the same version of the LSP (Lightweight Security Package) installed.



Warning If the software versions, intrusion rule update versions and vulnerability database update versions are not identical on both Firepower Management Centers, you cannot establish high availability.

License Requirements for FMC High Availability Configurations

All Licensing Types

No special license is required for Firepower Management Center hardware appliances in a high availability pair.

A device managed with Firepower Management Center hardware appliances in a high availability configuration requires the same number of feature licenses and subscriptions as a device managed by a single Firepower Management Center hardware appliance.

The system automatically replicates all feature licenses from active to standby Firepower Management Center when the high-availability pair is formed, and updates license changes during ongoing data synchronization, so the licenses are available on failover.

Smart Licensing

Each FTD device requires the same licenses whether managed by a single FMC or by FMCs in a high availability pair.

Example: If you want to enable advanced malware protection for two Firepower Threat Defense devices managed by a Firepower Management Center pair, buy two Malware licenses and two TM subscriptions, register the active Firepower Management Center with the Cisco Smart Software Manager, then assign the licenses to the two Firepower Threat Defense devices on the active Firepower Management Center.

Only the active Firepower Management Center is registered with Cisco Smart Software Manager. When failover occurs, the system communicates with Cisco Smart Software Manager to release the Smart License entitlements from the originally-active Firepower Management Center and assign them to the newly-active Firepower Management Center.

Classic Licensing

Each device requires the same licenses whether managed by a single FMC or by FMCs in a high availability pair.

Example: If you want to enable advanced malware protection for two devices managed by a Firepower Management Center pair, buy two Malware licenses and two TAM subscriptions, add those licenses to the Firepower Management Center, then assign the licenses to the two devices on the active Firepower Management Center.

Prerequisites for Firepower Management Center High Availability

Before establishing a Firepower Management Center high availability pair:

- Export required policies from the intended secondary Firepower Management Center to the intended primary Firepower Management Center. For more information, see Exporting Configurations, on page 175.
- Make sure that the intended secondary Firepower Management Center does not have any devices added to it. Delete devices from the intended secondary Firepower Management Center and register these devices to the intended primary Firepower Management Center. For more information see Delete a Device from the FMC, on page 231 and Add a Device to the FMC, on page 229.
- Import the policies into the intended primary Firepower Management Center. For more information, see Importing Configurations, on page 176.
- On the intended primary Firepower Management Center, verify the imported policies, edit them as needed and deploy them to the appropriate device. For more information, see Deploy Configuration Changes, on page 346.
- On the intended primary Firepower Management Center, associate the appropriate licenses to the newly added devices. For more information see Assign Licenses to Managed Devices from the Device Management Page, on page 134.

You can now proceed to establish high availability. For more information, see Establishing Firepower Management Center High Availability, on page 208.

Establishing Firepower Management Center High Availability

Establishing high availability can take a significant amount of time, even several hours, depending on the bandwidth between the peers and the number of policies. It also depends on the number of devices registered to the active Firepower Management Center, which need to be synced to the standby Firepower Management Center. You can view the High Availability page to check the status of the high availability peers.

Before you begin

- Confirm that both the Firepower Management Centers adhere to the high availability system requirements. For more information , see Requirements for Firepower Management Center High Availability, on page 205.
- Confirm that you completed the prerequisites for establishing high availability. For more information, see Prerequisites for Firepower Management Center High Availability, on page 207.

Procedure

Step 1 Log into the Firepower Management Center that you want to designate as the secondary.

Step 2 Choose **System** > **Integration**.

Step 3 Choose **High Availability**.

Step 4 Under Role for this Firepower Management Center, choose **Secondary**.

Step 5 Enter the hostname or IP address of the primary Firepower Management Center in the **Primary Firepower Management Center Host** text box.

You can leave this empty if the primary Firepower Management Center does not have an IP address reachable from the peer FMC (which can be public or private IP address). In this case, use both the **Registration Key** and the **Unique NAT ID** fields. You need to specify the IP address of at least one FMC to enable HA connection.

Step 6 Enter a one-time-use registration key in the **Registration Key** text box.

The registration key is any user-defined alphanumeric value up to 37 characters in length. This registration key will be used to register both -the secondary and the primary Firepower Management Centers.

Step 7 If you did not specify the primary IP address, or if you do not plan to specify the secondary IP address on the primary Firepower Management Center, then in the **Unique NAT ID** field, enter a unique alphanumeric ID. See NAT Environments, on page 883 for more information.

Step 8 Click **Register**.

Step 9 Using an account with Admin access, log into the Firepower Management Center that you want to designate as the primary.

Step 10 Choose **System** > **Integration**.

Step 11 Choose **High Availability**.

Step 12 Under Role for this Firepower Management Center, choose **Primary**.

Step 13 Enter the hostname or IP address of the secondary Firepower Management Center in the **Secondary Firepower Management Center Host** text box.

You can leave this empty if the secondary Firepower Management Center does not have an IP address reachable from the peer FMC (which can be public or private IP address). In this case, use both the **Registration Key** and the **Unique NAT ID** fields. You need to specify the IP address of at least one FMC to enable HA connection.

Step 14 Enter the same one-time-use registration key in the **Registration Key** text box you used in step 6.

Step 15 If required, enter the same NAT ID that you used in step 7 in the **Unique NAT ID** text box.

Step 16 Click **Register**.

What to do next

After establishing a Firepower Management Center high availability pair, devices registered to the active Firepower Management Center are automatically registered to the standby Firepower Management Center.

Note When a registered device has a NAT IP address, automatic device registration fails and the secondary Firepower Management Center High Availablity page lists the device as local, pending. You can then assign a different NAT IP address to the device on the standby Firepower Management Center High Availability page. If automatic registration otherwise fails on the standby Firepower Management Center, but the device appears to be registered to the active Firepower Management Center, see Using CLI to Resolve Device Registration in Firepower Management Center High Availability, on page 211.

Viewing Firepower Management Center High Availability Status

After you identify your active and standby Firepower Management Centers, you can view information about the local Firepower Management Center and its peer.

Note In this context, Local Peer refers to the appliance where you are viewing the system status. Remote Peer refers to the other appliance, regardless of active or standby status.

Procedure

Step 1 Log into one of the Firepower Management Centers that you paired using high availability.

Step 2 Choose **System** > **Integration**.

Step 3 Choose **High Availability**.

You can view:

Summary Information

- The health status of the high availability pair. The status of a correctly functioning system will oscillate between "Healthy" and "Synchronization task is in progress" as the standby unit receives configuration changes from the active unit.
- The current synchronization status of the high availability pair

- The IP address of the active peer and the last time it was synchronized
- The IP address of the standby peer and the last time it was synchronized

System Status

- The IP addresses for both peers
- The operating system for both peers
- The software version for both peers
- The appliance model of both peers

Configurations Synced on Firepower Management Center High Availability Pairs

When you establish high availability between two Firepower Management Centers, the following configuration data is synced between them:

- License entitlements
- Access control policies
- Intrusion rules
- Malware and file policies
- DNS policies
- Identity policies
- SSL policies
- Prefilter policies
- Network discovery rules
- Application detectors
- Correlation policy rules
- Alerts
- Scanners
- Response groups
- Remediation settings, although you must install custom modules on both Firepower Management Centers. For more information on remediation settings, see Managing Remediation Modules, on page 2043.

Configuring External Access to the FMC Database in a High Availability Pair

In a high availability setup, we recommend you to use only the active peer to configure the external access to the database. When you configure the standby peer for external database access, it leads to frequent disconnections. To restore the connectivity, you must Pausing Communication Between Paired Firepower Management Centers and Restarting Communication Between Paired Firepower Management Centers the synchronization of the standby peer. For information on how to enable external database access to Firepower Management Centers, see Enabling External Access to the Database, on page 878 .

Using CLI to Resolve Device Registration in Firepower Management Center High Availability

If automatic device registration fails on the standby Firepower Management Center, but appears to be registered to the active Firepower Management Center, complete the following steps:

Warning If you do an RMA of Secondary Firepower Management Center or add a Secondary Firepower Management Center, the managed FTDs are unregistered and as a result, their configuration may be deleted.

Procedure

Step 1 Unregister the device from the active Firepower Management Center. See *Delete (Unregister) a Device from the Firepower Management Center* in Cisco Secure Firewall Management Center Device Configuration Guide.

Step 2 Log in to the CLI for the affected device.

Step 3 Run the CLI command: **configure manager delete**.

This command disables and removes the current Firepower Management Center.

Step 4 Run the CLI command: **configure manager add**.

This command configures the device to initiate a connection to a Firepower Management Center.

Tip Configure remote management on the device, only for the active Firepower Management Center. When you establish high availability, the devices are automatically registered to the standby Firepower Management Center.

Step 5 Log in to the active Firepower Management Center and register the device.

Switching Peers in a Firepower Management Center High Availability Pair

Because the system restricts some functionality to the active Firepower Management Center, if that appliance fails, you must promote the standby Firepower Management Center to active:

Procedure

Step 1 Log into one of the Firepower Management Centers that you paired using high availability.

Step 2 Choose **System** > **Integration**.

Step 3 Choose **High Availability**.

Step 4 Choose **Switch Peer Roles** to change the local role from Active to Standby, or Standby to Active. With the Primary or Secondary designation unchanged, the roles are switched between the two peers.

Pausing Communication Between Paired Firepower Management Centers

If you want to temporarily disable high availability, you can disable the communications channel between the Firepower Management Centers. If you pause synchronization on the active peer, you can resume synchronization on either the standby or active peer. However, if you pause synchronization on the standby peer, you only can resume synchronization on the standby peer.

Procedure

Step 1 Log into one of the Firepower Management Centers that you paired using high availability.

Step 2 Choose **System** > **Integration**.

Step 3 Choose **High Availability**.

Step 4 Choose **Pause Synchronization**.

Restarting Communication Between Paired Firepower Management Centers

If you temporarily disable high availability, you can restart high availability by enabling the communications channel between the Firepower Management Centers. If you paused synchronization on the active unit, you can resume synchronization on either the standby or active unit. However, if you paused synchronization on the standby unit, you only can resume synchronization on the standby unit.

Procedure

Step 1 Log into one of the Firepower Management Centers that you paired using high availability.
Step 2 Choose **System** > **Integration**.
Step 3 Choose **High Availability**.
Step 4 Choose **Resume Synchronization**.

Changing the IP address of a Firepower Management Center in a High Availability Pair

Note If you landed on this topic while trying to edit remote management on a 7000 and 8000 Series managed device, see Editing Remote Management on a Managed Device, on page 493.

If the IP address for one of the high availability peers changes, high availability enters a degraded state. To recover high availability, you must manually change the IP address.

Procedure

Step 1 Log into one of the Firepower Management Centers that you paired using high availability.
Step 2 Choose **System** > **Integration**.
Step 3 Choose **High Availability**.
Step 4 Choose **Peer Manager**.
Step 5 Choose **Edit** ( ).
Step 6 Enter the display name of the appliance, which is used only within the context of the Firepower System.

Entering a different display name does not change the host name for the appliance.

Step 7 Enter the fully qualified domain name or the name that resolves through the local DNS to a valid IP address (that is, the host name), or the host IP address.
Step 8 Click **Save**.

Disabling Firepower Management Center High Availability

Procedure

Step 1 Log into one of the Firepower Management Centers in the high availability pair.

Step 2 Choose **System** > **Integration**.

Step 3 Choose **High Availability**.

Step 4 Choose **Break High Availability**.

Step 5 Choose one of the following options for handling managed devices:

- To control all managed devices with this Firepower Management Center, choose **Manage registered devices from this console**. All devices will be unregistered from the peer.
- To control all managed devices with the other Firepower Management Center, choose **Manage registered devices from peer console**. All devices will be unregistered from this Firepower Management Center.
- To stop managing devices altogether, choose **Stop managing registered devices from both consoles**. All devices will be unregistered from both Firepower Management Centers.

Note If you choose to manage the registered devices from the secondary Firepower Management Center, the devices will be unregistered from the primary Firepower Management Center. The devices are now registered to be managed by the secondary Firepower Management Center. However the licenses that were applied to these devices are deregistered on account of the high availability break operation. You must now proceed to re-register (enable) the licenses on the devices from the secondary Firepower Management Center. For more information see Move or Remove Licenses from FTD Devices, on page 121.

Step 6 Click **OK**.

Replacing FMCs in a High Availability Pair

If you need to replace a failed unit in a Firepower Management Center high availability pair, you must follow one of the procedures listed below. The table lists four possible failure scenarios and their corresponding replacement procedures.

Failure Status	Data Backup Status	Replacement Procedure
Primary FMC failed	Data backup successful	Replace a Failed Primary FMC (Successful Backup), on page 215
	Data backup not successful	Replace a Failed Primary FMC (Unsuccessful Backup), on page 216
Secondary FMC failed	Data backup successful	Replace a Failed Secondary FMC (Successful Backup), on page 217
	Data backup not successful	Replace a Failed Secondary FMC (Unsuccessful Backup), on page 218

Replace a Failed Primary FMC (Successful Backup)

Two Firepower Management Centers, FMC1 and FMC2, are part of a high availability pair. FMC1 is the primary and FMC2 is the secondary. This task describes the steps to replace a failed primary Firepower Management Center, FMC1, when data backup from the primary is successful.

Before you begin

Verify that the data backup from the failed primary Firepower Management Center is successful.

Procedure

Step 1 Contact Support to request a replacement for a failed Firepower Management Center - FMC1.

Step 2 When the primary Firepower Management Center - FMC1 fails, access the web interface of the secondary Firepower Management Center - FMC2 and switch peers. For more information, see Switching Peers in a Firepower Management Center High Availability Pair, on page 212.

This promotes the secondary Firepower Management Center - FMC2 to active.

You can use FMC2 as the active Firepower Management Center until the primary Firepower Management Center - FMC1 is replaced.

Caution Do not break Firepower Management Center High Availability from FMC2, since licenses that were synced to FMC2 from FMC1 (before failure), will be removed from FMC2 and you will be unable to perform any deploy actions from FMC2.

Step 3 Reimage the replacement Firepower Management Center with the same software version as FMC1.

Step 4 Restore the data backup retrieved from FMC1 to the new Firepower Management Center.

Step 5 Install required Firepower Management Center patches, geolocation database (GeoDB) updates, vulnerability database (VDB) updates and system software updates to match FMC2.

The new Firepower Management Center and FMC2 will now both be active peers, resulting in a high availability split-brain.

Step 6 When the Firepower Management Center web interface prompts you to choose an active appliance, select FMC2 as active.

This syncs the latest configuration from FMC2 to the newFirepower Management Center - FMC1.

Step 7 When the configuration syncs successfully, access the web interface of the secondary Firepower Management Center - FMC2 and switch roles to make the primaryFirepower Management Center - FMC1 active. For more information, see Switching Peers in a Firepower Management Center High Availability Pair, on page 212.

Step 8 Apply Classic licenses received with the new Firepower Management Center - FMC1 and delete the old licenses. For more information, see Generate a Classic License and Add It to the Firepower Management Center, on page 131.

Smart licenses work seamlessly.

What to do next

High availability has now been re-established and the primary and the secondary Firepower Management Centers will now work as expected.

Replace a Failed Primary FMC (Unsuccessful Backup)

Two Firepower Management Centers - FMC1 and FMC2 are part of a high availability pair. FMC1 is the primary and FMC2 is the secondary. This task describes the steps to replace a failed primary Firepower Management Center -FMC1 when data backup from the primary is unsuccessful.

Procedure

Step 1 Contact Support to request a replacement for a failed Firepower Management Center - FMC1.

Step 2 When the primary Firepower Management Center - FMC1 fails, access the web interface of the secondary Firepower Management Center - FMC2 and switch peers. For more information, see Switching Peers in a Firepower Management Center High Availability Pair, on page 212.

This promotes the secondary Firepower Management Center - FMC2 to active.

You can use FMC2 as the active Firepower Management Center until the primary Firepower Management Center - FMC1 is replaced.

Caution Do not break Firepower Management Center High Availability from FMC2, since classic and smart licenses that were synced to FMC2 from FMC1 (before failure), will be removed from FMC2 and you will be unable to perform any deploy actions from FMC2.

Step 3 Reimage the replacement Firepower Management Center with the same software version as FMC1.

Step 4 Install required Firepower Management Center patches, geolocation database (GeoDB) updates, vulnerability database (VDB) updates and system software updates to match FMC2.

Step 5 Deregister the Firepower Management Center - FMC2 from the Cisco Smart Software Manager. For more information, see Deregister a Firepower Management Center from the Cisco Smart Software Manager, on page 122.

Deregistering a Firepower Management Center from the Cisco Smart Software Manager removes the Management Center from your virtual account. All license entitlements associated with the Firepower Management Center release back to your virtual account. After deregistration, the Firepower Management Center enters Enforcement mode where no update or changes on licensed features are allowed.

Step 6 Access the web interface of the secondary Firepower Management Center - FMC2 and break Firepower Management Center high availability. For more information, see Disabling Firepower Management Center High Availability, on page 213. When prompted to select an option for handling managed devices, choose **Manage registered devices from this console**.

As a result, classic and smart licenses that were synced to the secondary Firepower Management Center-FMC2, will be removed and you cannot perform deployment activities from FMC2.

Step 7 Re-establish Firepower Management Center high availability, by setting up the Firepower Management Center - FMC2 as the primary and Firepower Management Center - FMC1 as the secondary. For more information , see Establishing Firepower Management Center High Availability, on page 208.

Step 8 Apply Classic licenses received with the new Firepower Management Center - FMC1 and delete the old licenses. For more information, see Generate a Classic License and Add It to the Firepower Management Center, on page 131.

Step 9 Register a Smart License to the primary Firepower Management Center - FMC2. For more information see Register Smart Licenses, on page 116.

What to do next

High availability has now been re-established and the primary and the secondary Firepower Management Centers will now work as expected.

Replace a Failed Secondary FMC (Successful Backup)

Two Firepower Management Centers - FMC1 and FMC2 are part of a high availability pair. FMC1 is the primary and FMC2 is the secondary. This task describes the steps to replace a failed secondary Firepower Management Center -FMC2 when data backup from the secondary is successful.

Before you begin

Verify that the data backup from the failed secondary Firepower Management Center is successful.

Procedure

Step 1 Contact Support to request a replacement for a failed Firepower Management Center - FMC2.

Step 2 Continue to use the primary Firepower Management Center - FMC1 as the active Firepower Management Center.

Step 3 Reimage the replacement Firepower Management Center with the same software version as FMC2.

Step 4 Restore the data backup from FMC2 to the new Firepower Management Center.

Step 5 Install required Firepower Management Center patches, geolocation database (GeoDB) updates, vulnerability database (VDB) updates and system software updates to match FMC1.

Step 6 Resume data synchronization (if paused) from the web interface of the new Firepower Management Center - FMC2, to synchronize the latest configuration from the primary Firepower Management Center - FMC1. For more information, see Restarting Communication Between Paired Firepower Management Centers, on page 212.

Classic and Smart Licenses work seamlessly.

What to do next

High availability has now been re-established and the primary and the secondary Firepower Management Centers will now work as expected.

Replace a Failed Secondary FMC (Unsuccessful Backup)

Two Firepower Management Centers - FMC1 and FMC2 are part of a high availability pair. FMC1 is the primary and FMC2 is the secondary. This task describes the steps to replace a failed secondary Firepower Management Center -FMC2 when data backup from the secondary is unsuccessful.

Procedure

Step 1 Contact Support to request a replacement for a failed Firepower Management Center - FMC2.

Step 2 Continue to use the primary Firepower Management Center - FMC1 as the active Firepower Management Center.

Step 3 Reimage the replacement Firepower Management Center with the same software version as FMC2.

Step 4 Install required Firepower Management Center patches, geolocation database (GeoDB) updates, vulnerability database (VDB) updates and system software updates to match FMC1.

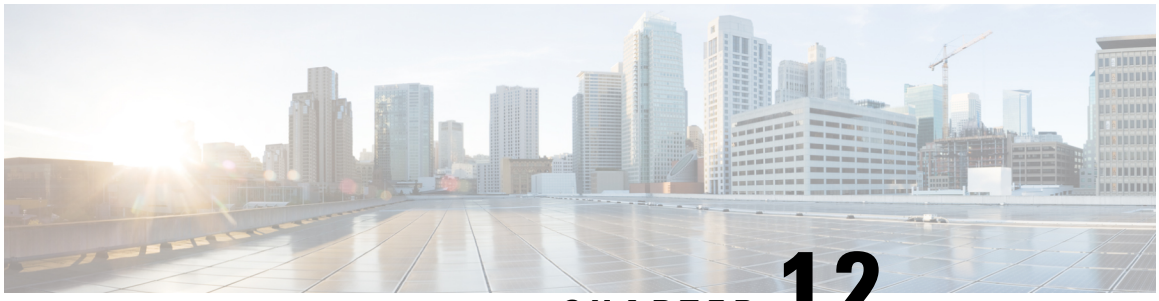
Step 5 Access the web interface of the primary Firepower Management Center - FMC1 and break Firepower Management Center high availability. For more information, see Disabling Firepower Management Center High Availability, on page 213. When prompted to select an option for handling managed devices, choose **Manage registered devices from this console**.

Step 6 Re-establish Firepower Management Center high availability, by setting up the Firepower Management Center - FMC1 as the primary and Firepower Management Center - FMC2 as the secondary. For more information , see Establishing Firepower Management Center High Availability, on page 208.

- When high availability is successfully established, the latest configuration from the primary Firepower Management Center - FMC1 is synchronized to the secondary Firepower Management Center - FMC2.
- Classic and Smart Licenses work seamlessly.

What to do next

High availability has now been re-established and the primary and the secondary Firepower Management Centers will now work as expected.



CHAPTER 12

Device Management Basics

The following topics describe how to manage devices in the Firepower System:

- About Device Management, on page 219
- Requirements and Prerequisites for Device Management, on page 226
- Complete the FTD Initial Configuration Using the CLI, on page 226
- Add a Device to the FMC, on page 229
- Delete a Device from the FMC, on page 231
- Add a Device Group, on page 232
- Configure Device Settings, on page 233
- Change the Manager for the Device, on page 244
- Viewing Device Information, on page 248

About Device Management

Use the Firepower Management Center to manage your devices.

About the Firepower Management Center and Device Management

When the Firepower Management Center manages a device, it sets up a two-way, SSL-encrypted communication channel between itself and the device. The Firepower Management Center uses this channel to send information to the device about how you want to analyze and manage your network traffic to the device. As the device evaluates the traffic, it generates events and sends them to the Firepower Management Center using the same channel.

By using the Firepower Management Center to manage devices, you can:

- configure policies for all your devices from a single location, making it easier to change configurations
- install various types of software updates on devices
- push health policies to your managed devices and monitor their health status from the Firepower Management Center

The Firepower Management Center aggregates and correlates intrusion events, network discovery information, and device performance data, allowing you to monitor the information that your devices are reporting in relation to one another, and to assess the overall activity occurring on your network.

You can use a Firepower Management Center to manage nearly every aspect of a device's behavior.

Note Although a Firepower Management Center can manage devices running certain previous releases as specified in the compatibility matrix available at http://www.cisco.com/c/en/us/support/security/defense-center/products-device-support-tables-list.html, new features are not available to these previous-release devices.

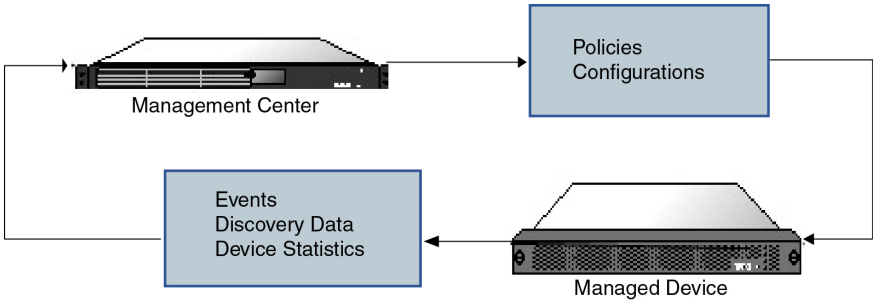
What Can Be Managed by a Firepower Management Center?

You can use the Firepower Management Center as a central management point in a Firepower System deployment to manage the following devices:

- 7000 and 8000 Series devices
- ASA FirePOWER modules
- NGIPSv devices
- Firepower Threat Defense (physical hardware and virtual)

When you manage a device, information is transmitted between the Firepower Management Center and the device over a secure, SSL-encrypted TCP tunnel.

The following illustration lists what is transmitted between a Firepower Management Center and its managed devices. Note that the types of events and policies that are sent between the appliances are based on the device type.



Beyond Policies and Events

In addition to deploying policies to devices and receiving events from them, you can also perform other device-related tasks on the Firepower Management Center.

Backing Up a Device

You **cannot** create or restore backup files for NGIPSv devices or ASA FirePOWER modules.

When you perform a backup of a physical managed device from the device itself, you back up the device configuration **only**. To back up configuration data and, optionally, unified files, perform a backup of the device using the managing Firepower Management Center.

To back up event data, perform a backup of the managing Firepower Management Center.

Updating Devices

From time to time, Cisco releases updates to the Firepower System, including:

- intrusion rule updates, which may contain new and updated intrusion rules
- vulnerability database (VDB) updates
- geolocation updates
- software patches and updates

You can use the Firepower Management Center to install an update on the devices it manages.

About Device Management Interfaces

Each device includes a single dedicated Management interface for communicating with the FMC.

You can perform initial setup on the management interface, or on the console port.

Management interfaces are also used to communicate with the Smart Licensing server, to download updates, and to perform other management functions.

Management Interfaces on

When you set up your device, you specify the FMC IP address that you want to connect to. Both management and event traffic go to this address at initial registration. Note: In some situations, the FMC might establish the *initial* connection on a different management interface; subsequent connections should use the management interface with the specified IP address.

If the FMC has a separate event-only interface, the managed device sends subsequent event traffic is sent to the FMC event-only interface if the network allows. In addition, some managed-device models include an additional management interface that you can configure for event-only traffic. If the event network goes down, then event traffic reverts to the regular management interfaces on the FMC and/or on the managed device.

Management Interface Support Per Device Model

See the hardware installation guide for your model for the management interface locations.



Note For the Firepower 4100/9300 chassis, the MGMT interface is for *chassis* management, not for Firepower Threat Defense logical device management. You must configure a separate NIC interface to be of type mgmt (and/or firepower-eventing), and then assign it to the Firepower Threat Defense logical device.



Note For Firepower Threat Defense on any chassis, the physical management interface is shared between the Diagnostic logical interface, which is useful for SNMP or syslog, and is configured along with data interfaces in the FMC, and the Management logical interface for FMC communication.

See the following table for supported management interfaces on each managed device model.

Table 23: Management Interface Support on Managed Devices

Model	Management Interface	Optional Event Interface
7000 series	eth0	No support
8000 series	eth0	eth1
ASA FirePOWER services module on the ASA 5585-X	eth0 **Note** eth0 is the internal name of the Management 1/0 interface.	eth1 **Note** eth1 is the internal name of the Management 1/1 interface.
Firepower Threat Defense on the Firepower 4100 and 9300	management0 **Note** management0 is the internal name of this interface, regardless of the physical interface ID.	management1 **Note** management1 is the internal name of this interface, regardless of the physical interface ID.
Firepower Threat Defense on the ASA 5506-X, 5508-X, or 5516-X	br1 **Note** br1 is the internal name of the Management 1/1 interface.	No support
Firepower Threat Defense on the 5512-X through 5555-X	br1 **Note** br1 is the internal name of the Management 0/0 interface.	No support
Firepower Threat Defense Virtual	br1	No support

Network Routes on Device Management Interfaces

Management interfaces (including event-only interfaces) support only static routes to reach remote networks. When you set up your managed device, the setup process creates a default route to the gateway IP address that you specify. You cannot delete this route; you can only modify the gateway address.



Note The routing for management interfaces is completely separate from routing that you configure for data interfaces.

The default route always uses the lowest-numbered management interface (e.g. management0).

At least one static route is recommended per management interface to access remote networks. We recommend placing each interface on a separate network to avoid potential routing problems, including routing problems from other devices to the FTD. If you do not experience problems with interfaces on the same network, then be sure to configure static routes correctly. For example, both management0 and management1 are on the same network, but the FMC management and event interfaces are on different networks. The gateway is 192.168.45.1. If you want management1 to connect to the FMC's event-only interface at 10.6.6.1/24, you can create a static route for 10.6.6.0/24 through management1 with the same gateway of 192.168.45.1. Traffic to 10.6.6.0/24 will hit this route before it hits the default route, so management1 will be used as expected.

Another example includes separate management and event-only interfaces on both the FMC and the managed device. The event-only interfaces are on a separate network from the management interfaces. In this case, add a static route through the event-only interface for traffic destined for the remote event-only network, and vice versa.

NAT Environments

Network address translation (NAT) is a method of transmitting and receiving network traffic through a router that involves reassigning the source or destination IP address. The most common use for NAT is to allow private networks to communicate with the internet. Static NAT performs a 1:1 translation, which does not pose a problem for FMC communication with devices, but port address translation (PAT) is more common. PAT lets you use a single public IP address and unique ports to access the public network; these ports are dynamically assigned as needed, so you cannot initiate a connection to a device behind a PAT router.

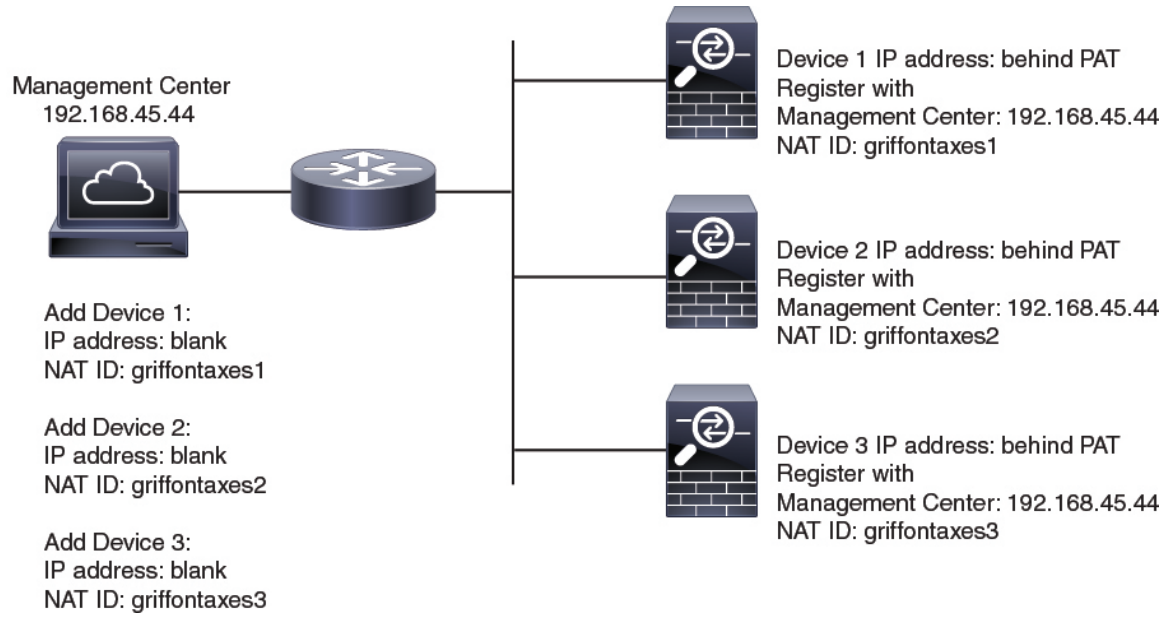
Normally, you need both IP addresses (along with a registration key) for both routing purposes and for authentication: the FMC specifies the device IP address when you add a device, and the device specifies the FMC IP address. However, if you only know one of the IP addresses, which is the minimum requirement for routing purposes, then you must also specify a unique NAT ID on both sides of the connection to establish trust for the initial communication and to look up the correct registration key. The FMC and device use the registration key and NAT ID (instead of IP addresses) to authenticate and authorize for initial registration.

For example, you add a device to the FMC, and you do not know the device IP address (for example, the device is behind a PAT router), so you specify only the NAT ID and the registration key on the FMC; leave the IP address blank. On the device, you specify the FMC IP address, the same NAT ID, and the same registration key. The device registers to the FMC's IP address. At this point, the FMC uses the NAT ID instead of IP address to authenticate the device.

Although the use of a NAT ID is most common for NAT environments, you might choose to use the NAT ID to simplify adding many devices to the FMC. On the FMC, specify a unique NAT ID for each device you want to add while leaving the IP address blank, and then on each device, specify both the FMC IP address and the NAT ID. Note: The NAT ID must be unique per device.

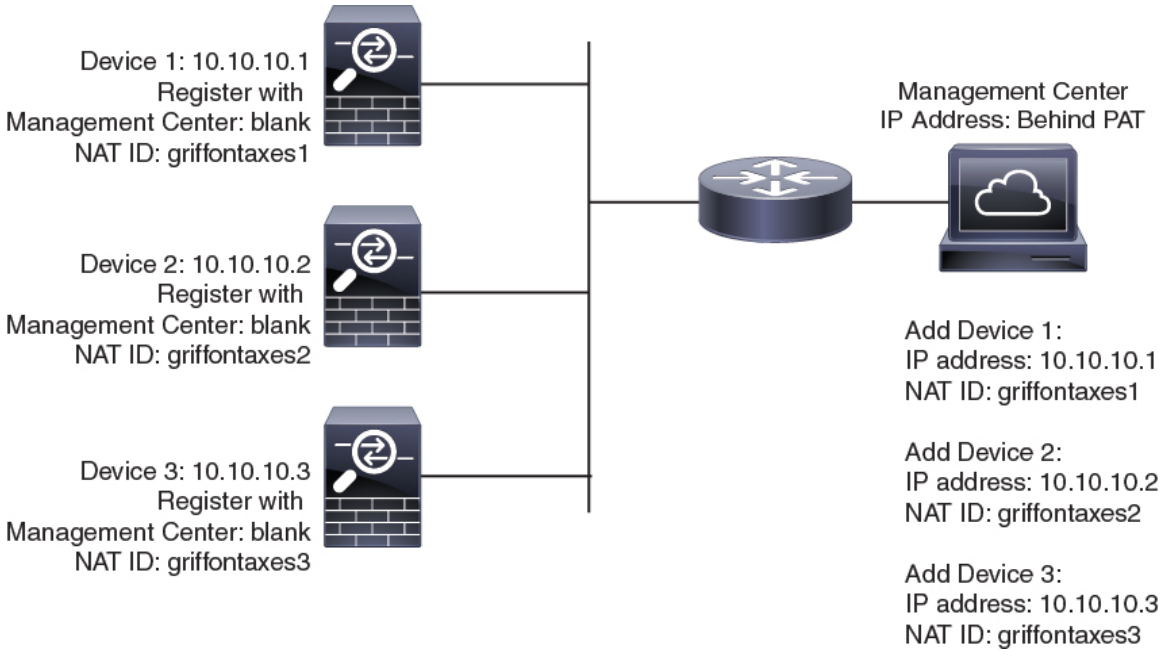
The following example shows three devices behind a PAT IP address. In this case, specify a unique NAT ID per device on both the FMC and the devices, and specify the FMC IP address on the devices.

Figure 1: NAT ID for Managed Devices Behind PAT



The following example shows the FMC behind a PAT IP address. In this case, specify a unique NAT ID per device on both the FMC and the devices, and specify the device IP addresses on the FMC.

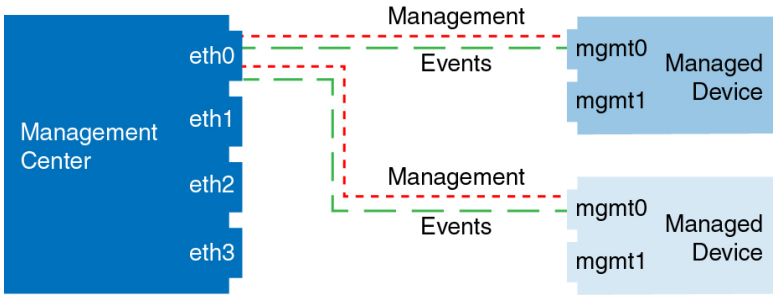
Figure 2: NAT ID for FMC Behind PAT



Management and Event Traffic Channel Examples

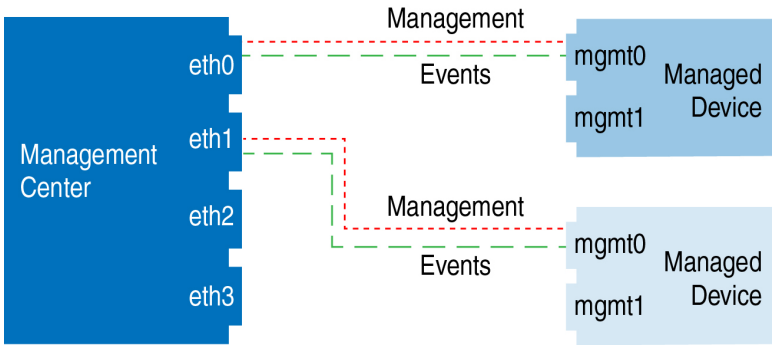
The following example shows the Firepower Management Center and managed devices using only the default management interfaces.

Figure 3: Single Management Interface on the Firepower Management Center



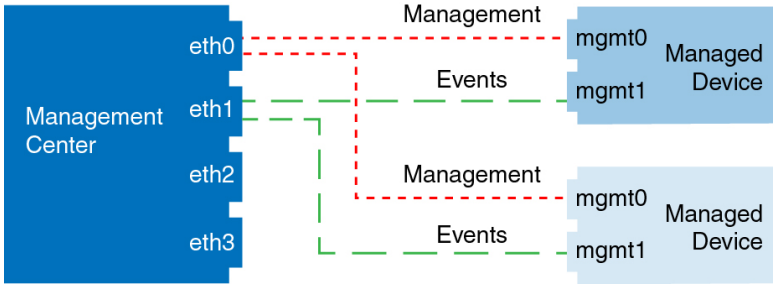
The following example shows the Firepower Management Center using separate management interfaces for devices; and each managed device using 1 management interface.

Figure 4: Mutliple Management Interfaces on the Firepower Management Center



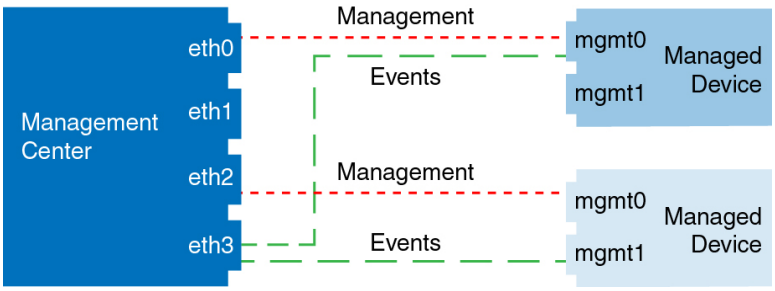
The following example shows the Firepower Management Center and managed devices using a separate event interface.

Figure 5: Separate Event Interface on the Firepower Management Center and Managed Devices



The following example shows a mix of multiple management interfaces and a separate event interface on the Firepower Management Center and a mix of managed devices using a separate event interface, or using a single management interface.

Figure 6: Mixed Management and Event Interface Usage



Requirements and Prerequisites for Device Management

Model Support

Any managed device; unless noted in the procedure.

Supported Domains

The domain in which the device resides.

User Roles

- Admin
- Network Admin

Complete the FTD Initial Configuration Using the CLI

Connect to the FTD CLI to perform initial setup, including setting the Management IP address, gateway, and other basic networking settings using the setup wizard. The dedicated Management interface is a special interface with its own network settings. You will also configure FMC communication settings. You can only configure the Management interface settings; you must configure data interface settings in FMC.

Before you begin

This procedure applies to all FTD devices except for the Firepower 4100/9300.

Procedure

Step 1 Connect to the FTD CLI, either from the console port or using SSH to the Management interface, which obtains an IP address from a DHCP server by default. If you intend to change the network settings, we recommend using the console port so you do not get disconnected.

Step 2 Log in with the username **admin** and the password **Admin123**.

Note If the password was already changed, and you do not know it, you must reimage the device to reset the password to the default.

Step 3 The first time you log in to FTD, you are prompted to accept the End User License Agreement (EULA) and to change the admin password. You are then presented with the CLI setup script.

Note You cannot repeat the CLI setup wizard unless you clear the configuration; for example, by reimaging. However, all of these settings can be changed later at the CLI using **configure network** commands. See the FTD command reference.

Defaults or previously entered values appear in brackets. To accept previously entered values, press **Enter**.

See the following guidelines:

- **Enter the IPv4 default gateway for the management interface**—The **data-interfaces** setting applies only to Firepower Device Manager management; you should set a gateway IP address for Management 1/1 when using FMC. In the edge deployment example shown in the network deployment section, the inside interface acts as the management gateway. In this case, you should set the gateway IP address to be the *intended* inside interface IP address; you must later use FMC to set the inside IP address.
- **If your networking information has changed, you will need to reconnect**—If you are connected with SSH but you change the IP address at initial setup, you will be disconnected. Reconnect with the new IP address and password. Console connections are not affected. Note also that the DHCP server on Management will be disabled if you change the IP address.
- **Manage the device locally?**—Enter **no** to use FMC. A **yes** answer means you will use Firepower Device Manager instead. Note also that the DHCP server on Management 1/1 will be disabled if it wasn't already.
- **Configure firewall mode?**—We recommend that you set the firewall mode at initial configuration. Changing the firewall mode after initial setup erases your running configuration.

Example:

```
You must accept the EULA to continue.
Press <ENTER> to display the EULA:
End User License Agreement
[...]

Please enter 'YES' or press <ENTER> to AGREE to the EULA:

System initialization in progress.  Please stand by.
You must change the password for 'admin' to continue.
Enter new password: ********
Confirm new password: ********
You must configure the network to continue.
You must configure at least one of IPv4 or IPv6.
Do you want to configure IPv4? (y/n) [y]:
Do you want to configure IPv6? (y/n) [n]:
Configure IPv4 via DHCP or manually? (dhcp/manual) [manual]:
Enter an IPv4 address for the management interface [192.168.45.45]: 10.10.10.15
Enter an IPv4 netmask for the management interface [255.255.255.0]: 255.255.255.192
Enter the IPv4 default gateway for the management interface [data-interfaces]: 10.10.10.1
Enter a fully qualified hostname for this system [firepower]: ftd-1.cisco.com
Enter a comma-separated list of DNS servers or 'none' [208.67.222.222,208.67.220.220]:
Enter a comma-separated list of search domains or 'none' []:
If your networking information has changed, you will need to reconnect.
DHCP Server Disabled
The DHCP server has been disabled. You may re-enable with configure network ipv4
dhcp-server-enable
For HTTP Proxy configuration, run 'configure network http-proxy'

Manage the device locally? (yes/no) [yes]: no
DHCP Server Disabled
```

```
Configure firewall mode? (routed/transparent) [routed]:
Configuring firewall mode ...

Update policy deployment information
    - add device configuration
    - add network discovery
    - add system policy

You can register the sensor to a Firepower Management Center and use the
Firepower Management Center to manage it. Note that registering the sensor
to a Firepower Management Center disables on-sensor Firepower Services
management capabilities.

When registering the sensor to a Firepower Management Center, a unique
alphanumeric registration key is always required.  In most cases, to register
a sensor to a Firepower Management Center, you must provide the hostname or
the IP address along with the registration key.
'configure manager add [hostname | ip address ] [registration key ]'

However, if the sensor and the Firepower Management Center are separated by a
NAT device, you must enter a unique NAT ID, along with the unique registration
key.
'configure manager add DONTRESOLVE [registration key ] [ NAT ID ]'

Later, using the web interface on the Firepower Management Center, you must
use the same registration key and, if necessary, the same NAT ID when you add
this sensor to the Firepower Management Center.
>
```

Step 4 Identify the FMC that will manage this FTD.

configure manager add {*hostname* | *IPv4_address* | *IPv6_address* | **DONTRESOLVE**} *reg_key* [*nat_id*]

- {*hostname* | *IPv4_address* | *IPv6_address* | **DONTRESOLVE**}—Specifies either the FQDN or IP address of the FMC. If the FMC is not directly addressable, use **DONTRESOLVE** and also specify the *nat_id*. At least one of the devices, either the FMC or the FTD, must have a reachable IP address to establish the two-way, SSL-encrypted communication channel between the two devices. If you specify **DONTRESOLVE** in this command, then the FTD must have a reachable IP address or hostname.
- *reg_key*—Specifies a one-time registration key of your choice that you will also specify on the FMC when you register the FTD. The registration key must not exceed 37 characters. Valid characters include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-).
- *nat_id*—Specifies a unique, one-time string of your choice that you will also specify on the FMC when you register the FTD when one side does not specify a reachable IP address or hostname. It is required if you set the FMC to **DONTRESOLVE**. The NAT ID must not exceed 37 characters. Valid characters include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-). This ID cannot be used for any other devices registering to the FMC.

Example:

```
> configure manager add MC.example.com 123456
Manager successfully configured.
```

If the FMC is behind a NAT device, enter a unique NAT ID along with the registration key, and specify DONTRESOLVE instead of the hostname, for example:

Example:

```
> configure manager add DONTRESOLVE regk3y78 natid90
Manager successfully configured.
```

If the FTD is behind a NAT device, enter a unique NAT ID along with the FMC IP address or hostname, for example:

Example:

```
> configure manager add 10.70.45.5 regk3y78 natid56
Manager successfully configured.
```

What to do next

Register your device to a FMC.

Add a Device to the FMC

Use this procedure to add a single device to the FMC. If you plan to link devices for redundancy or performance, you must still use this procedure, keeping in mind the following points:

- 8000 Series stacks—Use this procedure to add each device to the Firepower Management Center, then establish the stack; see Establishing Device Stacks, on page 536.
- 7000 and 8000 Series high availability—Use this procedure to add each device to the Firepower Management Center, then establish high availability; see Establishing Firepower 7000/8000 Series High Availability, on page 521. For high availability stacks, first stack the devices, then establish high availability between the stacks.
- Firepower Threat Defense high availability—Use this procedure to add each device to the Firepower Management Center, then establish high availability; see Add a Firepower Threat Defense High Availability Pair, on page 699.
- Firepower Threat Defense clusters—Make sure cluster units are in a successfully formed cluster on FXOS, then use this procedure to add each unit to the Firepower Management Center as a separate managed device. Finally, cluster the units on the Firepower Management Center. For more information, see FMC: Add a Cluster, on page 727.

Note If you have established or will establish FMC high availability, add devices *only* to the active (or intended active) FMC. When you establish high availability, devices registered to the active FMC are automatically registered to the standby.

Before you begin

- Set up the device to be managed by the FMC. See:
  - Firepower Threat Defense devices: Complete the FTD Initial Configuration Using the CLI, on page 226

  - 7000 and 8000 Series devices: Configuring Remote Management on a Managed Device, on page 492
  - Other device types: The getting started guide for your model

- If you are adding an FTD device, the FMC must be registered for Smart Licensing. A valid evaluation license is sufficient, but if it expires, you will not be able to add new devices until you successfully register.
- If you registered a FMC and a device using IPv4 and want to convert them to IPv6, you must delete and reregister the device.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 From the **Add** drop-down menu, choose **Device**.

Step 3 In the **Host** field, enter the IP address or the hostname of the device you want to add.

The hostname of the device is the fully qualified domain name or the name that resolves through the local DNS to a valid IP address. Use a hostname rather than an IP address if your network uses DHCP to assign IP addresses.

In a NAT environment, you may not need to specify the IP address or hostname of the device, if you already specified the IP address or hostname of the FMC when you configured the device to be managed by the FMC. For more information, see NAT Environments, on page 223.

Step 4 In the **Display Name** field, enter a name for the device as you want it to display in the FMC.

Step 5 In the **Registration Key** field, enter the same registration key that you used when you configured the device to be managed by the FMC. The registration key is a one-time-use shared secret. The key can include alphanumeric characters and hyphens (-).

Step 6 In a multidomain deployment, regardless of your current domain, assign the device to a leaf **Domain**.

If your current domain is a leaf domain, the device is automatically added to the current domain. If your current domain is not a leaf domain, post-registration, you must switch to the leaf domain to configure the device.

Step 7 (Optional) Add the device to a device **Group**.

Step 8 Choose an initial **Access Control Policy** to deploy to the device upon registration, or create a new policy.

If the device is incompatible with the policy you choose, deploying will fail. This incompatibility could occur for multiple reasons, including licensing mismatches, model restrictions, passive vs inline issues, and other misconfigurations. After you resolve the issue that caused the failure, manually deploy configurations to the device.

Step 9 Choose licenses to apply to the device.

If you registered the FMC to use Smart Licensing, then this dialog box only shows available Smart Licenses.

Smart Licensing

Assign the Smart Licenses you need for the features you want to deploy:

- **Malware** (if you intend to use AMP malware inspection)

- **Threat** (if you intend to use intrusion prevention)
- **URL** (if you intend to implement category-based URL filtering)

Note You can apply an AnyConnect remote access VPN license after you add the device, from the **System** > **Licenses** > **Smart Licenses** page.

Classic Licensing

If you registered the FMC to use Smart Licensing, then this dialog box only shows available Smart Licenses. For classic licenses, go to the **Devices** > **Device Management** > **Device** > **License** area to assign licenses.

- Control, Malware, and URL Filtering licenses require a Protection license.
- VPN licenses require a 7000 or 8000 Series device.
- Control licenses are supported on NGIPSv andASA FirePOWER devices, but do *not* allow you to configure 8000 Series fastpath rules, switching, routing, stacking, or device high availability.

Step 10 If you used a NAT ID during device setup, expand in the **Advanced** section and enter the same NAT ID in the **Unique NAT ID** field. The NAT ID can include alphanumeric characters and hyphens (-).

Step 11 Check the **Transfer Packets** check box to allow the device to transfer packets to the Firepower Management Center.

This option is enabled by default. When events like IPS or Snort are triggered with this option enabled, the device sends event metadata information and packet data to the FMC for inspection. If you disable it, only event information will be sent to the FMC but packet data is not sent.

Step 12 Click **Register**.

It may take up to two minutes for the FMC to verify the device's heartbeat and establish communication. If the registration succeeds, the device is added to the list. If it fails, you will see an error message. If the device fails to register, check the following items:

- Ping—Access the device CLI, and ping the FMC IP address using the following command:

  ping system *ip_address*

  If the ping is not successful, check your network settings using the **show network** command. If you need to change the device IP address, use the **configure network** {**ipv4** | **ipv6**} **manual** command.

- Registration key, NAT ID, and FMC IP address—Make sure you are using the same registration key, and if used, NAT ID, on both devices. You can set the registration key and NAT ID on the device using the **configure manager add** command.

For more troubleshooting information, see https://cisco.com/go/fmc-reg-error.

Delete a Device from the FMC

If you no longer want to manage a device, you can delete it from the FMC. Deleting a device:

- Severs all communication between the FMC and the device.

- Removes the device from the Device Management page.
- Returns the device to local time management if the device is configured using the platform settings policy to receive time from the FMC using NTP.

After deleting the device from the FMC:

- The FTD continues to process the traffic after you delete it from the FMC.
  - Policies, such as NAT and VPN, ACLs, and the interface configurations remain intact.
- Registering the FTD again to the same or a different FMC, the FTD configuration is removed from the FTD.
  - The ACLs that are selected during registration replace the earlier ACLs and the interface configuration remains intact.
- To manage the device later, re-add it to the FMC.



Note When a device is deleted and then re-added, the FMC web interface prompts you to re-apply your access control policies. However, there is no option to re-apply the NAT and VPN policies during registration. Any previously applied NAT or VPN configuration will be removed during registration and must be re-applied after registration is complete.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to delete, click **Delete** ().

Step 3 Confirm that you want to delete the device.

Add a Device Group

The Firepower Management Center allows you to group devices so you can easily deploy policies and install updates on multiple devices. You can expand and collapse the list of devices in the group.

In a multidomain deployment, you can create device groups within a leaf domain only. When you configure a Firepower Management Center for multitenancy, existing device groups are removed; you can re-add them at the leaf domain level.

If you add the primary device in a stack or a high-availability pair to a group, both devices are added to the group. If you unstack the devices or break the high-availability pair, both devices remain in that group.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 From the **Add** drop-down menu, choose **Add Group**.

To edit an existing group, click **Edit** ( ) for the group you want to edit.

Step 3 Enter a **Name**.

Step 4 Under **Available Devices**, choose one or more devices to add to the device group. Use Ctrl or Shift while clicking to choose multiple devices.

Step 5 Click **Add** to include the devices you chose in the device group.

Step 6 Optionally, to remove a device from the device group, click **Delete** ( ) next to the device you want to remove.

Step 7 Click **OK** to add the device group.

Configure Device Settings

After you add a device, you can configure some settings on the device's **Device** page.

Managing System Shut Down

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any except ASA FirePOWER	Leaf only	Admin/Network Admin

Note You cannot shut down or restart the ASA FirePOWER with the Firepower System user interface. See the ASA documentation for more information on how to shut down the respective devices.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device that you want to restart, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Device**.

Tip For stacked devices, you shut down or restart an individual device on the Devices page of the appliance editor.

Step 4 To shut down the device, click **Shut Down Device** ( ) in the **System** section.

Step 5 When prompted, confirm that you want to shut down the device.

Step 6 To restart the device, click **Restart Device** ( ).

Step 7 When prompted, confirm that you want to restart the device.

Edit Management Settings

You can edit management settings in the **Management** area.

Update the Hostname or IP Address in FMC

If you edit the hostname or IP address of a device after you added it to the FMC (using the device's CLI, for example), you need to use the procedure below to manually update the hostname or IP address on the managing FMC.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to modify management options, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Device**, and view the **Management** area.

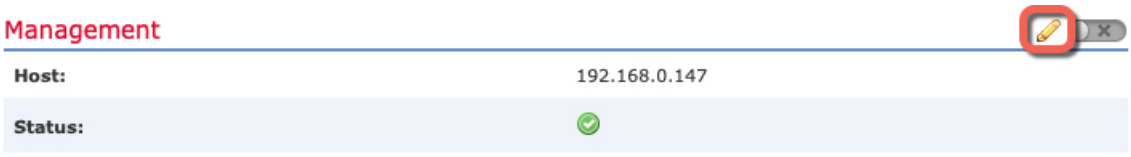
Tip For stacked devices, you modify management options on an individual device on the **Device** page of the appliance editor.

Step 4 Disable management temporarily by clicking the slider so it is disabled ( ).

You are prompted to proceed with disabling management; click **Yes**.

Disabling management blocks the connection between the Firepower Management Center and the device, but does **not** delete the device from the Firepower Management Center.

Step 5 Edit the **Host** IP address or hostname by clicking **Edit** ( ).



Step 6 In the **Management** dialog box, modify the name or IP address in the **Host** field, and click **Save**.

Step 7 Reenable management by clicking the slider so it is enabled ( ).

Modify Management Interfaces at the CLI

Modify the management interface settings on the managed device using the CLI. Many of these settings are ones that you set when you performed the initial setup; this procedure lets you change those settings, and set additional settings such as enabling an event interface if your model supports it, or adding static routes.

For information about the Firepower Threat Defense CLI, see the FTD command reference.

For information about the classic device CLI, see Classic Device Command Line Reference, on page 2431 in this guide.

The Firepower Threat Defense and classic devices use the same commands for management interface configuration. Other commands may differ between the platforms.

Note When using SSH, be careful when making changes to the management interface; if you cannot re-connect because of a configuration error, you will need to access the device console port.

Note If you change the device management IP address, then see the following tasks for FMC connectivity depending on how you identified the FMC during initial device setup using the **configure manager add** command (see Identify a New FMC, on page 245):

- **IP address—No action.** If you identified the FMC using a reachable IP address, then the management connection will be reestablished automatically after several minutes. We recommend that you also change the device IP address shown in FMC to keep the information in sync; see Update the Hostname or IP Address in FMC, on page 234. This action can help the connection reestablish faster. **Note**: If you specified an unreachable FMC IP address, then see the procedure for NAT ID below.
- **NAT ID only—Manually reestablish the connection.** If you identified the FMC using only the NAT ID, then the connection cannot be automatically reestablished. In this case, change the device management IP address in FMC according to Update the Hostname or IP Address in FMC, on page 234.

Before you begin

- For Firepower Threat Defense devices, you can create user accounts that can log into the CLI using the **configure user add** command.
- For the 7000 & 8000 Series devices, you can create user accounts at the web interface as described in Creating a User Account, on page 62.

Procedure

Step 1 Connect to the device CLI, either from the console port or using SSH.

Step 2 Log in with the Admin username and password.

Step 3 Enable an event-only interface (for supported models; see Management Interface Support Per Device Model, on page 221).

configure network management-interface enable *management_interface*

configure network management-interface disable-management-channel *management_interface*

Example:

This example is for a Firepower 4100 or 9300 device; valid interface names differ by device type.

```
> configure network management-interface enable management1
Configuration updated successfully

> configure network management-interface disable-management-channel management1
Configuration updated successfully

>
```

The Firepower Management Center event-only interface cannot accept management channel traffic, so you should simply disable the management channel on the device event interface.

You can optionally disable events for the management interface using the **configure network management-interface disable-events-channel** command. In either case, the device will try to send events on the event-only interface, and if that interface is down, it will send events on the management interface even if you disable the event channel.

You cannot disable both event and management channels on an interface.

Step 4 Configure the network settings of the management interface and/or event interface:

If you do not specify the *management_interface* argument, then you change the network settings for the default management interface. When configuring an event interface, be sure to specify the *management_interface* argument. The event interface can be on a separate network from the management interface, or on the same network. If you are connected to the interface you are configuring, you will be disconnected. You can re-connect to the new IP address.

a) Configure the IPv4 address:

- Manual configuration:

  configure network ipv4 manual *ip_address netmask gateway_ip* [*management_interface*]

  Note that the *gateway_ip* in this command is used to create the default route for the device. If you configure an event-only interface, then you must enter the *gateway_ip* as part of the command; however, this entry just configures the default route to the value you specify and does not create a separate static route for the eventing interface. If you are using an event-only interface on a different network from the management interface, we recommend that you set the *gateway_ip* for use with the management interface, and then create a static route separately for the event-only interface using the **configure network static-routes** command.

  Example:

  ```
  > configure network ipv4 manual 10.10.10.45 255.255.255.0 10.10.10.1 management1
  Setting IPv4 network configuration.
  Network settings changed.

  >
  ```

- DHCP (supported on the default management interface only):

  configure network ipv4 dhcp

b) Configure the IPv6 address:

- Stateless autoconfiguration:

  configure network ipv6 router [*management_interface*]

  Example:

```
> configure network ipv6 router management0
Setting IPv6 network configuration.
Network settings changed.

>
```

- Manual configuration:

  configure network ipv6 manual *ip6_address ip6_prefix_length* [*ip6_gateway_ip*] [*management_interface*]

  Note that the *ipv6_gateway_ip* in this command is used to create the default route for the device. If you configure an event-only interface, then you must enter the *ipv6_gateway_ip* as part of the command; however, this entry just configures the default route to the value you specify and does not create a separate static route for the eventing interface. If you are using an event-only interface on a different network from the management interface, we recommend that you set the *ipv6_gateway_ip* for use with the management interface, and then create a static route separately for the event-only interface using the **configure network static-routes** command.

  Example:

```
> configure network ipv6 manual 2001:0DB8:BA98::3210 64 management1
Setting IPv6 network configuration.
Network settings changed.

>
```

- DHCPv6 (supported on the default management interface only):

  configure network ipv6 dhcp

Step 5 (Firepower Threat Defense only) Enable a DHCP server on the default management interface to provide IP addresses to connected hosts:

configure network ipv4 dhcp-server-enable *start_ip_address end_ip_address*

Example:

```
> configure network ipv4 dhcp-server-enable 10.10.10.200 10.10.10.254
DHCP Server Enabled

>
```

You can only configure a DHCP server when you set the management interface IP address manually. This command is not supported on the Firepower Threat Defense Virtual. To display the status of the DHCP server, enter **show network-dhcp-server**:

```
> show network-dhcp-server
```

```
DHCP Server Enabled
10.10.10.200-10.10.10.254
```

Step 6 Add a static route for the event-only interface if the Firepower Management Center is on a remote network; otherwise, all traffic will match the default route through the management interface.

configure network static-routes {**ipv4** | **ipv6**} **add** *management_interface destination_ip netmask_or_prefix gateway_ip*

For the *default* route, do not use this command; you can only change the default route gateway IP address when you use the **configure network ipv4** or **ipv6** commands (see step 4).

For information about routing, see Network Routes on Device Management Interfaces, on page 222.

Example:

```
> configure network static-routes ipv4 add management1 192.168.6.0 255.255.255.0 10.10.10.1
Configuration updated successfully

> configure network static-routes ipv6 add management1 2001:0DB8:AA89::5110 64
2001:0DB8:BA98::3211
Configuration updated successfully

>
```

To display static routes, enter **show network-static-routes** (the default route is not shown):

```
> show network-static-routes
---------------[ IPv4 Static Routes ]---------------
Interface                : management1
Destination              : 192.168.6.0
Gateway                  : 10.10.10.1
Netmask                  : 255.255.255.0
[…]
```

Step 7 Set the hostname:

configure network hostname *name*

Example:

```
> configure network hostname farscape1.cisco.com
```

Syslog messages do not reflect a new hostname until after a reboot.

Step 8 Set the search domains:

configure network dns searchdomains *domain_list*

Example:

```
> configure network dns searchdomains example.com,cisco.com
```

Set the search domain(s) for the device, separated by commas. These domains are added to hostnames when you do not specify a fully-qualified domain name in a command, for example, **ping system**. The domains are used only on the management interface, or for commands that go through the management interface.

Step 9 Set up to 3 DNS servers, separated by commas:

configure network dns servers *dns_ip_list*

Example:

```
> configure network dns servers 10.10.6.5,10.20.89.2,10.80.54.3
```

Step 10 Set the remote management port for communication with the FMC:

configure network management-interface tcpport *number*

Example:

```
> configure network management-interface tcpport 8555
```

The FMC and managed devices communicate using a two-way, SSL-encrypted communication channel, which by default is on port 8305.

Note Cisco **strongly** recommends that you keep the default settings for the remote management port, but if the management port conflicts with other communications on your network, you can choose a different port. If you change the management port, you must change it for **all** devices in your deployment that need to communicate with each other.

Step 11 Configure an HTTP proxy. The device is configured to directly-connect to the internet on ports TCP/443 (HTTPS) and TCP/80 (HTTP). You can use a proxy server, to which you can authenticate via HTTP Digest. After issuing the command, you are prompted for the HTTP proxy address and port, whether proxy authentication is required, and if it is required, the proxy username, proxy password, and confirmation of the proxy password.

configure network http-proxy

Example:

```
> configure network http-proxy
Manual proxy configuration
Enter HTTP Proxy address: 10.100.10.10
Enter HTTP Proxy Port: 80
Use Proxy Authentication? (y/n) [n]: Y
Enter Proxy Username: proxyuser
Enter Proxy Password: proxypassword
Confirm Proxy Password: proxypassword
```

Step 12 If you change the device management IP address, then see the following tasks for FMC connectivity depending on how you identified the FMC during initial device setup using the **configure manager add** command (see Identify a New FMC, on page 245):

- **IP address—No action.** If you identified the FMC using a reachable IP address, then the management connection will be reestablished automatically after several minutes. We recommend that you also change the device IP address shown in FMC to keep the information in sync; see Update the Hostname or IP Address in FMC, on page 234. This action can help the connection reestablish faster. **Note**: If you specified an unreachable FMC IP address, then you must manually reestablish the connection using Update the Hostname or IP Address in FMC, on page 234.

- **NAT ID only—Manually reestablish the connection.** If you identified the FMC using only the NAT ID, then the connection cannot be automatically reestablished. In this case, change the device management IP address in FMC according to Update the Hostname or IP Address in FMC, on page 234.

Edit General Settings

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Device**.

Step 4 In the **General** section, click **Edit** ( ).

Step 5 Enter a **Name** for the managed device.

Tip For stacked devices, you edit the assigned device name for the stack on the Stack page of the appliance editor. You can edit the assigned device name for an individual device on the Devices page of the appliance editor.

Step 6 Change the **Transfer Packets** setting:

- Check the check box to allow packet data to be stored with events on the Firepower Management Center.
- Clear the check box to prevent the managed device from sending packet data with the events.

Step 7 Click **Force Deploy** to force deployment of current policies and device configuration to the device.

Step 8 Click **Deploy**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Edit License Settings

You can enable licenses on your device if you have available licenses on your Firepower Management Center.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to enable or disable licenses, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Device**.

Tip For stacked devices, you enable or disable the licenses for the stack on the Stack page of the appliance editor.

Step 4 In the **License** section, click **Edit** ( ).

Step 5 Check or clear the check box next to the license you want to enable or disable for the managed device.

Step 6 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Edit Advanced Settings

The following topics explain how to edit the advanced device settings.

Note For information about the Transfer Packets setting, see Edit General Settings, on page 240.

Configure Automatic Application Bypass

Automatic Application Bypass (AAB) allows packets to bypass detection if Snort is down or if a packet takes too long to process. AAB causes Snort to restart within ten minutes of the failure, and generates troubleshooting data that can be analyzed to investigate the cause of the Snort failure.

Caution AAB activation partially restarts the Snort process, which temporarily interrupts the inspection of a few packets. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

AAB limits the time allowed to process packets through an interface. You balance packet processing delays with your network's tolerance for packet latency.

The feature functions with any deployment; however, it is most valuable in inline deployments.

Typically, you use Rule Latency Thresholding in the intrusion policy to fast-path packets after the latency threshold value is exceeded. Rule Latency Thresholding does not shut down the engine or generate troubleshooting data.

If detection is bypassed, the device generates a health monitoring alert.

By default the AAB is disabled; to enable AAB follow the steps described.

Before you begin

Model Support—ASA FirePOWER, 7000 & 8000 Series, and NGIPSv

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to edit advanced device settings, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Device** (or **Stack** for stacked devices), then click **Edit** ( ) in the **Advanced** section.

Step 4 Check **Automatic Application Bypass**.

Step 5 Enter a **Bypass Threshold** from 250 ms to 60,000 ms. The default setting is 3000 milliseconds (ms).

Step 6 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Inspect Local Router Traffic

If locally-bound traffic matches a Monitor rule in a Layer 3 deployment, that traffic may bypass inspection. To ensure inspection of the traffic, enable Inspect Local Router Traffic.

Before you begin

Model Support—7000 & 8000 Series

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to edit advanced device settings, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Device** (or **Stack** for stacked devices), then click **Edit** ( ) in the **Advanced Settings** section.

Step 4 Check **Inspect Local Router Traffic** to inspect exception traffic when a 7000 or 8000 Series device is deployed as a router.

Step 5 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configure Fastpath Rules (8000 Series)

As a form of early traffic handling, 8000 Series fastpath rules can send traffic directly through an 8000 Series device without further inspection or logging. (In a passive deployment, 8000 Series fastpath rules simply stop analysis.) Each 8000 Series fastpath rule applies to a specific security zone or inline interface set. Because 8000 Series fastpath rules function at the hardware level, you can use only the following simple, outer-header criteria to fastpath traffic:

- Initiator and responder IP address or address block
- Protocol, and for TCP and UDP, initiator and responder port
- VLAN ID

By default, 8000 Series fastpath rules affect connections from specified initiators to specified responders. To fastpath all connections that meets the rule's criteria, regardless of which host is the initiator and which is the responder, you can make the rule bidirectional.

Note Although they perform a similar function, 8000 Series fastpath rules are not related to the Fastpath tunnel or prefilter rules that you configure in prefilter policies.

Note When you specify a port other than *Any* for TCP or UDP traffic, only the first fragment in matching fragmented traffic is fastpathed. All other fragments are forwarded for further inspection. This is because the 8000 Series only fastpaths fragmented traffic when the IP header in each fragment contains all the IP header information needed to match the fastpath rule, and subsequent fragments do not contain the field that identifies the port.

Before you begin

Model Support—8000 Series

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the 8000 Series device where you want to configure the rule, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Device** (or **Stack** for stacked devices), then click **Edit** ( ) in the Advanced Settings section.

Step 4 Click **New IPv4 Rule** or **New IPv6 Rule**.

Step 5 From the **Domain** drop-down list, choose an inline set or passive security zone.

Step 6 Configure the traffic you want to fastpath. Traffic must meet all the conditions to be fastpathed.

- Initiator and Responder (required): Enter IP addresses or address blocks for initiators and responders.
- Protocol: Choose a protocol, or choose **All**.

- Initiator Port and Responder Port: For TCP and UDP traffic, enter initiator and responder ports. Leave the fields blank or enter **Any** to match all TCP or UDP traffic. You can enter a comma-separated list of ports, but you cannot enter port ranges.
- VLAN: Enter a VLAN ID. Leave the field blank or enter **Any** to match all traffic regardless of VLAN tag.

Step 7 (Optional) Make the rule **Bidirectional**.

Step 8 Click **Save**, then **Save** again.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Change the Manager for the Device

You might need to change the manager on a device in the following circumstances:

- Reestablish the Management Connection if You Change the FMC IP Address, on page 244—If you change the FMC IP address or hostname, reestablishing the management connection depends on how you added the device to the FMC.
- Identify a New FMC, on page 245—After you delete the device from the old FMC, if present, you can configure the device for the new FMC, and then add it to the FMC.
- Switch from Firepower Device Manager to FMC, on page 246—You cannot use both FDM and FMC at the same time for the same device. If you change from FDM to FMC, the FTD configuration will be erased, and you will need to start over.
- Switch from FMC to Firepower Device Manager, on page 247—You cannot use both FDM and FMC at the same time for the same device. If you change from FMC to FDM, the FTD configuration will be erased, and you will need to start over.

Reestablish the Management Connection if You Change the FMC IP Address

When you change the FMC IP address, there is not a command on the device to change the FMC IP address to the new address. Reestablishing the management connection depends on how you added the device to the FMC.

Before you begin

Model Support—FTD

Procedure

Depending on how you added the device to the FMC, see the following tasks:

- **IP address—No action.** If you added the device to the FMC using a reachable device IP address, then the management connection will be reestablished automatically after several minutes even though the

IP address identified on the FTD is the old IP address. **Note**: If you specified a device IP address that is unreachable, then you must contact Cisco TAC, who can advise you how to restore connectivity for your devices.

- **NAT ID only—Contact Cisco TAC.** If you added the device using only the NAT ID, then the connection cannot be reestablished. In this case, you must contact Cisco TAC, who can advise you how to restore connectivity for your devices.

Identify a New FMC

This procedure shows how to identify a new FMC for the managed device. You should perform these steps even if the new FMC uses the old FMC's IP address.

Procedure

Step 1 On the old FMC, if present, delete the managed device.

You cannot change the FMC IP address if you have an active connection with an FMC.

Step 2 Connect to the device CLI, for example using SSH.

Step 3 Configure the new FMC.

configure manager add {*hostname* | *IPv4_address* | *IPv6_address* | **DONTRESOLVE** } *regkey* [*nat_id*]

- {*hostname* | *IPv4_address* | *IPv6_address*}—Sets the FMC hostname, IPv4 address, or IPv6 address.
- **DONTRESOLVE**—If the FMC is not directly addressable, use **DONTRESOLVE** instead of a hostname or IP address. If you use **DONTRESOLVE** , then a *nat_id* is required. When you add this device to the FMC, make sure that you specify both the device IP address and the *nat_id*; one side of the connection needs to specify an IP address, and both sides need to specify the same, unique NAT ID.
- *regkey*—Make up a registration key to be shared between the FMC and the device during registration. You can choose any text string for this key between 1 and 37 characters; you will enter the same key on the FMC when you add the FTD.
- *nat_id*—Make up an alphanumeric string from 1 to 37 characters used only during the registration process between the FMC and the device when one side does not specify an IP address. This NAT ID is a one-time password used only during registration. Make sure the NAT ID is unique, and not used by any other devices awaiting registration. Specify the same NAT ID on the FMC when you add the FTD.

Example:

```
> configure manager add DONTRESOLVE abc123 efg456
Manager successfully configured.
Please make note of reg_key as this will be required while adding Device in FMC.

>
```

Step 4 Add the device to the FMC.

Switch from Firepower Device Manager to FMC

This procedure describes how to change your manager from Firepower Device Manager (FDM), a local device manager, to FMC. You can switch between FDM and FMC without reinstalling the software. You cannot use both FDM and FMC at the same time for the same device. If you change from FDM to FMC, the FTD configuration will be erased, and you will need to start over.



Caution Changing the manager resets the Firepower Threat Defense configuration to the factory default. However, the management bootstrap configuration is maintained.

Before you begin

Model Support—FTD

Procedure

Step 1 In FDM, for High Availability, break the high availability configuration. Ideally, break HA from the active unit.

Step 2 In FDM, unregister the device from the Smart Licensing server.

Step 3 Connect to the device CLI, for example using SSH.

Step 4 Remove the current management setting.

configure manager delete

Caution Deleting the local manager resets the Firepower Threat Defense configuration to the factory default. However, the management bootstrap configuration is maintained.

Example:

```
> configure manager delete

If you enabled any feature licenses, you must disable them in
Firepower Device Manager before deleting the local manager.
Otherwise, those licenses remain assigned to the device in
Cisco Smart Software Manager.
Do you want to continue[yes/no]:yes

DHCP Server Disabled
>
```

Step 5 Configure the new FMC.

configure manager add {*hostname* | *IPv4_address* | *IPv6_address* | **DONTRESOLVE** } *regkey* [*nat_id*]

- {*hostname* | *IPv4_address* | *IPv6_address*}—Sets the FMC hostname, IPv4 address, or IPv6 address.
- **DONTRESOLVE**—If the FMC is not directly addressable, use **DONTRESOLVE** instead of a hostname or IP address. If you use **DONTRESOLVE** , then a *nat_id* is required. When you add this device to the FMC, make sure that you specify both the device IP address and the *nat_id*; one side of the connection needs to specify an IP address, and both sides need to specify the same, unique NAT ID.

- *regkey*—Make up a registration key to be shared between the FMC and the device during registration. You can choose any text string for this key between 1 and 37 characters; you will enter the same key on the FMC when you add the FTD.
- *nat_id*—Make up an alphanumeric string from 1 to 37 characters used only during the registration process between the FMC and the device when one side does not specify an IP address. This NAT ID is a one-time password used only during registration. Make sure the NAT ID is unique, and not used by any other devices awaiting registration. Specify the same NAT ID on the FMC when you add the FTD.

Example:

```
> configure manager add DONTRESOLVE abc123 efg456
Manager successfully configured.
Please make note of reg_key as this will be required while adding Device in FMC.

>
```

Step 6 Add the device to the FMC.

Switch from FMC to Firepower Device Manager

This procedure describes how to change your manager from FMC to Firepower Device Manager (FDM), a local device manager. You can switch between FDM and FMC without reinstalling the software. You cannot use both FDM and FMC at the same time for the same device. If you change from FMC to FDM, the FTD configuration will be erased, and you will need to start over.



Caution Changing the manager resets the Firepower Threat Defense configuration to the factory default. However, the management bootstrap configuration is maintained.

Before you begin

Model Support—FTD

Procedure

Step 1 In FMC, for High Availability, break the high availability configuration. Ideally, break HA from the active unit.

Step 2 In FMC, delete the managed device.

You cannot change the manager if you have an active connection with an FMC.

Step 3 Connect to the device CLI, for example using SSH.

Step 4 Remove the current management setting.

configure manager delete

Caution Deleting the local manager resets the Firepower Threat Defense configuration to the factory default. However, the management bootstrap configuration is maintained.

Example:

```
> configure manager delete

If you enabled any feature licenses, you must disable them in
Firepower Device Manager before deleting the local manager.
Otherwise, those licenses remain assigned to the device in
Cisco Smart Software Manager.
Do you want to continue[yes/no]:yes

DHCP Server Disabled
>
```

Step 5 Configure the new FMC.

configure manager add {*hostname* | *IPv4_address* | *IPv6_address* | **DONTRESOLVE** } *regkey* [*nat_id*]

- {*hostname* | *IPv4_address* | *IPv6_address*}—Sets the FMC hostname, IPv4 address, or IPv6 address.
- **DONTRESOLVE**—If the FMC is not directly addressable, use **DONTRESOLVE** instead of a hostname or IP address. If you use **DONTRESOLVE** , then a *nat_id* is required. When you add this device to the FMC, make sure that you specify both the device IP address and the *nat_id*; one side of the connection needs to specify an IP address, and both sides need to specify the same, unique NAT ID.
- *regkey*—Make up a registration key to be shared between the FMC and the device during registration. You can choose any text string for this key between 1 and 37 characters; you will enter the same key on the FMC when you add the FTD.
- *nat_id*—Make up an alphanumeric string from 1 to 37 characters used only during the registration process between the FMC and the device when one side does not specify an IP address. This NAT ID is a one-time password used only during registration. Make sure the NAT ID is unique, and not used by any other devices awaiting registration. Specify the same NAT ID on the FMC when you add the FTD.

Example:

```
> configure manager add DONTRESOLVE abc123 efg456
Manager successfully configured.
Please make note of reg_key as this will be required while adding Device in FMC.

>
```

Step 6 Add the device to the FMC.

Viewing Device Information

In a multidomain deployment, ancestor domains can view information about all devices in descendant domains. You must be in a leaf domain to edit a device.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **Edit** ( ) next to the device you want to view.

In a multidomain deployment, if you are in an ancestor domain, you can click **View** ( ) to view a device from a descendant domain in read-only mode.

Step 3 Click **Device**.

Step 4 You can view the following information:

- General — Displays general settings for the device; see General Information, on page 249.
- License — Displays license information for the device; see License Information, on page 250.
- System — Displays system information about the device; see System Information, on page 250.
- Health — Displays information about the current health status of the device; see Health Information, on page 250.
- Management — Displays information about the communication channel between the Firepower Management Center and the device; see Management Information, on page 251.
- Advanced — Displays information about advanced feature configuration; see Advanced Settings, on page 251.

Device Management Page Information

The Device Management page provides you with range of information and options to manage Firepower devices:

- View By—Use this option to view the devices based on group, licenses, model, or access control policy.
- Device State—You can also view the devices based on its state. You can click on a state icon to view the devices belonging to it. The number of devices belonging to the states are provided within brackets.
- Search—You can search for a configured device by providing the device name, host name, or the IP address.
- Add options—You can use the add options to configure device, high availability, FTD cluster, stack, and group.
- Edit and other actions—Against each configured device, use the **Edit** ( ) icon to edit the device parameters and attributes.

When you click on the device, the device properties page appears with several tabs. You can use the tabs to view the device information, and configure routing, interfaces, inline sets, and DHCP.

General Information

The General section of the **Device** tab displays the settings described in the table below.

Table 24: General Section Table Fields

Field	Description
Name	The display name of the device on the Firepower Management Center.

Field	Description
Transfer Packets	This displays whether or not the managed device sends packet data with the events to the Firepower Management Center.
Mode	The displays the mode of the management interface for the device: **routed** or **transparent**. **Note** The **Mode** field is displayed only for Firepower Threat Defense devices.
Compliance Mode	This displays the security certifications compliance for a device. Valid values are CC, UCAPL and None.

License Information

The License section of the **Device** page displays the licenses enabled for the device.

System Information

The System section of the **Device** page displays a read-only table of system information, as described in the following table.

Table 25: System Section Table Fields

Field	Description
Model	The model name and number for the managed device.
Serial	The serial number of the chassis of the managed device.
Time	The current system time of the device. This is always in UTC.
Version	The version of the software currently installed on the managed device.
Policy	A link to the platform settings policy currently deployed to the managed device.

You can also shut down or restart the device.

Health Information

The Health section of the **Device** page displays the information described in the table below.

Table 26: Health Section Table Fields

Field	Description
Status	An icon that represents the current health status of the device. Clicking the icon displays the Health Monitor for the appliance.
Policy	A link to a read-only version of the health policy currently deployed at the device.

Field	Description
Blacklist	A link to the Health Blacklist page, where you can enable and disable health blacklist modules.

Management Information

The **Management** section of the **Device** page displays the fields described in the table below.

Table 27: Management Section Table Fields

Field	Description
Host	The IP address or hostname of the device. To change the hostname or IP Address of the device, see Edit Management Settings, on page 234.
Status	An icon indicating the status of the communication channel between the Firepower Management Center and the managed device. You can hover over the status icon to view the last time the Firepower Management Center contacted the device.

Advanced Settings

The **Advanced** section of the **Device** page displays a table of advanced configuration settings, as described below. You can edit any of these settings.

Table 28: Advanced Section Table Fields

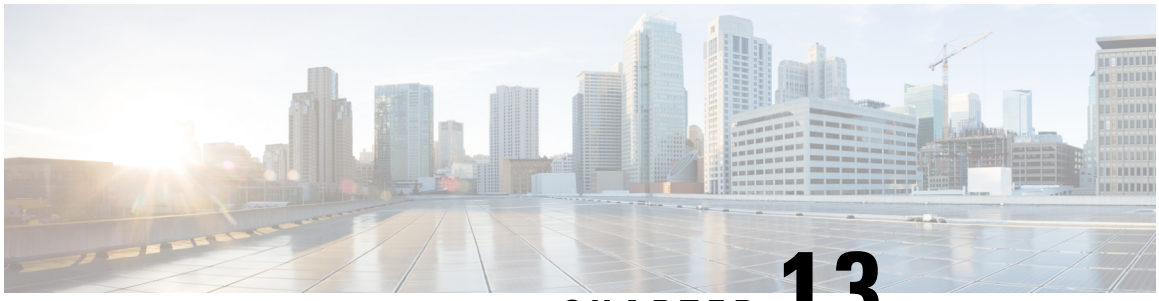
Field	Description	Supported Devices
Application Bypass	The state of Automatic Application Bypass on the device.	7000 & 8000 Series NGIPSv ASA FirePOWER
Bypass Threshold	The Automatic Application Bypass threshold, in milliseconds.	
Inspect Local Router Traffic	Whether the device inspects traffic received on routed interfaces that is destined for itself, such as ICMP, DHCP, and OSPF traffic.	7000 & 8000 Series
Fast-Path Rules	The number of 8000 Series fastpath rules that have been created on the device.	8000 Series



PART III

System Monitoring and Troubleshooting

- Dashboards, on page 255
- Health Monitoring, on page 275
- Monitoring the System, on page 301
- Troubleshooting the System, on page 311



CHAPTER 13

Dashboards

The following topics describe how to use dashboards in the Firepower System:

- About Dashboards, on page 255
- Firepower System Dashboard Widgets, on page 256
- Managing Dashboards, on page 268

About Dashboards

Firepower System dashboards provide you with at-a-glance views of current system status, including data about the events collected and generated by the system. You can also use dashboards to see information about the status and overall health of the appliances in your deployment. Keep in mind that the information the dashboard provides depends on how you license, configure, and deploy the system.

Tip The dashboard is a complex, highly customizable monitoring feature that provides exhaustive data. For a broad, brief, and colorful picture of your monitored network, use the Context Explorer.

Dashboards are available on the Firepower Management Center and 7000 & 8000 Series devices.

A dashboard uses tabs to display widgets: small, self-contained components that provide insight into different aspects of the system. For example, the predefined Appliance Information widget tells you the appliance name, model, and currently running version of the Firepower System software. The system constrains widgets by the dashboard time range, which you can change to reflect a period as short as the last hour or as long as the last year.

The system is delivered with several predefined dashboards, which you can use and modify. If your user role has access to dashboards (Administrator, Maintenance User, Security Analyst, Security Analyst [Read Only], and custom roles with the Dashboards permission), by default your home page is the predefined Summary Dashboard. However, you can configure a different default home page, including non-dashboards. You can also change the default dashboard. Note that if your user role cannot access dashboards, your default home page is relevant to the role; for example, a Discovery Admin sees the Network Discovery page.

You can also use predefined dashboards as the base for custom dashboards, which you can either share or restrict as private. Unless you have Administrator access, you cannot view or modify private dashboards created by other users.

Note Some drill-down pages and table views of events include a **Dashboard** toolbar link that you can click to view a relevant predefined dashboard. If you delete a predefined dashboard or tab, the associated toolbar links do not function.

In a multidomain deployment, you cannot view dashboards from ancestor domains; however, you can create new dashboards that are copies of the higher-level dashboards.

Firepower System Dashboard Widgets

A dashboard has one or more tabs, each of which can display one or more widgets in a three-column layout. The Firepower System is delivered with many predefined dashboard widgets, each of which provides insight into a different aspect of the Firepower System. Widgets are grouped into three categories:

- *Analysis & Reporting widgets* display data about the events collected and generated by the Firepower System.
- *Miscellaneous widgets* display neither event data nor operations data. Currently, the only widget in this category displays an RSS feed.
- *Operations widgets* display information about the status and overall health of the Firepower System.

The dashboard widgets that you can view depend on:

- the type of appliance you are using
- your user role
- your current domain (in a multidomain deployment)

In addition, each dashboard has a set of preferences that determines its behavior.

You can minimize and maximize widgets, add and remove widgets from tabs, as well as rearrange the widgets on a tab.

Note For widgets that display event counts over a time range, the total number of events may not reflect the number of events for which detailed data is available in the tables on pages under the Analysis menu. This occurs because the system sometimes prunes older event details to manage disk space usage. To minimize the occurrence of event detail pruning, you can fine-tune event logging to log only those events most important to your deployment.

Widget Availability

The dashboard widgets that you can view depend on the type of appliance you are using, your user role, and your current domain (in a multidomain deployment).

In a multidomain deployment, if you do not see a widget that you expect to see, switch to the Global domain. See Switching Domains on the Firepower Management Center, on page 10.

Note that:

- An *invalid* widget is one that you cannot view because you are using the wrong type of appliance.
- An *unauthorized* widget is one that you cannot view because your user account does not have the necessary privileges.

For example, the Appliance Status widget is available only on the FMC for users with Administrator, Maintenance User, Security Analyst, or Security Analyst (Read Only) account privileges.

Although you cannot add an unauthorized or invalid widget to a dashboard, an imported dashboard may contain unauthorized or invalid widgets. For example, such widgets can be present if the imported dashboard:

- Was created by a user with different access privileges, or
- Belongs to an ancestor domain.

Unavailable widgets are disabled and display error messages that indicate why you cannot view them.

Individual widgets also display error messages when those widgets have timed out or are otherwise experiencing problems.



Note You can delete or minimize unauthorized and invalid widgets, as well as widgets that display no data, keeping in mind that modifying a widget on a shared dashboard modifies it for all users of the appliance.

Dashboard Widget Availability by User Role

The following table lists the user account privileges required to view each widget. Only user accounts with Administrator, Maintenance User, Security Analyst, or Security Analyst (Read Only) access can use dashboards.

Users with custom roles may have access to any combination of widgets, or none at all, as their user roles permit.

Table 29: User Roles and Dashboard Widget Availability

Widget	Administrator	Maintenance User	Security Analyst	Security
Appliance Information	yes	yes	yes	yes
Appliance Status	yes	yes	yes	no
Correlation Events	yes	no	yes	yes
Current Interface Status	yes	yes	yes	yes
Current Sessions	yes	no	no	no
Custom Analysis	yes	no	yes	yes
Disk Usage	yes	yes	yes	yes
Interface Traffic	yes	yes	yes	yes
Intrusion Events	yes	no	yes	yes

Widget	Administrator	Maintenance User	Security Analyst	Security Ana
Network Compliance	yes	no	yes	yes
Product Licensing	yes	yes	no	no
Product Updates	yes	yes	no	no
RSS Feed	yes	yes	yes	yes
System Load	yes	yes	yes	yes
System Time	yes	yes	yes	yes
White List Events	yes	no	yes	yes

Predefined Dashboard Widgets

The Firepower System is delivered with several predefined widgets that, when used on dashboards, can provide you with at-a-glance views of current system status. These views include:

- data about the events collected and generated by the system
- information about the status and overall health of the appliances in your deployment



Note The dashboard widgets you can view depend on the type of appliance you are using, your user role, and your current domain in a multidomain deployment.

The Appliance Information Widget

The Appliance Information widget provides a snapshot of the appliance. It appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard. The widget provides:

- the name, IPv4 address, IPv6 address, and model of the appliance
- the versions of the Firepower System software, operating system, Snort, rule update, rule pack, module pack, vulnerability database (VDB), and geolocation update installed on the appliances with dashboards, except for virtual Firepower Management Centers
- for managed appliances, the name and status of the communications link with the managing appliance

You can configure the widget to display more or less information by modifying the widget preferences to display a simple or an advanced view; the preferences also control how often the widget updates.

The Appliance Status Widget

The Appliance Status widget indicates the health of the appliance and of any appliances it is managing. Note that because the Firepower Management Center does not automatically apply a health policy to managed devices, you must manually apply a health policy to devices or their status appears as `Disabled`. This widget appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard.

You can configure the widget to display appliance status as a pie chart or in a table by modifying the widget preferences.

The preferences also control how often the widget updates.

You can click a section on the pie chart or one of the numbers on the appliance status table to go to the Health Monitor page and view the compiled health status of the appliance and of any appliances it is managing.

The Correlation Events Widget

The Correlation Events widget shows the average number of correlation events per second, by priority, over the dashboard time range. It appears by default on the Correlation tab of the Detailed Dashboard.

You can configure the widget to display correlation events of different priorities by modifying the widget preferences, as well as to choose a linear (incremental) or logarithmic (factor of ten) scale.

Check one or more **Priorities** check boxes to display separate graphs for events of specific priorities, including events that do not have a priority. Choose **Show All** to display an additional graph for all correlation events, regardless of priority. The preferences also control how often the widget updates.

You can click a graph to view correlation events of a specific priority, or click the **All** graph to view all correlation events. In either case, the events are constrained by the dashboard time range; accessing correlation events via the dashboard changes the events (or global) time window for the appliance.

The Current Interface Status Widget

The Current Interface Status widget shows the status of all interfaces on the appliance, enabled or unused. On a Firepower Management Center, you can display the management (`eth0`, `eth1`, and so on) interfaces. On a managed device, you can choose to show only sensing (`s1p1` and so on) interfaces or both management and sensing interfaces. Interfaces are grouped by type: management, inline, passive, switched, routed, stacked, and unused.

For each interface, the widget provides:

- the name of the interface
- the link state of the interface
- the link mode (for example, 100Mb full duplex, or 10Mb half duplex) of the interface
- the type of interface, that is, copper or fiber
- the amount of data received (Rx) and transmitted (Tx) by the interface

The color of the ball representing link state indicates the current status, as follows:

- green: link is up and at full speed
- yellow: link is up but not at full speed
- red: link is not up
- gray: link is administratively disabled
- blue: link state information is not available (for example, ASA)

The widget preferences control how often the widget updates.

The Current Sessions Widget

The Current Sessions widget shows which users are currently logged into the appliance, the IP address associated with the machine where the session originated, and the last time each user accessed a page on the appliance (based on the local time for the appliance). The user that represents you, that is, the user currently viewing the widget, is marked with a **User icon** and rendered in bold type. Sessions are pruned from this widget's data within one hour of logoff or inactivity. This widget appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard.

On the Current Sessions widget, you can:

- click any user name to manage user accounts on the User Management page.
- click the **Host icon** or **Compromised Host icon** next to any IP address to view the host profile for the associated machine.
- click any IP address or access time to view the audit log constrained by that IP address and by the time that the user associated with that IP address logged on to the web interface.

The widget preferences control how often the widget updates.

The Custom Analysis Widget

The Custom Analysis widget is a highly customizable widget that allows you to display detailed information on the events collected and generated by the Firepower System.

The widget is delivered with multiple presets that provide quick access to information about your deployment. The predefined dashboards make extensive use of these presets. You can use these presets or create a custom configuration. At a minimum, a custom configuration specifies the data you are interested in (table and field), and an aggregation method for that data. You can also set other display-related preferences, including whether you want to show events as relative occurences (bar graph) or over time (line graph).

The widget displays the last time it updated, based on local time. The widget updates with a frequency that depends on the dashboard time range. For example, if you set the dashboard time range to an hour, the widget updates every five minutes. On the other hand, if you set the dashboard time range to a year, the widget updates once a week. To determine when the dashboard will update next, hover your pointer over the **Last updated** notice in the bottom left corner of the widget.



Note A red-shaded Custom Analysis widget indicates that its use is harming system performance. If the widget continues to stay red over time, remove the widget. You can also disable all Custom Analysis widgets from the Dashboard settings in your system configuration (**System** > **Configuration** > **Dashboard**)

Displaying Relative Occurrences of Events (Bar Graphs)

For bar graphs in the Custom Analysis widget, the colored bars in the widget background show the relative number of occurrences of each event. Read the bars from right to left.

The **Direction icon** indicates and controls the sort order of the display. A downward-pointing icon indicates descending order; an upward-pointing icon indicates ascending order. To change the sort order, click the icon.

Next to each event, the widget can display one of three icons to indicate any changes from the most recent results:

- The new event icon **Add** () signifies that the event is new to the results.

- The **Up Arrow icon** indicates that the event has moved up in the standings since the last time the widget updated. A number indicating how many places the event has moved up appears next to the icon.
- The **Down Arrow icon** indicates that the event has moved down in the standings since the last time the widget updated. A number indicating how many places the event has moved down appears next to the icon.

Displaying Events Over Time (Line Graphs)

If you want information on events or other collected data over time, you can configure the Custom Analysis widget to display a line graph, such as one that displays the total number of intrusion events generated in your deployment over time.

Limitations to the Custom Analysis Widget

A Custom Analysis widget may indicate that you are unauthorized to view the data that is configured to display. For example, Maintenance Users are not authorized to view discovery events. As another example, the widget does not display information related to unlicensed features. However, you (and any other users who share the dashboard) can modify the widget preferences to display data that you can see, or even delete the widget. If you want to make sure that this does not happen, save the dashboard as private.

When viewing user data, the system displays only authoritative users.

When viewing URL category information, the system does not display uncategorized URLs.

When viewing intrusion events aggregated by **Count**, the count includes reviewed events for intrusion events; if you view the count in tables on pages under the Analysis menus, the count will not include reviewed events.



Note In a multidomain deployment, the system builds a separate network map for each leaf domain. As a result, a leaf domain can contain an IP address that is unique within its network, but identical to an IP address in another leaf domain. When you view Custom Analysis widgets in an ancestor domain, multiple instances of that repeated IP address can be displayed. At first glance, they might appear to be duplicate entries. However, if you drill down to the host profile information for each IP address, the system shows that they belong to different leaf domains.

How to Create Dashboard Widgets for a Device

Any widgets that show events from devices can be configured to use a filter that limits the display of events for a given device or a set of devices.

1. Create and save a search: Go to **Analysis** > **Search** and enter the search parameters to match the specific device names.



Note You must provide exact text match as there is no drop-down listing the deployed device names.

2. Go to **Overview** > **Dashboards** > **Add Widgets** to create a **Custom Analysis** widget.
3. Return to **Overview** > **Dashboards** and modify the new widget to customize with the scope of search.

Example: Configuration of Custom Analysis Widget

You can configure the Custom Analysis widget to display a list of recent intrusion events by configuring the widget to display data from the **Intrusion Events** table. Choosing the **Classification** field and aggregating this data by **Count** displays the number of events that were generated for each type.

On the other hand, aggregating by **Unique Events** displays the number of unique intrusion events of each type (for example, how many detections of network trojans, potential violations of corporate policy, attempted denial-of-service attacks, and so on).

You can further customize the widget using a saved search, either one of the predefined searches delivered with your appliance or a custom search that you created. For example, constraining the first example (intrusion events using the **Classification** field, aggregated by **Count**) using the **Dropped Events** search displays the number of intrusion events that were dropped for each type.

Related Topics

Modifying Dashboard Time Settings, on page 272

Custom Analysis Widget Preferences

The following table describes the preferences you can set in the Custom Analysis widget.

Different preferences appear depending on how you configure the widget. For example, a different set of preferences appears if you configure the widget to show relative occurrences of events (a bar graph) vs a graph over time (a line graph). Some preferences, such as Filter, only appear if you choose a specific table from which to display data.

Table 30: Custom Analysis Widget Preferences

Preference	Details
Title	If you do not specify a title for the widget, the system uses the configured event type as the title.
Preset	Custom Analysis presets provide quick access to information about your deployment. The predefined dashboards make extensive use of these presets. You can use these presets or you can create a custom configuration.
Table (required)	The table of events or assets that contains the data the widget displays.
Field (required)	The specific field of the event type you want to display. To show data over time (line graphs), choose **Time**. To show relative occurrences of events (bar graphs), choose another option.
Aggregate (required)	The aggregation method configures how the widget groups the data it displays. For most event types, the default option is **Count**.
Filter	You can use application filters to constrain data from the Application Statistics and Intrusion Event Statistics by Application tables.

Preference	Details
Search	You can use a saved search to constrain the data that the widget displays. You do not have to specify a search, although some presets use predefined searches. Only you can access searches that you have saved as private. If you configure the widget on a shared dashboard and constrain its events using a private search, the widget resets to not using the search when another user logs in. This affects your view of the widget as well. If you want to make sure that this does not happen, save the dashboard as private. Only fields that constrain connection summaries can constrain Custom Analysis dashboard widgets based on connection events. Invalid saved searches are dimmed. If you constrain a Custom Analysis widget using a saved search, then edit the search, the widget does not reflect your changes until the next time it updates.
Show	Choose whether you want to display the most (**Top**) or the least (**Bottom**) frequently occurring events.
Results	Choose the number of result rows to display.
Show Movers	Choose whether you want to display the icons that indicate changes from the most recent results.
Time Zone	Choose the time zone you want to use to display results.
Color	You can change the color of the bars in the widget's bar graph.

Related Topics

Configuring Widget Preferences, on page 270

Viewing Associated Events from the Custom Analysis Widget

From a Custom Analysis widget, you can invoke an event view (workflow) that provides detailed information about the events displayed in the widget. The events appear in the default workflow for that event type, constrained by the dashboard time range. This also changes the appropriate time window on the Firepower Management Center, depending on how many time windows you configured and on the event type.

For example:

- If you configure multiple time windows, then access health events from a Custom Analysis widget, the events appear in the default health events workflow, and the health monitoring time window changes to the dashboard time range.
- If you configure a single time window and then access any type of event from the Custom Analysis widget, the events appear in the default workflow for that event type, and the global time window changes to the dashboard time range.

Procedure

You have the following choices:

- On any Custom Analysis widget, click **View** () in the lower right corner of the widget to view all associated events, constrained by the widget preferences.

- On a Custom Analysis widget showing relative occurrences of events (bar graph), click any event to view associated events constrained by the widget preferences, as well as by that event.

The Disk Usage Widget

The Disk Usage widget displays the percentage of space used on the hard drive, based on disk usage category. It also indicates the percentage of space used on and capacity of each partition of the appliance's hard drive. The Disk Usage widget displays the same information for the malware storage pack if installed in the device, or if the Firepower Management Center manages a device containing a malware storage pack. This widget appears by default on the Status tabs of the Default Dashboard and the Summary Dashboard.

The By Category stacked bar displays each disk usage category as a proportion of the total available disk space used. The following table describes the available categories.

Table 31: Disk Usage Categories

Disk Usage Category	Description
Events	all events logged by the system
Files	all files stored by the system
Backups	all backup files
Updates	all files related to updates, such as rule updates and system updates
Other	system troubleshooting files and other miscellaneous files
Free	free space remaining on the appliance

You can hover your pointer over a disk usage category in the By Category stacked bar to view the percentage of available disk space used by that category, the actual storage space on the disk, and the total disk space available for that category. Note that if you have a malware storage pack installed, the total disk space available for the Files category is the available disk space on the malware storage pack.

You can configure the widget to display only the By Category stacked bar, or you can show the stacked bar plus the admin (`/`), `/Volume`, and `/boot` partition usage, as well as the `/var/storage` partition if the malware storage pack is installed, by modifying the widget preferences.

The widget preferences also control how often the widget updates, as well as whether it displays the current disk usage or collected disk usage statistics over the dashboard time range.

The Interface Traffic Widget

The Interface Traffic widget shows the rate of traffic received (Rx) and transmitted (Tx) on the appliance's management interface. For 7000 & 8000 Series devices, the widget also shows information on the sensing interfaces. The widget does not appear by default on any of the predefined dashboards.

Devices with Malware licenses enabled periodically attempt to connect to the AMP cloud even if you have not configured dynamic analysis. Because of this, these devices show transmitted traffic; this is expected behavior. Outbound (transmitted) traffic includes flow control packets. Because of this, passive sensing interfaces on 7000 & 8000 Series devices may show transmitted traffic; this is also expected behavior.

The widget preferences control how often the widget updates. On 7000 & 8000 Series devices, the preferences also control whether the widget displays the traffic rate for unused interfaces (by default, the widget only displays the traffic rate for active interfaces).

The Intrusion Events Widget

The Intrusion Events widget shows the intrusion events that occurred over the dashboard time range, organized by priority. This includes statistics on intrusion events with dropped packets and different impacts. This widget appears by default on the Intrusion Events tab of the Summary Dashboard.

In the widget preferences, you can choose:

- **Event Flags** to display separate graphs for events with dropped packets, would have dropped packets, or specific impacts. Choose **All** to display an additional graph for all intrusion events, regardless of impact or rule state.

  For explanations of the icons, see Working with Intrusion Events, on page 2251. The arrow (if any) that appears above the impact level numbers describes the inline result and is defined as follows:

  Table 32: Inline Result Field Contents in Workflow and Table Views

This Icon	Indicates
A black down arrow	The system dropped the packet that triggered the rule.
A gray down arrow	IPS would have dropped the packet if you enabled the **Drop when Inline** intrusion policy option (in an inline deployment), or if a Drop and Generate rule generated the event while the system was pruning.
No icon (blank)	The triggered rule was not set to Drop and Generate Events

  In a passive deployment, the system does not drop packets, including when an inline interface is in tap mode, regardless of the rule state or the inline drop behavior of the intrusion policy.

- **Show** to specify **Average Events Per Second** (EPS) or **Total Events**.
- **Vertical Scale** to specify **Linear** (incremental) or **Logarithmic** (factor of ten) scale.
- How often the widget updates.

On the widget, you can:

- Click a graph corresponding to dropped packets, to would have dropped packets, or to a specific impact to view intrusion events of that type.
- Click the graph corresponding to dropped events to view dropped events.
- Click the graph corresponding to would have dropped events to view would have dropped events.
- Click the **All** graph to view all intrusion events.

The resulting event view is constrained by the dashboard time range; accessing intrusion events via the dashboard changes the events (or global) time window for the appliance. Note that packets in a passive deployment are not dropped, regardless of intrusion rule state or the inline drop behavior of the intrusion policy.

The Network Compliance Widget

The Network Compliance widget summarizes your hosts' compliance with the white lists you configured. By default, the widget displays a pie chart that shows the number of hosts that are compliant, non-compliant, and that have not been evaluated, for all compliance white lists in active correlation policies. This widget appears by default on the Correlation tab of the Detailed Dashboard.

You can configure the widget to display network compliance either for all white lists or for a specific white list by modifying the widget preferences.

If you choose to display network compliance for all white lists, the widget considers a host to be non-compliant if it is not compliant with any white list in an active correlation policy.

You can also use the widget preferences to specify which of three different styles you want to use to display network compliance.

The **Network Compliance** style (the default) displays a pie chart that shows the number of hosts that are compliant, non-compliant, and that have not been evaluated. You can click the pie chart to view the host violation count, which lists the hosts that violate at least one white list.

The **Network Compliance over Time (%)** style displays a stacked area graph showing the relative proportion of hosts that are compliant, non-compliant, and that have not yet been evaluated, over the dashboard time range.

The **Network Compliance over Time** style displays a line graph that shows the number of hosts that are compliant, non-compliant, and that have not yet been evaluated, over the dashboard time range.

The preferences control how often the widget updates. You can check the **Show Not Evaluated** box to hide events which have not been evaluated.

The Product Licensing Widget

The Product Licensing widget shows the device and feature licenses currently installed on the Firepower Management Center. It also indicates the number of items licensed and the number of remaining licensed items allowed. It does not appear by default on any of the predefined dashboards.

The top section of the widget displays all device and feature licenses installed on the Firepower Management Center, including temporary licenses, while the Expiring Licenses section displays only temporary and expired licenses.

The bars in the widget background show the percentage of each type of license that is being used; you should read the bars from right to left. Expired licenses are marked with a strikethrough.

You can configure the widget to display either the features that are currently licensed, or all the features that you can license, by modifying the widget preferences. The preferences also control how often the widget updates.

You can click any of the license types to go to the License page of the local configuration and add or delete feature licenses.

The Product Updates Widget

The Product Updates widget provides you with a summary of the software currently installed on the appliance as well as information on updates that you have downloaded, but not yet installed. This widget appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard.

Because the widget uses scheduled tasks to determine the latest version, it displays `Unknown` until you configure a scheduled task to download, push or install updates.

You can configure the widget to hide the latest versions by modifying the widget preferences. The preferences also control how often the widget updates.

The widget also provides you with links to pages where you can update the software. You can:

- Manually update an appliance by clicking the current version.
- Create a scheduled task to download an update by clicking the latest version.

The RSS Feed Widget

The RSS Feed widget adds an RSS feed to a dashboard. By default, the widget shows a feed of Cisco security news. It appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard.

You can also configure the widget to display a preconfigured feed of company news, the Snort.org blog, or the Cisco Threat Research blog, or you can create a custom connection to any other RSS feed by specifying its URL in the widget preferences.

Feeds update every 24 hours (although you can manually update the feed), and the widget displays the last time the feed was updated based on the local time of the appliance. Keep in mind that the appliance must have access to the web site (for the two preconfigured feeds) or to any custom feed you configure.

When you configure the widget, you can also choose how many stories from the feed you want to show in the widget, as well as whether you want to show descriptions of the stories along with the headlines; keep in mind that not all RSS feeds use descriptions.

On the RSS Feed widget, you can:

- click one of the stories in the feed to view the story
- click the **more** link to go to the feed's web site
- click **Update** () to manually update the feed

The System Load Widget

The System Load widget shows the CPU usage (for each CPU), memory (RAM) usage, and system load (also called the load average, measured by the number of processes waiting to execute) on the appliance, both currently and over the dashboard time range. It appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard.

You can configure the widget to show or hide the load average by modifying the widget preferences. The preferences also control how often the widget updates.

The System Time Widget

The System Time widget shows the local system time, uptime, and boot time for the appliance. It appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard.

You can configure the widget to hide the boot time by modifying the widget preferences. The preferences also control how often the widget synchronizes with the appliance's clock.

The White List Events Widget

The White List Events widget shows the average events per second by priority, over the dashboard time range. It appears by default on the Correlation tab of the Default Dashboard.

You can configure the widget to display white list events of different priorities by modifying the widget preferences.

In the widget preferences, you can:

- choose one or more **Priorities** check boxes to display separate graphs for events of specific priorities, including events that do not have a priority
- choose **Show All** to display an additional graph for all white list events, regardless of priority
- choose **Vertical Scale** to choose **Linear** (incremental) or **Logarithmic** (factor of ten) scale

The preferences also control how often the widget updates.

You can click a graph to view white list events of a specific priority, or click the **All** graph to view all white list events. In either case, the events are constrained by the dashboard time range; accessing white list events via the dashboard changes the events (or global) time window for the Firepower Management Center.

Managing Dashboards

Procedure

Step 1 Choose **Overview** > **Dashboards**, and then choose the dashboard you want to modify from the menu.

Step 2 Manage your dashboards:

- Create Dashboards — Create a custom dashboard; see Creating Custom Dashboards, on page 270.
- Delete Dashboards — To delete a dashboard, click **Delete** () next to the dashboard you want to delete. If you delete your default dashboard, you must define a new default or the appliance prompts you to choose a dashboard every time you attempt to view a dashboard.
- Edit Options — Edit custom dashboard options; see Editing Dashboards Options, on page 272.
- Modify Time Constraints — Modify the time display or pause/unpause the dashboard as described in Modifying Dashboard Time Settings, on page 272.

Step 3 Add (see Adding a Dashboard, on page 269), Delete (click **Close** ()), and Rename (see Renaming a Dashboard, on page 273) dashboards.

Note You cannot change the order of dashboards.

Step 4 Manage dashboard widgets:

- Add Widgets — Add widgets to a dashboard; see Adding Widgets to a Dashboard, on page 269.
- Configure Preferences — Configure widget preferences; see Configuring Widget Preferences, on page 270.
- Customize Display — Customize the widget display; see Customizing the Widget Display, on page 271.
- View Events — View associated events from the Custom Analysis Widget; see Viewing Associated Events from the Custom Analysis Widget, on page 263.

Tip Every configuration of the Custom Analysis widget in the Cisco predefined dashboards corresponds to a system preset for that widget. If you change or delete one of these widgets, you can restore it by creating a new Custom Analysis widget based on the appropriate preset.

Adding a Dashboard

Procedure

Step 1 View the dashboard you want to modify; see Viewing Dashboards, on page 274.

Step 2 Click **Add** (+).

Step 3 Enter a name.

Step 4 Click **OK**.

Adding Widgets to a Dashboard

Each tab can display one or more widgets in a three-column layout. When adding a widget to a dashboard, you choose the tab to which you want to add the widget. The system automatically adds it to the column with the fewest widgets. If all columns have an equal number of widgets, the new widget is added to the leftmost column. You can add a maximum of 15 widgets to a dashboard tab.

Tip After you add widgets, you can move them to any location on the tab. You cannot, however, move widgets from tab to tab.

The dashboard widgets you can view depend on the type of appliance you are using, your user role, and your current domain (in a multidomain deployment). Keep in mind that because not all user roles have access to all dashboard widgets, users with fewer permissions viewing a dashboard created by a user with more permissions may not be able to use all of the widgets on the dashboard. Although the unauthorized widgets still appear on the dashboard, they are disabled.

Procedure

Step 1 View the dashboard where you want to add a widget; see Viewing Dashboards, on page 274.

Step 2 Click the tab where you want to add the widget.

Step 3 Click **Add Widgets**. You can view the widgets in each category by clicking on the category name, or you can view all widgets by clicking **All Categories**.

Step 4 Click **Add** next to the widgets you want to add. The Add Widgets page indicates how many widgets of each type are on the tab, including the widget you want to add.

Tip To add multiple widgets of the same type (for example, you may want to add multiple RSS Feed widgets, or multiple Custom Analysis widgets), click **Add** again.

Step 5 When you are finished adding widgets, click **Done** to return to the dashboard.

What to do next

- If you added a Custom Analysis widget, configure the widget preferences; see Configuring Widget Preferences, on page 270.

Related Topics

Widget Availability, on page 256

Configuring Widget Preferences

Each widget has a set of preferences that determines its behavior.

Procedure

Step 1 On the title bar of the widget whose preferences you want to change, click **Show Preferences** ( ).

Step 2 Make changes as needed.

Step 3 On the widget title bar, click **Hide Preferences** ( ) to hide the preferences section.

Creating Custom Dashboards

Tip Instead of creating a new dashboard, you can export a dashboard from another appliance, then import it onto your appliance. You can then edit the imported dashboard to suit your needs.

Procedure

Step 1 Choose **Overview** > **Dashboards** > **Management**.

Step 2 Click **Create Dashboard**.

Step 3 Modify the custom dashboard options as described in Custom Dashboard Options, on page 270.

Step 4 Click **Save**.

Custom Dashboard Options

The table below describes options you can use when creating or editing custom dashboards.

Table 33: Custom Dashboard Options

Option	Description
Copy Dashboard	When you create a custom dashboard, you can choose to base it on any existing dashboard, whether user-created or system-defined. This option makes a copy of the preexisting dashboard, which you can modify to suit your needs. Optionally, you can create a blank new dashboard by choosing **None**. This option is available only when you create a new dashboard. In a multidomain deployment, you can copy any non-private dashboards from ancestor domains.
Name	A unique name for the custom dashboard.
Description	A brief description of the custom dashboard.
Change Tabs Every	Specifies (in minutes) how often the dashboard should cycle through its tabs. Unless you pause the dashboard or your dashboard has only one tab, this setting advances your view to the next tab at the interval you specify. To disable tab cycling, enter `0` in the **Change Tabs Every** field.
Refresh Page Every	Determines how often the entire dashboard page automatically refreshes. Refreshing the entire dashboard allows you to see any preference or layout changes that were made to a shared dashboard by another user, or that you made to a private dashboard on another computer, since the last time the dashboard refreshed. A frequent refresh can be useful, for example, in a networks operations center (NOC) where a dashboard is displayed at all times. If you make changes to the dashboard at a local computer, the dashboard in the NOC automatically refreshes at the interval you specify, and no manual refresh is required. This refresh does not update the data, and you do not need to refresh the entire dashboard to see data updates; individual widgets update according to their preferences. This value must be greater than the **Change Tabs Every** setting. Unless you pause the dashboard, this setting will refresh the entire dashboard at the interval you specify. To disable the periodic page refresh, enter `0` in the **Refresh Page Every** field. **Note** This setting is separate from the update interval available on many individual widgets; although refreshing the dashboard page resets the update interval on individual widgets, widgets will update according to their individual preferences even if you disable the **Refresh Page Every** setting.
Save As Private	Determines whether the custom dashboard can be viewed and modified by all users of the appliance or is associated with your user account and reserved solely for your own use. Keep in mind that any user with dashboard access, regardless of role, can modify shared dashboards. If you want to make sure that only you can modify a particular dashboard, save it as private.

Customizing the Widget Display

You can minimize and maximize widgets, as well as rearrange the widgets on a tab.

Procedure

Step 1 View a dashboard; see Viewing Dashboards, on page 274.

Step 2 Customize the widget display:

- To rearrange a widget on a tab, click the title bar of the widget you want to move, then drag it to its new location.

  Note You cannot move widgets from tab to tab. If you want a widget to appear on a different tab, you must delete it from the existing tab and add it to the new tab.

- To minimize or maximize a widget on the dashboard, click **Minimize** ( ) or **Maximize** ( ) in a widget's title bar.
- To delete a widget if you no longer want to view it on a tab, click **Close** ( ) in the title bar of the widget.

Editing Dashboards Options

Procedure

Step 1 View the dashboard you want to edit; see Viewing Dashboards, on page 274.

Step 2 Click **Edit** ( ).

Step 3 Change the options as described in Custom Dashboard Options, on page 270.

Step 4 Click **Save**.

Modifying Dashboard Time Settings

You can change the time range to reflect a period as short as the last hour (the default) or as long as the last year. When you change the time range, the widgets that can be constrained by time automatically update to reflect the new time range.

The maximum number of data points in any graph is 300, and the time setting determines how much time is summarized within each data point. Following is the number of data points, and the time span covered, in the dashboards for each time range:

- 1 hour = 12 data points, 5 minutes each
- 6 hours = 72 data points, 5 minutes each
- 1 day = 288 data points, 5 minutes each
- 1 week = 300 data points, 33.6 minutes each
- 2 weeks = 300 data points, 67.2 minutes each
- 30 days = 300 data points, 144 minutes each
- 90 days = 300 data points, 432 minutes each
- 180 days = 300 data points, 864 minutes each
- 1 year = 300 data points, 1752 minutes each

Note that not all widgets can be constrained by time. For example, the dashboard time range has no effect on the Appliance Information widget, which provides information that includes the appliance name, model, and current version of the Firepower System software.

Keep in mind that for enterprise deployments of the Firepower System, changing the time range to a long period may not be useful for widgets like the Custom Analysis widget, depending on how often newer events replace older events.

You can also pause a dashboard, which allows you to examine the data provided by the widgets without the display changing and interrupting your analysis. Pausing a dashboard has the following effects:

- Individual widgets stop updating, regardless of any **Update Every** widget preference.
- Dashboard tabs stop cycling, regardless of the **Cycle Tabs Every** setting in the dashboard properties.
- Dashboard pages stop refreshing, regardless of the **Refresh Page Every** setting in the dashboard properties.
- Changing the time range has no effect.

When you are finished with your analysis, you can unpause the dashboard. Unpausing the dashboard causes all appropriate widgets on the page to update to reflect the current time range. In addition, dashboard tabs resume cycling and the dashboard page resumes refreshing according to the settings you specified in the dashboard properties.

If you experience connectivity problems or other issues that interrupt the flow of system information to the dashboard, the dashboard automatically pauses and an error notice appears until the problem is resolved.



Note Your session normally logs you out after 1 hour of inactivity (or another configured interval), regardless of whether the dashboard is paused. If you plan to passively monitor the dashboard for long periods of time, consider exempting some users from session timeout, or changing the system timeout settings.

Procedure

Step 1 View the dashboard where you want to add a widget; see Viewing Dashboards, on page 274.

Step 2 Optionally, to change the dashboard time range, choose a time range from the **Show the Last** drop-down list.

Step 3 Optionally, pause or unpause the dashboard on the time range control, using **Pause** () or **Play** ().

Renaming a Dashboard

Procedure

Step 1 View the dashboard you want to modify; see Viewing Dashboards, on page 274.

Step 2 Click the dasboard title you want to rename.

Step 3 Type a name.

Step 4 Click **OK**.

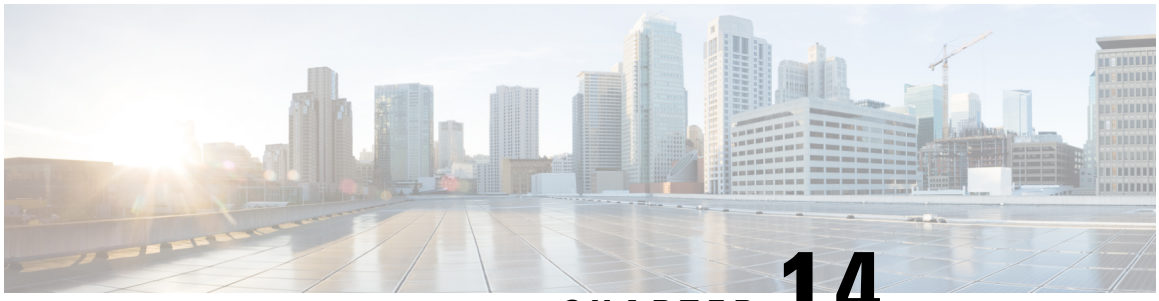
Viewing Dashboards

By default, the home page for your appliance displays the default dashboard. If you do not have a default dashboard defined, the home page shows the Dashboard Management page, where you can choose a dashboard to view.

Procedure

At any time, you can do one of the following:

- To view the default dashboard for your appliance, choose **Overview** > **Dashboards**.
- To view a specific dashboard, choose **Overview** > **Dashboards**, and choose the dashboard from the menu.
- To view all available dashboards, choose **Overview** > **Dashboards** > **Management**. You can then choose **View** ( ) next to an individual dashboard to view it.



CHAPTER 14

Health Monitoring

The following topics describe how to use health monitoring in the Firepower System:

- Requirements and Prerequisites for Health Monitoring, on page 275
- About Health Monitoring, on page 275
- Health Policies, on page 283
- The Health Monitor Blocklist, on page 286
- Health Monitor Alerts, on page 288
- Using the Health Monitor, on page 291
- Viewing Appliance Health Monitors, on page 292
- Health Event Views, on page 295

Requirements and Prerequisites for Health Monitoring

Model Support

Any

Supported Domains

Any

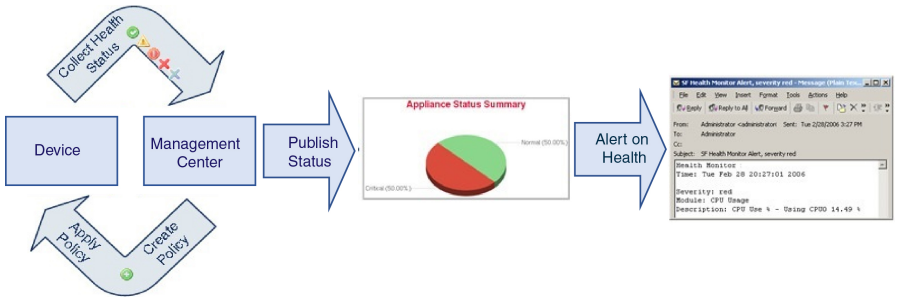
User Roles

Admin

Maintenance User

About Health Monitoring

The health monitor on the Firepower Management Center tracks a variety of health indicators to ensure that the hardware and software in the Firepower System are working correctly. You can use the health monitor to check the status of critical functionality across your Firepower System deployment.



You can use the health monitor to create a collection of tests, referred to as a *health policy*, and apply the health policy to one or more appliances. The tests, referred to as *health modules*, are scripts that test for criteria you specify. You can modify a health policy by enabling or disabling tests or by changing test settings, and you can delete health policies that you no longer need. You can also suppress messages from selected appliances by blocking them.

The tests in a health policy run automatically at the interval you configure. You can also run all tests, or a specific test, on demand. The health monitor collects health events based on the test conditions configured.

Note All appliances automatically report their hardware status via the Hardware Alarms health module. The Firepower Management Center also automatically reports status using the modules configured in the default health policy. Some health modules, such as the Appliance Heartbeat module, run on the Firepower Management Center and report the status of the Firepower Management Center's managed devices. Some health modules do not provide managed device status unless you apply a health policy configured with those modules to a device.

You can use the health monitor to access health status information for the entire system, for a particular appliance, or, in a multidomain deployment, a particular domain. Pie charts and status tables on the Health Monitor page provide a visual summary of the status of all appliances on your network, including the Firepower Management Center. Individual appliance health monitors let you drill down into health details for a specific appliance.

Fully customizable event views allow you to quickly and easily analyze the health status events gathered by the health monitor. These event views allow you to search and view event data and to access other information that may be related to the events you are investigating. For example, if you want to see all the occurrences of CPU usage with a certain percentage, you can search for the CPU usage module and enter the percentage value.

You can also configure email, SNMP, or syslog alerting in response to health events. A *health alert* is an association between a standard alert and a health status level. For example, if you need to make sure an appliance never fails due to hardware overload, you can set up an email alert. You can then create a health alert that triggers that email alert whenever CPU, disk, or memory usage reaches the Warning level you configure in the health policy applied to that appliance. You can set alerting thresholds to minimize the number of repeating alerts you receive.

Note The health monitoring can take 5-6 minutes from the occurrence of the health event to generate a health alert.

You can also generate troubleshooting files for an appliance if you are asked to do so by Support.

Because health monitoring is an administrative activity, only users with administrator user role privileges can access system health data.

High Availability Pair

In a FMC high-availability deployment running Version 6.7 or higher, the active FMC creates a health monitor page that uses REST APIs to show detailed metric-based information. The standby FMC creates the health monitor page that shows the alert information and provide a visual summary of the status of all appliances on your network using pie charts and status tables. The standby FMC does not display the metric-based information.

Health Modules

Health modules, or *health tests,* test for the criteria you specify in a health policy.

Table 34: Health Modules

Module	Platforms	Description
AMP for Endpoints Status	FMC	The module alerts if the FMC cannot connect to the AMP cloud or Cisco AMP Private Cloud after an initial successful connection, or if the private cloud cannot contact the public AMP cloud. It also alerts if you deregister an AMP cloud connection using the AMP for Endpoints management console.
AMP for Firepower Status (AMP for Networks Status)	FMC	This module alerts if: • The FMC cannot contact the AMP cloud (public or private) or the Cisco Threat Grid public cloud or on-premises appliance, or the AMP private cloud cannot contact the public AMP cloud. • The encryption keys used for the connection are invalid. • A device cannot contact the Cisco Threat Grid cloud or an Cisco Threat Grid on-premises appliance to submit files for dynamic analysis. • An excessive number of files are detected in network traffic based on the file policy configuration. If your FMC loses connectivity to the Internet, the system may take up to 30 minutes to generate a health alert.
Appliance Heartbeat	Any	This module determines if an appliance heartbeat is being heard from the appliance and alerts based on the appliance heartbeat status.
Automatic Application Bypass Status	Firepower 7000/8000 series	This module determines if an appliance has been bypassed because it did not respond within the number of seconds set in the bypass threshold, and alerts when a bypass occurs.
Backlog Status	FMC	This module alerts if the backlog of event data awaiting transmission from the device to the FMC has grown continuously for more than 30 minutes. To reduce the backlog, evaluate your bandwidth and consider logging fewer events.
Card Reset	Any	This module checks for network cards which have restarted due to hardware failure and alerts when a reset occurs.

Module	Platforms	Description
Classic License Monitor	FMC	This module determines if sufficient Classic licenses remain. It also alerts when devices in a stack have mismatched license sets. It alerts based on a warning level automatically configured for the module. You cannot change the configuration of this module.
Cluster/Failover Status	FTD	This module monitors the status of device clusters. The module alerts if: • A new primary unit is elected to a cluster. • A new secondary unit joins a cluster. • A primary or secondary unit leaves a cluster.
CPU Usage	Any	This module checks that the CPU on the appliance is not overloaded and alerts when CPU usage exceeds the percentages configured for the module.
Disk Status	Any	This module examines performance of the hard disk, and malware storage pack (if installed) on the appliance. This module generates a Warning (yellow) health alert when the hard disk and RAID controller (if installed) are in danger of failing, or if an additional hard drive is installed that is not a malware storage pack. This module generates an Alert (red) health alert when an installed malware storage pack cannot be detected.
Disk Usage	Any	This module compares disk usage on the appliance's hard drive and malware storage pack to the limits configured for the module and alerts when usage exceeds the percentages configured for the module. This module also alerts when the system excessively deletes files in monitored disk usage categories, or when disk usage excluding those categories reaches excessive levels, based on module thresholds. See Disk Usage and Drain of Events Health Monitor Alerts, on page 319 for information about troubleshooting scenarios for Disk Usage alerts. Use the Disk Usage health status module to monitor disk usage for the `/` and `/volume` partitions on the appliance and track draining frequency. Although the disk usage module lists the `/boot` partition as a monitored partition, the size of the partition is static so the module does not alert on the boot partition. **Attention** If you receive alerts for high unmanaged disk usage for the partition `/volume` even though the usage is below the critical or warning threshold specified in the health policy, this could indicate that there are files which need to be deleted manually from the system. Contact TAC if you receive these alerts.
Hardware Alarms	Firepower 7000/8000 series FTD hardware	This module determines if hardware needs to be replaced on a physical managed device and alerts based on the hardware status. The module also reports on the status of hardware-related daemons and on the status of 7000 and 8000 Series devices in high-availability deployments.

Module	Platforms	Description
HA Status	FMC	This module monitors and alerts on the high availability status of the FMC. If you have not established FMC high availability, the HA Status is `Not in HA`. This module does not monitor or alert on the high availability status of managed devices, regardless of whether they are paired. The HA Status for a managed device is always `Not in HA`. Use the device management page **Devices** > **Device Management** to monitor devices in high availability pairs.
Health Monitor Process	Any	This module monitors the status of the health monitor itself and alerts if the number of minutes since the last health event received by the FMC exceeds the Warning or Critical limits.
Host Limit	FMC	This module determines if the number of hosts the FMC can monitor is approaching the limit and alerts based on the warning level configured for the module. For more information, see Firepower System Host Limit, on page 1832.
Inline Link Mismatch Alarms	Any managed device except ASA FirePOWER	This module monitors the ports associated with inline sets and alerts if the two interfaces of an inline pair negotiate different speeds.
Interface Status	Any	This module determines if the device currently collects traffic and alerts based on the traffic status of physical interfaces and aggregate interfaces. For physical interfaces, the information includes interface name, link state, and bandwidth. For aggregate interfaces, the information includes interface name, number of active links, and total aggregate bandwidth. **Note** This module also monitors the HA standby device traffic flow. Though it is known that the standby device would not be receiving any traffic, yet, the FMC alerts that the interface is not receiving any traffic. The same alerting principle is applied when traffic is not received by some of the subinterfaces on a port channel. If you use the **show interface** CLI command to know the interface statistics of your device, the input and output rates in the CLI command result can be different from the traffic rates that appear in this interface module. This module displays the traffic rates according to the values from Snort performance monitoring. Sampling intervals of snort performance monitoring and the FTD interface statistics are different. This difference in sampling interval results in different throughput values in the FMC GUI and in the FTD **show interface** CLI command result. For ASA FirePOWER, interfaces labeled DataPlaneInterfacex, where x is a numerical value, are internal interfaces (not user-defined) and involve packet flow within the system.

Module	Platforms	Description
Intrusion and File Event Rate	Any managed device	This module compares the number of intrusion events per second to the limits configured for this module and alerts if the limits are exceeded. If the Intrusion and File Event Rate is zero, the intrusion process may be down or the managed device may not be sending events. Select **Analysis** > **Intrusions** > **Events** to check if events are being received from the device. Typically, the event rate for a network segment averages 20 events per second. For a network segment with this average rate, Events per second (Critical) should be set to `50` and Events per second (Warning) should be set to `30`. To determine limits for your system, find the Events/Sec value on the Statistics page for your device (**System** > **Monitoring** > **Statistics**), then calculate the limits using these formulas: • Events per second (Critical) = Events/Sec * 2.5 • Events per second (Warning) = Events/Sec * 1.5 The maximum number of events you can set for either limit is 999, and the Critical limit must be higher than the Warning limit.
Link State Propagation	Firepower 7000/8000 series ASA 5500-X series	This module determines when a link in a paired inline set fails and triggers the link state propagation mode. If a link state propagates to the pair, the status classification for that module changes to Critical and the state reads: `Module Link State Propagation: ethx_ethy is Triggered` where `x` and `y` are the paired interface numbers.
Local Malware Analysis	Any	This module alerts if a device is configured for local malware analysis and fails to download local malware analysis engine signature updates from the AMP cloud.
Memory Usage	Any	This module compares memory usage on the appliance to the limits configured for the module and alerts when usage exceeds the levels configured for the module. For appliances with more than 4 GB of memory, the preset alert thresholds are based on a formula that accounts for proportions of available memory likely to cause system problems. On >4 GB appliances, because the interval between Warning and Critical thresholds may be very narrow, Cisco recommends that you manually set the **Warning Threshold %** value to `50`. This will further ensure that you receive memory alerts for your appliance in time to address the issue. See Memory Usage Thresholds for Health Monitor Alerts, on page 318 for additional information about how thresholds are calculated. Complex access control policies and rules can command significant resources and negatively affect performance. Some lower-end ASA devices with FirePOWER Services Software may generate intermittent memory usage warnings, as the device's memory allocation is being used to the fullest extent possible.

Module	Platforms	Description
Power Supply	FMC hardware Firepower 7000/8000 series	This module determines if power supplies on the device require replacement and alerts based on the power supply status. **Note** If an 8000 Series device experiences a power failure, it may take up to 20 minutes to generate an alert.
Process Status	Any	This module determines if processes on the appliance exit or terminate outside of the process manager. If a process is deliberately exited outside of the process manager, the module status changes to Warning and the health event message indicates which process exited, until the module runs again and the process has restarted. If a process terminates abnormally or crashes outside of the process manager, the module status changes to Critical and the health event message indicates the terminated process, until the module runs again and the process has restarted.
Reconfiguring Detection	Any managed device	This module alerts if a device reconfiguration has failed.
RRD Server Process	FMC	This module determines if the round robin data server that stores time series data is running properly. The module will alert If the RRD server has restarted since the last time it updated; it will enter Critical or Warning status if the number of consecutive updates with an RRD server restart reaches the numbers specified in the module configuration.
Security Intelligence	FMC and devices running Version 6.2.3 or earlier	This module alerts if Security Intelligence is in use and: • The FMC cannot update a feed, or feed data is corrupt or contains no recognizable IP addresses. • A device running Version 6.2.3 or earlier had a problem receiving updated Security Intelligence data from the FMC. • A device running Version 6.2.3 or earlier cannot load all of the Security Intelligence data provided to it by the FMC due to memory issues; see Troubleshooting Memory Use, on page 1260.
Smart License Monitor	FMC	This module alerts if: • There is a communication error between the Smart Licensing Agent (Smart Agent) and the Smart Software Manager. • The Product Instance Registration Token has expired. • The Smart License usage is out of compliance. • The Smart License authorization or evaluation mode has expired.
Time Series Data Monitor	FMC	This module tracks the presence of corrupt files in the directory where time series data (such as correlation event counts) are stored and alerts when files are flagged as corrupt and removed.

Module	Platforms	Description
Time Synchronization Status	Any	This module tracks the synchronization of a device clock that obtains time using NTP with the clock on the NTP server and alerts if the difference in the clocks is more than ten seconds.
URL Filtering Monitor	FMC	This module alerts if the FMC fails to: • Communicate with, or retrieve a URL threat intelligence data update from, Cisco Collective Security Intelligence (CSI). • Push URL threat data to devices.
User Agent Status	FMC	This module alerts when heartbeats are not detected for any User Agents connected to the FMC.
VPN Status	FMC	This module alerts when one or more VPN tunnels between Firepower devices are down. This module tracks: • VPN for 7000/8000 series devices) • Site-to-site VPN for FTD **Attention** Site-to-site VPN tunnels created with Virtual Tunnel Interfaces (VTIs) do not generate health alerts when the tunnel goes down. If you experience packet loss over a VPN with VTIs, check your VPN configuration.

Configuring Health Monitoring

Procedure

Step 1 Determine which health modules you want to monitor as discussed in #unique_230.

You can set up specific policies for each kind of appliance you have in your Firepower System, enabling only the appropriate tests for that appliance.

Tip To quickly enable health monitoring without customizing the monitoring behavior, you can apply the default policy provided for that purpose.

Step 2 Apply a health policy to each appliance where you want to track health status as discussed in Creating Health Policies, on page 283.

Step 3 (Optional.) Configure health monitor alerts as discussed in Creating Health Monitor Alerts, on page 289.

You can set up email, syslog, or SNMP alerts that trigger when the health status level reaches a particular severity level for specific health modules.

Health Policies

A health policy contains configured health test criteria for several modules. You can control which health modules run against each of your appliances and configure the specific limits used in the tests run by each module.

When you configure a health policy, you decide whether to enable each health module for that policy. You also select the criteria that control which health status each enabled module reports each time it assesses the health of a process.

You can create one health policy that can be applied to every appliance in your system, customize each health policy to the specific appliance where you plan to apply it, or use the default health policy provided for you. In a multidomain deployment, administrators in ancestor domains can apply health policies to devices in descendant domains, which descendant domains can use or replace with customized local policies.

Default Health Policy

The Firepower Management Center setup process creates and applies an initial health policy, in which most—but not all—available health modules are enabled. The system also applies this initial policy to devices added to the Firepower Management Center.

This *initial* health policy is based on a *default* health policy, which you can neither view nor edit, but which you can copy when you create a custom health policy.

Upgrades and the Default Health Policy

When you upgrade the FMC, any new health modules are added to all health policies, including the initial health policy, default health policy, and any other custom health policies. Usually, new health modules are added in an enabled state.



Note For a new health module to begin monitoring and alerting, reapply health policies after upgrade.

Creating Health Policies

If you want to customize a health policy to use with your appliances, you can create a new policy. The settings in the policy initially populate with the settings from the health policy you choose as a basis for the new policy. You can enable or disable modules within the policy and change the alerting criteria for each module as needed.

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain. Administrators in ancestor domains can apply health policies to devices in descendant domains, which descendant domains can use or replace with customized local policies.

Procedure

Step 1 Choose **System** > **Health** > **Policy** .

Step 2 Click **New Policy**.

Step 3 Choose the existing policy that you want to use as the basis for the new policy from the **Copy Policy** drop-down list**.**

Step 4 Enter a name for the policy.

Step 5 Enter a description for the policy.

Step 6 Choose **Save** to save the policy information.

Step 7 Choose the module you want to use.

Step 8 Choose **On** for the **Enabled** option to enable use of the module for health status testing.

Step 9 Where appropriate, set the **Critical** and **Warning** criteria.

Step 10 Configure any additional settings for the module. Repeat steps 7-10 for each module.

Step 11 You have three choices:

- To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
- To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
- To temporarily save your changes to this module and switch to another module's settings to modify, choose the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

What to do next

- Apply the health policy to each appliance as described in Applying Health Policies, on page 284. This applies your changes and updates the policy status for all affected policies.

Applying Health Policies

When you apply a health policy to an appliance, the health tests for all the modules you enabled in the policy automatically monitor the health of the processes and hardware on the appliance. Health tests then continue to run at the intervals you configured in the policy, collecting health data for the appliance and forwarding that data to the Firepower Management Center.

If you enable a module in a health policy and then apply the policy to an appliance that does not require that health test, the health monitor reports the status for that health module as disabled.

If you apply a policy with all modules disabled to an appliance, it removes all applied health policies from the appliance so no health policy is applied.

When you apply a different policy to an appliance that already has a policy applied, expect some latency in the display of new data based on the newly applied tests.

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain. Administrators in ancestor domains can apply health policies to devices in descendant domains, which descendant domains can use or replace with customized local policies.

Procedure

Step 1 Choose **System** > **Health** > **Policy** .

Step 2 Click the **Apply** ( ) next to the policy you want to apply.

Tip The **Status** ( ) next to the Health Policy column indicates the current health status for the appliance.

Step 3 Choose the appliances where you want to apply the health policy.

Step 4 Click **Apply** to apply the policy to the appliances you chose.

What to do next

- Optionally, monitor the task status; see Viewing Task Messages, on page 316.

  Monitoring of the appliance starts as soon as the policy is successfully applied.

Editing Health Policies

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain. Administrators in ancestor domains can apply health policies to devices in descendant domains, which descendant domains can use or replace with customized local policies.

Procedure

Step 1 Choose **System** > **Health** > **Policy** .

Step 2 Click **Edit** ( ) next to the policy you want to modify.

Step 3 Edit the **Policy Name** or **Policy Description** fields as desired.

Step 4 Click the health module you want to modify.

Step 5 Modify settings as described in Health Modules, on page 277.

Step 6 You have three options:

- To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
- To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
- To temporarily save your changes to this module and switch to another module's settings to modify, choose the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

Step 7 Apply the health policy to your appliance as described in Applying Health Policies, on page 284.

Apply the health policy to each appliance where you want to track health status. When you apply the health policy to an appliance, all the modules you enabled in the policy monitor the health of the processes and hardware on the appliance, and forwards that data to the FMC.

What to do next

- Reapply the health policy as described in Applying Health Policies, on page 284. This applies your changes and updates the policy status for all affected policies.

Deleting Health Policies

You can delete health policies that you no longer need. If you delete a policy that is still applied to an appliance, the policy settings remain in effect until you apply a different policy. In addition, if you delete a health policy that is applied to a device, any health monitoring alerts in effect for the device remain active until you disable the underlying associated alert response.

In a multidomain deployment, you can only delete health policies created in the current domain.

Tip To stop health monitoring for an appliance, create a health policy with all modules disabled and apply it to the appliance.

Procedure

Step 1 Choose **System** > **Health** > **Policy** .

Step 2 Click **Delete** ( ) next to the policy you want to delete.
A message appears, indicating if the deletion was successful.

The Health Monitor Blocklist

In the course of normal network maintenance, you disable appliances or make them temporarily unavailable. Because those outages are deliberate, you do not want the health status from those appliances to affect the summary health status on your Firepower Management Center.

You can use the health monitor blocklistfeature to disable health monitoring status reporting on an appliance or module. For example, if you know that a segment of your network will be unavailable, you can temporarily disable health monitoring for a managed device on that segment to prevent the health status on the Firepower Management Center from displaying a warning or critical state because of the lapsed connection to the device.

When you disable health monitoring status, health events are still generated, but they have a disabled status and do not affect the health status for the health monitor. If you remove the appliance or module from the blocklist, the events that were generated during the blocklisting continue to show a status of disabled.

To temporarily disable health events from an appliance, go to the blocklist configuration page and add an appliance to the blocklist. After the setting takes effect, the system no longer includes the blocklisted appliance when calculating the overall health status. The Health Monitor Appliance Status Summary lists the appliance as disabled.

You can also disable an individual health module. For example, when you reach the host limit on a Firepower Management Center, you can disable Host Limit status messages.

Note that on the main Health Monitor page you can distinguish between appliances that are blocklisted if you expand to view the list of appliances with a particular status by clicking the arrow in that status row.

A **Blocklist** ( ) icon and a notation are visible after you expand the view for a blocklisted or partially blocklisted appliance.

Note On a Firepower Management Center, Health Monitor blocklist settings are local configuration settings. Therefore, if you blocklist a device, then delete it and later re-register it with the Firepower Management Center, the blocklist settings remain persistent. The newly re-registered device remains blocklisted.

In a multidomain deployment, administrators in ancestor domains can blocklist an appliance or health module in descendant domains. However, administrators in the descendant domains can override the ancestor configuration and clear the blocklist for devices in their domain.

Blocklisting Appliances

You can blocklist appliances individually or by group, model, or associated health policy.

If you need to set the events and health status for an individual appliance to disabled, you can blocklist the appliance. After the blocklist settings take effect, the appliance shows as disabled in the Health Monitor Appliance Module Summary, and health events for the appliance have a status of disabled.

In a multidomain deployment, blocklisting an appliance in an ancestor domain blocklists it for all descendant domains. Descendant domains can override this inherited configuration and clear the blocklist. You can only blocklist the Firepower Management Center at the Global level.

Procedure

Step 1 Choose **System** > **Health** > **Blacklist**.

Step 2 Use the drop-down list on the right to sort the list by group, model, or by health policy.

Tip The status icon next to the Health Policy column **Status** ( ) indicates the current health status for the appliance. The status icon next to the System Policy column **Status** ( ) indicates the communication status between the Firepower Management Center and the device.

Step 3 You have two choices:

- To blocklist all appliances in a group, model, or policy category, check the check box for the category, then click **Blacklist Selected Devices**.
- To clear blocklisting from all appliances in a group, model, or policy category, check the check box for the category, then click **Clear Blacklist on Selected Devices**.

What to do next

To blocklist individual health policy module on appliances, see Blocklisting Health Policy Modules, on page 288.

Blocklisting Health Policy Modules

You can blocklist individual health policy modules on appliances. You may want to do this to prevent events from the module from changing the status for the appliance to warning or critical.

After the blocklist settings take effect, the appliance shows as **Partially Blocklisted** or **All Modules Blocklisted** on the Blocklist page and in the Appliance Health Monitor Module Status Summary, but only in expanded views on the main Appliance Status Summary page.

Tip Make sure that you keep track of individually blocklisted modules so you can reactivate them when you need them. You may miss necessary warning or critical messages if you accidentally leave a module disabled.

In a multidomain deployment, administrators in ancestor domains can blocklist health modules in descendant domains. However, administrators in descendant domains can override this ancestor configuration and clear the blocklisting for policies applied in their domains. You can only blocklist Firepower Management Center health modules at the Global level.

Procedure

Step 1 Choose **System** > **Health** > **Blacklist**.

Step 2 Click **Edit** ( ) next to the appliance you want to modify.

Step 3 Check the check boxes next to the health policy modules you want to blocklist. Certain modules are applicable to specific devices only; for more information, see Health Modules, on page 277.

Step 4 Click **OK**.

Step 5 In the device exclusion main page, click **Apply**.

Health Monitor Alerts

You can set up alerts to notify you through email, through SNMP, or through the system log when the status changes for the modules in a health policy. You can associate an existing alert response with health event levels to trigger and alert when health events of a particular level occur.

For example, if you are concerned that your appliances may run out of hard disk space, you can automatically send an email to a system administrator when the remaining disk space reaches the warning level. If the hard drive continues to fill, you can send a second email when the hard drive reaches the critical level.

In a multidomain deployment, you can view and modify health monitor alerts created in the current domain only.

Health Monitor Alert Information

The alerts generated by the health monitor contain the following information:

- Severity, which indicates the severity level of the alert.

- Module, which specifies the health module whose test results triggered the alert.
- Description, which includes the health test results that triggered the alert.

The table below describes these severity levels.

Table 35: Alert Severities

Severity	Description
Critical	The health test results met the criteria to trigger a Critical alert status.
Warning	The health test results met the criteria to trigger a Warning alert status.
Normal	The health test results met the criteria to trigger a Normal alert status.
Error	The health test did not run.
Recovered	The health test results met the criteria to return to a normal alert status, following a Critical or Warning alert status.

Creating Health Monitor Alerts

You must be an Admin user to perform this procedure.

When you create a health monitor alert, you create an association between a severity level, a health module, and an alert response. You can use an existing alert or configure a new one specifically to report on system health. When the severity level occurs for the selected module, the alert triggers.

If you create or update a threshold in a way that duplicates an existing threshold, you are notified of the conflict. When duplicate thresholds exist, the health monitor uses the threshold that generates the fewest alerts and ignores the others. The timeout value for the threshold must be between 5 and 4,294,967,295 minutes.

In a multidomain deployment, you can view and modify health monitor alerts created in the current domain only.

Before you begin

- Configure an alert response that governs the Firepower Management Center's communication with the SNMP, syslog, or email server where you send the health alert; see Firepower Management Center Alert Responses, on page 2075.

Procedure

Step 1 Choose **System** > **Health** > **Monitor Alerts**.

Step 2 Enter a name for the health alert in the **Health Alert Name** field.

Step 3 From the **Severity** list, choose the severity level you want to use to trigger the alert.

Step 4 From the **Module** list, choose the health policy modules for which you want the alert to apply.

Step 5 From the **Alert** list, choose the alert response that you want to trigger when the specified severity level is reached.

Step 6 Optionally, in the **Threshold Timeout** field, enter the number of minutes that should elapse before each threshold period ends and the threshold count resets.

Even if the policy run time interval value is less than the threshold timeout value, the interval between two reported health events from a given module is always greater. For example, if you change the threshold timeout to 8 minutes and the policy run time interval is 5 minutes, there is a 10-minute interval (5 x 2) between reported events.

Step 7 Click **Save** to save the health alert.

Editing Health Monitor Alerts

You must be an Admin user to perform this procedure.

You can edit existing health monitor alerts to change the severity level, health module, or alert response associated with the health monitor alert.

In a multidomain deployment, you can view and modify health monitor alerts created in the current domain only.

Procedure

Step 1 Choose **System** > **Health** > **Monitor Alerts**.

Step 2 Choose the alert you want to modify from the **Active Health Alerts** list.

Step 3 Click **Load** to load the configured settings for the alert you chose.

Step 4 Modify settings as needed.

Step 5 Click **Save** to save the modified health alert.
A message indicates if the alert configuration was successfully saved.

Deleting Health Monitor Alerts

In a multidomain deployment, you can view and modify health monitor alerts created in the current domain only.

Procedure

Step 1 Choose **System** > **Health** > **Monitor Alerts**.

Step 2 Choose the active health alerts you want to delete, then click **Delete**.

What to do next

- Disable or delete the underlying alert response to ensure that alerting does not continue; see Firepower Management Center Alert Responses, on page 2075.

Using the Health Monitor

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

The health monitor provides the compiled health status for all devices managed by the Firepower Management Center, plus the Firepower Management Center. The health monitor is composed of: The health summary is shown when hovering on the hexagon that representing the device health.

- The status table — Provides a count of the managed appliances for this Firepower Management Center by overall health status.
- The pie chart — Indicates the percentage of appliances currently in each health status category.
- The appliance list — Provides details on the health of the managed devices.

In a multidomain deployment, the health monitor in an ancestor domain displays data from all descendant domains. In the descendant domains, it displays data from the current domain only.

Procedure

Step 1 Choose **System** > **Health** > **Monitor**.

Step 2 Choose the appropriate status in the **Status** column of the table or the appropriate portion of the pie chart to the list appliances with that status.

Tip If the arrow in the row for a status level points down, the appliance list for that status shows in the lower table. If the arrow points right, the appliance list is hidden.

Step 3 You have the following choices:

- View appliance health monitors; see Viewing Appliance Health Monitors, on page 292.
- Create health policies; see Creating Health Policies, on page 283.
- Create health monitor alerts; see Creating Health Monitor Alerts, on page 289.

Health Monitor Status Categories

Available status categories are listed by severity in the table below.

Table 36: Health Status Indicator

Status Level	Status Icon	Status Color in Pie Chart	Description
Error	**Error** (✖)	Black	Indicates that at least one health monitoring module has failed on the appliance and has not been successfully re-run since the failure occurred. Contact your technical support representative to obtain an update to the health monitoring module.

Status Level	Status Icon	Status Color in Pie Chart	Description
Critical	**Critical** ()	Red	Indicates that the critical limits have been exceeded for at least one health module on the appliance and the problem has not been corrected.
Warning	**Warning** ()	Yellow	Indicates that warning limits have been exceeded for at least one health module on the appliance and the problem has not been corrected. This status also indicates a transitionary state, where the required data is temporarily unavailable or could not be processed because of changes in the device configuration. Depending on the monitoring cycle, this transitionary state is auto-corrected.
Normal	**Normal** ()	Green	Indicates that all health modules on the appliance are running within the limits configured in the health policy applied to the appliance.
Recovered	**Recovered** ()	Green	Indicates that all health modules on the appliance are running within the limits configured in the health policy applied to the appliance, including modules that were in a Critical or Warning state.
Disabled	**Disabled** ()	Blue	Indicates that an appliance is disabled or blocked, that the appliance does not have a health policy applied to it, or that the appliance is currently unreachable.

Viewing Appliance Health Monitors

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

The Appliance Health Monitor provides a detailed view of the health status of an appliance.

In a multidomain deployment, you can view the health status of appliances in descendant domains.



Tip Your session normally logs you out after 1 hour of inactivity (or another configured interval). If you plan to passively monitor health status for long periods of time, consider exempting some users from session timeout, or changing the system timeout settings. See User Account Login Options, on page 64 and Configure Session Timeouts, on page 917 for more information.

Procedure

Step 1 Choose **System** > **Health** > **Monitor**.

Step 2 Expand the appliance list. To show appliances with a particular status, click the arrow in that status row. Alternatively, in the **Appliance Status Summary** graph, click the color for the appliance status category you want to view.

Tip If the arrow in the row for a status level points down, the appliance list for that status shows in the lower table. If the arrow points right, the appliance list is hidden.

Step 3 In the **Appliance** column of the appliance list, click the name of the appliance for which you want to view details.

Tip In the **Module Status Summary** graph, click the color for an event status category to toggle display of Alert Details for that status category.

What to do next

- If you want to run all health modules for the appliance, see Running All Modules for an Appliance, on page 293
- If you want to run a specific health module for an appliance, see Running a Specific Health Module, on page 294
- If you want to generate health module alert graphs for the appliance, see Generating Health Module Alert Graphs, on page 294
- If you want to produce troubleshooting files for the appliance, see Downloading Advanced Troubleshooting Files, on page 324
- If you want to download advanced troubleshooting files for the appliance, see Downloading Advanced Troubleshooting Files, on page 324
- If you want to execute Firepower Threat Defense CLI commands from the Firepower Management Center web interface, see Using the Firepower Threat Defense CLI from the Web Interface, on page 325

Running All Modules for an Appliance

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

Health module tests run automatically at the policy run time interval you configure when you create a health policy. However, you can also run all health module tests on demand to collect up-to-date health information for the appliance.

In a multidomain deployment, you can run health module tests for appliances in the current domain and in any descendant domains.

Procedure

Step 1 View the health monitor for the appliance; see Viewing Appliance Health Monitors, on page 292.

Step 2 Click **Run All Modules**. The status bar indicates the progress of the tests, then the Health Monitor Appliance page refreshes.

Note When you manually run health modules, the first refresh that automatically occurs may not reflect the data from the manually run tests. If the value has not changed for a module that you just ran manually, wait a few seconds, then refresh the page by clicking the device name. You can also wait for the page to refresh again automatically.

Running a Specific Health Module

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

Health module tests run automatically at the policy run time interval you configure when you create a health policy. However, you can also run a health module test on demand to collect up-to-date health information for that module.

In a multidomain deployment, you can run health module tests for appliances in the current domain and in any descendant domains.

Procedure

Step 1 View the health monitor for the appliance; see Viewing Appliance Health Monitors, on page 292.

Step 2 In the **Module Status Summary** graph, click the color for the health alert status category you want to view.

Step 3 In the **Alert Detail** row for the alert for which you want to view a list of events, click **Run**.

The status bar indicates the progress of the test, then the Health Monitor Appliance page refreshes.

Note When you manually run health modules, the first refresh that automatically occurs may not reflect the data from the manually run tests. If the value has not changed for a module that you just manually ran, wait a few seconds, then refresh the page by clicking the device name. You can also wait for the page to refresh automatically again.

Generating Health Module Alert Graphs

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

You can graph the results over a period of time of a particular health test for a specific appliance.

Procedure

Step 1 View the health monitor for the appliance; see Viewing Appliance Health Monitors, on page 292.

Step 2 In the **Module Status Summary** graph of the Health Monitor Appliance page, click the color for the health alert status category you want to view.

Step 3 In the **Alert Detail** row for the alert for which you want to view a list of events, click **Graph**.

Tip If no events appear, you may need to adjust the time range.

Health Event Views

The Health Event View page allows you to view health events logged by the health monitor on the Firepower Management Center logs health events. The fully customizable event views allow you to quickly and easily analyze the health status events gathered by the health monitor. You can search event data to easily access other information that may be related to the events you are investigating. If you understand what conditions each health module tests for, you can more effectively configure alerting for health events.

You can perform many of the standard event view functions on the health event view pages.

Viewing Health Events

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

The Table View of Health Events page provides a list of all health events on the specified appliance.

When you access health events from the Health Monitor page on your Firepower Management Center, you retrieve all health events for all managed appliances.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.



Tip You can bookmark this view to allow you to return to the page in the health events workflow containing the Health Events table of events. The bookmarked view retrieves events within the time range you are currently viewing, but you can then modify the time range to update the table with more recent information if needed.

Procedure

Choose **System** > **Health** > **Events**.

Tip If you are using a custom workflow that does not include the table view of health events, click **(switch workflow)**. On the Select Workflow page, click **Health Events**.

Note If no events appear, you may need to adjust the time range.

Viewing Health Events by Module and Appliance

Procedure

Step 1 View the health monitor for the appliance; see Viewing Appliance Health Monitors, on page 292.

Step 2 In the **Module Status Summary** graph, click the color for the event status category you want to view.

The Alert Detail list toggles the display to show or hide events.

Step 3 In the **Alert Detail** row for the alert for which you want to view a list of events, click **Events**.

The Health Events page appears, containing results for a query with the name of the appliance and the name of the specified health alert module as constraints. If no events appear, you may need to adjust the time range.

Step 4 If you want to view all health events for the specified appliance, expand **Search Constraints**, and click the **Module Name** constraint to remove it.

Viewing the Health Events Table

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **System** > **Health** > **Events**.

Step 2 You have the following choices:

- Bookmark — To bookmark the current page so that you can quickly return to it, click **Bookmark This Page**, provide a name for the bookmark, and click **Save**.
- Change Workflow — To choose another health events workflow, click **(switch workflow)**.
- Delete Events — To delete health events, check the check box next to the events you want to delete, and click **Delete**. To delete all the events in the current constrained view, click **Delete All**, then confirm you want to delete all the events.
- Generate Reports — Generate a report based on data in the table view — click **Report Designer**.
- Modify — Modify the time and date range for events listed in the Health table view. Note that events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.
- Navigate — Navigate through event view pages.
- Navigate Bookmark — To navigate to the bookmark management page, click **View Bookmarks** from any event view.
- Navigate Other — Navigate to other event tables to view associated events.
- Sort — Sort the events that appear, change what columns display in the table of events, or constrain the events that appear
- View All — To view event details for all events in the view, click **View All**.

- View Details — To view the details associated with a single health event, click the down arrow link on the left side of the event.
- View Multiple — To view event details for multiple health events, choose the check box next to the rows that correspond with the events you want to view details for and then click **View**.
- View Status — To view all events of a particular status, click status in the Status column for an event with that status.

Hardware Alert Details for 7000 and 8000 Series Devices



Note The 8350 hardware platform has six fans, which display as FAN2 through FAN7. This is expected behavior. If you receive a hardware alert related to FAN1 or fan numbering in general on the 8350 platform, you can disregard the alert.

Table 37: Conditions Monitored for 7000 and 8000 SeriesDevices

Condition Monitored	Causes of Yellow or Red Error Conditions
Device high availability status	If 7000 or 8000 Series devices in a high-availability pair are no longer communicating with each other (due, for example, to a cabling problem), the Hardware Alarms module changes to red.
`ftwo` daemon status	If the `ftwo` daemon goes down, health status for the Hardware Alarms module changes to red and message details include a reference to the daemon.
NFE cards detected	Indicates the number of NFE cards detected on the system. If this value does not match the appliance's expected NFE count, the Hardware Alarms module changes to red.
NFE hardware status	If one or more NFE cards are not communicating, the Hardware Alarms module changes to red and the applicable card appears in the message details.
NFE heartbeat	If the system detects no NFE heartbeat, the Hardware Alarms module changes to red and message details include a reference to the relevant card(s).
NFE internal link status	If the link between the NMSB and NFE card(s) goes down, the Hardware Alarms module changes to red and message details include a reference to the relevant ports.
NFE Message daemon	If the NFE Message daemon goes down, health status for the Hardware Alarms module changes to red and the message details include a reference to the daemon (and, if applicable, the NFE card number).

Condition Monitored	Causes of Yellow or Red Error Conditions
NFE temperature	If NFE temperature exceeds 97 degrees Celsius, health status for the Hardware Alarms module changes to yellow and message details include a reference to the NFE temperature (and, if applicable, the NFE card number). If NFE temperature exceeds 102 degrees Celsius, health status for the Hardware Alarms module changes to red and message details include a reference to the NFE temperature. (and, if applicable, the NFE card number).
NFE temperature status	Indicates the current temperature status of the given NFE card. The Hardware Alarms module indicates green for OK, yellow for Warning, and red for Critical (and, if applicable, the NFE card number).
NFE `TCAM` daemon	If the NFE `TCAM` daemon goes down, health status for the Hardware Alarms module changes to red and message details include a reference to the daemon (and, if applicable, the NFE card number).
`nfm_ipfragd` (host frag) daemon	If the `nfm_ipfragd` daemon goes down, health status for the Hardware Alarms module changes to red and message details include a reference to the daemon (and, if applicable, the NFE card number).
NFE Platform daemon	If the NFE Platform daemon goes down, health status for the Hardware Alarms module changes to red and message details include a reference to the daemon (and, if applicable, the NFE card number).
NMSB communications	If the Media assembly is not present or not communicating, health status for the Hardware Alarms module changes to red and message details include a reference to the NFE temperature (and, if applicable, the NFE card number).
`psls` daemon status	If the `psls` daemon goes down, health status for the Hardware Alarms module changes to red and message details include a reference to the daemon.
`Rulesd` (host rules) daemon	If the `Rulesd` daemon goes down, health status for the Hardware Alarms module changes to yellow and message details include a reference to the daemon (and, if applicable, the NFE card number).
`scmd` daemon status	If the `scmd` daemon goes down, health status for the Hardware Alarms module changes to red and message details include a reference to the daemon.

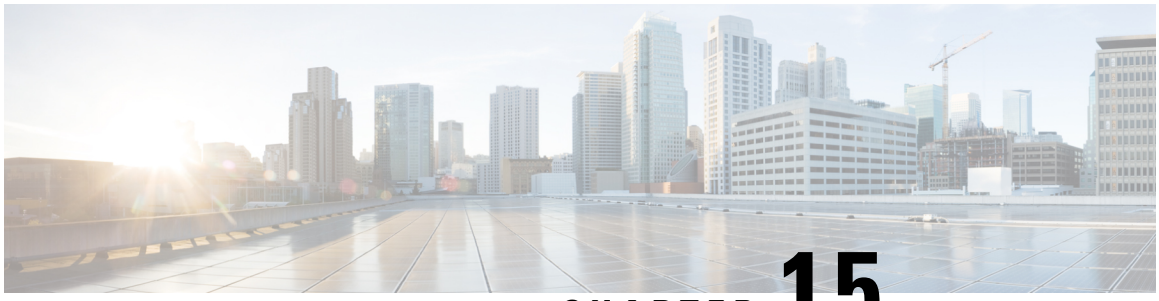
The Health Events Table

The Health Monitor modules you choose to enable in your health policy run various tests to determine appliance health status. When the health status meets criteria that you specify, a health event is generated.

The table below describes the fields that can be viewed and searched in the health events table.

Table 38: Health Event Fields

Field	Description
Module Name	Specify the name of the module which generated the health events you want to view. For example, to view events that measure CPU performance, type `CPU`. The search should retrieve applicable CPU Usage and CPU temperature events.
Test Name (Search only)	The name of the health module that generated the event.
Time (Search only)	The timestamp for the health event.
Description	The description of the health module that generated the event. For example, health events generated when a process was unable to execute are labeled `Unable to Execute`.
Value	The value (number of units) of the result obtained by the health test that generated the event. For example, if the Firepower Management Center generates a health event whenever a device it is monitoring is using 80 percent or more of its CPU resources, the value could be a number from 80 to 100.
Units	The units descriptor for the result. You can use the asterisk (*) to create wildcard searches. For example, if the Firepower Management Center generates a health event when a device it is monitoring is using 80 percent or more of its CPU resources, the units descriptor is a percentage sign (%).
Status	The status (Critical, Yellow, Green, or Disabled) reported for the appliance.
Domain	For health events reported by managed devices, the domain of the device that reported the health event. For health events reported by the Firepower Management Center, `Global`. This field is only present in a multidomain deployment.
Device	The appliance where the health event was reported.



CHAPTER 15

Monitoring the System

The following topics describe how to monitor the Firepower System:

- About System Statistics, on page 301
- The Host Statistics Section, on page 301
- The Disk Usage Section, on page 302
- The Processes Section, on page 302
- The SFDataCorrelator Process Statistics Section, on page 308
- The Intrusion Event Information Section, on page 308
- Viewing System Statistics, on page 309

About System Statistics

You can view system statistics for the Firepower Management Center and 7000 & 8000 Series devices.

The Statistics page lists the current status of general appliance statistics, including disk usage and system processes, Data Correlator statistics (FMC only), and intrusion event information (FMC only).

The Host Statistics Section

The following table describes the host statistics listed on the Statistics page.

Table 39: Host Statistics

Category	Description
Time	The current time on the system.
Uptime	The number of days (if applicable), hours, and minutes since the system was last started.
Memory Usage	The percentage of system memory that is being used.
Load Average	The average number of processes in the CPU queue for the past 1 minute, 5 minutes, and 15 minutes.
Disk Usage	The percentage of the disk that is being used. Click the arrow to view more detailed host statistics.

Category	Description
Processes	A summary of the processes running on the system.

The Disk Usage Section

The Disk Usage section of the Statistics page provides a quick synopsis of disk usage, both by category and by partition status. If you have a malware storage pack installed on a device, you can also check its partition status. You can monitor this page from time to time to ensure that enough disk space is available for system processes and the database.



Tip You can also use the Disk Usage health monitor on the Firepower Management Center to monitor disk usage and alert on low disk space conditions.

The Processes Section

The Processes section of the Statistics page allows you to see the processes that are currently running on an appliance. It provides general process information and specific information for each running process. You can use the Firepower Management Center's web interface to view the process status for any managed device.

Note that there are two different types of processes that run on an appliance: daemons and executable files. Daemons always run, and executable files are run when required.

Process Status Fields

When you expand the Processes section of the Statistics page, you can also view the following:

Cpu(s)

Lists the following CPU usage information:

- user process usage percentage
- system process usage percentage
- nice usage percentage (CPU usage of processes that have a negative nice value, indicating a higher priority). Nice values indicate the scheduled priority for system processes and can range between -20 (highest priority) and 19 (lowest priority).
- idle usage percentage

Mem

Lists the following memory usage information:

- total number of kilobytes in memory
- total number of used kilobytes in memory

- total number of free kilobytes in memory
- total number of buffered kilobytes in memory

Swap

Lists the following swap usage information:

- total number of kilobytes in swap
- total number of used kilobytes in swap
- total number of free kilobytes in swap
- total number of cached kilobytes in swap

The following table describes each column that appears in the Processes section.

Table 40: Process List Columns

Column	Description
Pid	The process ID number
Username	The name of the user or group running the process
Pri	The process priority
Nice	The *nice* value, which is a value that indicates the scheduling priority of a process. Values range between -20 (highest priority) and 19 (lowest priority)
Size	The memory size used by the process (in kilobytes unless the value is followed by `m`, which indicates megabytes)
Res	The amount of resident paging files in memory (in kilobytes unless the value is followed by `m`, which indicates megabytes)
State	The process state: • D — process is in uninterruptible sleep (usually Input/Output) • N — process has a positive nice value • R — process is runnable (on queue to run) • S — process is in sleep mode • T — process is being traced or stopped • W — process is paging • X — process is dead • Z — process is defunct • < — process has a negative nice value
Time	The amount of time (in hours:minutes:seconds) that the process has been running

Column	Description
Cpu	The percentage of CPU that the process is using
Command	The executable name of the process

Related Topics

System Daemons, on page 304
Executables and System Utilities, on page 305

System Daemons

Daemons continually run on an appliance. They ensure that services are available and spawn processes when required. The following table lists daemons that you may see on the Process Status page and provides a brief description of their functionality.



Note The table below is not an exhaustive list of all processes that may run on an appliance.

Table 41: System Daemons

Daemon	Description
crond	Manages the execution of scheduled commands (cron jobs)
dhclient	Manages dynamic host IP addressing
fpcollect	Manages the collection of client and server fingerprints
httpd	Manages the HTTP (Apache web server) process
httpsd	Manages the HTTPS (Apache web server with SSL) service, and checks for working SSL certificate authentication; runs in the background to provide secure web access to the appl
keventd	Manages Linux kernel event notification messages
klogd	Manages the interception and logging of Linux kernel messages
kswapd	Manages Linux kernel swap memory
kupdated	Manages the Linux kernel update process, which performs disk synchronization
mysqld	Manages database processes
ntpd	Manages the Network Time Protocol (NTP) process
pm	Manages all Firepower System processes, starts required processes, restarts any process th unexpectedly
reportd	Manages reports

Daemon	Description
safe_mysqld	Manages safe mode operation of the database; restarts the database daemon if an error logs runtime information to a file
SFDataCorrelator	Manages data transmission
sfestreamer (FMC only)	Manages connections to third-party client applications that use the Event Streamer
sfmgr	Provides the RPC service for remotely managing and configuring an appliance using a connection to the appliance
SFRemediateD (FMC only)	Manages remediation responses
sftimeserviced (FMC only)	Forwards time synchronization messages to managed devices
sfmbservice	Provides access to the sfmb message broker process running on a remote appliance, usi connection to the appliance. Currently used only by health monitoring to send health eve from a managed device to a Firepower Management Center.
sftroughd	Listens for connections on incoming sockets and then invokes the correct executable (Cisco message broker, sfmb) to handle the request
sftunnel	Provides the secure communication channel for all processes requiring communication appliance
sshd	Manages the Secure Shell (SSH) process; runs in the background to provide SSH acce appliance
syslogd	Manages the system logging (syslog) process

Executables and System Utilities

There are a number of executables on the system that run when executed by other processes or through user action. The following table describes the executables that you may see on the Process Status page.

Table 42: System Executables and Utilities

Executable	Description
awk	Utility that executes programs written in the `awk` programming language
bash	GNU Bourne-Again Shell
cat	Utility that reads files and writes content to standard output
chown	Utility that changes user and group file permissions
chsh	Utility that changes the default login shell
SFDataCorrelator (FMC only)	Analyzes binary files created by the system to generate events, connection data, and network maps
cp	Utility that copies files

Executable	Description
df	Utility that lists the amount of free space on the appliance
echo	Utility that writes content to standard output
egrep	Utility that searches files and folders for specified input; supports extended set of regular expressions not supported in standard grep
find	Utility that recursively searches directories for specified input
grep	Utility that searches files and directories for specified input
halt	Utility that stops the server
httpsdctl	Handles secure Apache Web processes
hwclock	Utility that allows access to the hardware clock
ifconfig	Indicates the network configuration executable. Ensures that the MAC address stays constant
iptables	Handles access restriction based on changes made to the Access Configuration page.
iptables-restore	Handles iptables file restoration
iptables-save	Handles saved changes to the iptables
kill	Utility that can be used to end a session and process
killall	Utility that can be used to end all sessions and processes
ksh	Public domain version of the Korn shell
logger	Utility that provides a way to access the syslog daemon from the command line
md5sum	Utility that prints checksums and block counts for specified files
mv	Utility that moves (renames) files
myisamchk	Indicates database table checking and repairing
mysql	Indicates a database process; multiple instances may appear
openssl	Indicates authentication certificate creation
perl	Indicates a perl process
ps	Utility that writes process information to standard output
sed	Utility used to edit one or more text files
sfheartbeat	Identifies a heartbeat broadcast, indicating that the appliance is active; heartbeat used to maintain contact between a device and Firepower Management Center

Executable	Description
sfmb	Indicates a message broker process; handles communication between Firepower Management Centers and device.
sh	Public domain version of the Korn shell
shutdown	Utility that shuts down the appliance
sleep	Utility that suspends a process for a specified number of seconds
smtpclient	Mail client that handles email transmission when email event notification functionality is enabled
snmptrap	Forwards SNMP trap data to the SNMP trap server specified when SNMP notification functionality is enabled
snort	Indicates that Snort is running
ssh	Indicates a Secure Shell (SSH) connection to the appliance
sudo	Indicates a sudo process, which allows users other than admin to run executables
top	Utility that displays information about the top CPU processes **Note** The CPU usage output of this utility is a split up of different types of usages of the CPU core. You must add both user and system processes usage to know the actual total CPU usage. For example, if the output of top command is: `%Cpu(s): 76.6 us, 22.1 sy, 0.0 ni, 0.0 id, 0.0 wa, 0.0 hi, 1.3 si, 0.0 st` Here, 76.6% of CPU time are used by user processes, 22.1% of CPU time is used by system(kernel) processes. The total CPU usage is 98.7%. Thus, the CPU usage reported in this utility appear to be different from the Health Monitor dashboard. In addition, this utility uses a three seconds interval to calculate the CPU usage. Whereas, the management center health monitor uses one-second intervals.
touch	Utility that can be used to change the access and modification times of specified files
vim	Utility used to edit text files
wc	Utility that performs line, word, and byte counts on specified files

Related Topics

Configure an Access List, on page 896

The SFDataCorrelator Process Statistics Section

On a Firepower Management Center, you can view statistics about the Data Correlator and network discovery processes for the current day. As the managed devices perform data acquisition, decoding, and analysis, the network discovery process correlates the data with the fingerprint and vulnerability databases, then produces binary files that are processed by the Data Correlator running on the Firepower Management Center. The Data Correlator analyzes the information from the binary files, generates events, and creates network maps.

The statistics that appear for network discovery and the Data Correlator are averages for the current day, using statistics gathered between 12:00 AM and 11:59 PM for each device.

The following table describes the statistics displayed for the Data Correlator process.

Table 43: Data Correlator Process Statistics

Category	Description
Events/Sec	Number of discovery events that the Data Correlator receives and processes per second
Connections/Sec	Number of connections that the Data Correlator receives and processes per second
CPU Usage — User (%)	Average percentage of CPU time spent on user processes for the current day
CPU Usage — System (%)	Average percentage of CPU time spent on system processes for the current day
VmSize (KB)	Average size of memory allocated to the Data Correlator for the current day, in kilobytes
VmRSS (KB)	Average amount of memory used by the Data Correlator for the current day, in kilobytes

The Intrusion Event Information Section

On both the Firepower Management Center and managed devices, you can view summary information about intrusion events on the Statistics page. This information includes the date and time of the last intrusion event, the total number of events that have occurred in the past hour and the past day, and the total number of events in the database.



Note The information in the Intrusion Event Information section of the Statistics page is based on intrusion events stored on the managed device rather than those sent to the Firepower Management Center. No intrusion event information is listed on this page if the managed device cannot (or is configured not to) store intrusion events locally.

The following table describes the statistics displayed in the Intrusion Event Information section of the Statistics page.

Table 44: Intrusion Event Information

Statistic	Description
Last Alert Was	The date and time that the last event occurred
Total Events Last Hour	The total number of events that occurred in the past hour
Total Events Last Day	The total number of events that occurred in the past twenty-four hours
Total Events in Database	The total number of events in the events database

Viewing System Statistics

On the Firepower Management Center, the web interface displays statistics for the FMC and any devices it manages. On 7000 and 8000 Series devices, the system displays statistics for that device only.

Before you begin

You must be an Admin or Maintenance user and be in the Global domain to view system statistics.

Procedure

Step 1 Choose **System** > **Monitoring** > **Statistics**.

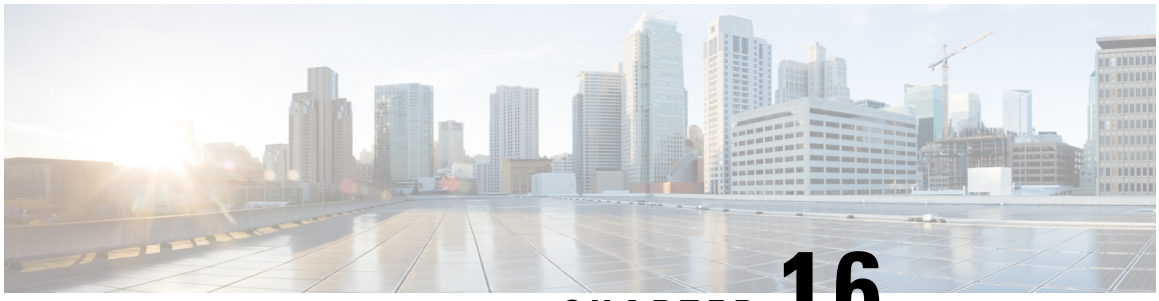
Step 2 (FMC only) Choose a device from the **Select Device(s)** list, and click **Select Devices**.

Step 3 View available statistics.

Step 4 In the Disk Usage section, you can:

- Hover your pointer over a disk usage category in the **By Category** stacked bar to view (in order):
  - the percentage of available disk space used by that category
  - the actual storage space on the disk
  - the total disk space available for that category
- Click the down arrow next to **By Partition** to expand it. If you have a malware storage pack installed, the `/var/storage` partition usage is displayed.

Step 5 (Optional) Click the arrow next to **Processes** to view the information described in Process Status Fields, on page 302.



CHAPTER 16

Troubleshooting the System

The following topics describe ways to diagnose problems you may encounter with the Firepower System:

- First Steps for Troubleshooting, on page 311
- System Messages, on page 311
- View Basic System Information, on page 314
- Managing System Messages, on page 314
- Memory Usage Thresholds for Health Monitor Alerts, on page 318
- Disk Usage and Drain of Events Health Monitor Alerts, on page 319
- Health Monitor Reports for Troubleshooting, on page 322
- Advanced Troubleshooting for the Firepower Threat Defense Device, on page 325
- Feature-Specific Troubleshooting, on page 330

First Steps for Troubleshooting

- Before you make changes to try to fix a problem, generate a troubleshooting file to capture the original problem. See Health Monitor Reports for Troubleshooting, on page 322 and its subsections.

  You may need this troubleshooting file if you need to contact Cisco TAC for support.

- Start your investigation by looking at error and warning messages in the Message Center. See System Messages, on page 311

- Look for applicable Tech Notes and other troubleshooting resources under the "Troubleshoot and Alerts" heading on the product documentation page for your product. See Top-Level Documentation Listing Pages for FMC Deployments, on page 13.

System Messages

When you need to track down problems occurring in the Firepower System, the Message Center is the place to start your investigation. This feature allows you to view the messages that the Firepower System continually generates about system activities and status.

To open the Message Center, click on the System Status icon, located to the immediate right of the Deploy button in the main menu. This icon can take one of the following forms, depending on the system status:

- — Indicates one or more errors and any number of warnings are present on the system.

- — Indicates one or more warnings and no errors are present on the system.
- — Indicates no warnings or errors are present on the system.

If a number is displayed with the icon, it indicates the total current number of error or warning messages.

To close the Message Center, click anywhere outside of it within the Firepower System web interface.

In addition to the Message Center, the web interface displays pop-up notifications in immediate response to your activities and ongoing system activities. Some pop-up notifications automatically disappear after five seconds, while others are "sticky," meaning they display until you explicitly dismiss them by clicking **Dismiss** ( ). Click the **Dismiss** link at the top of the notifications list to dismiss all notifications at once.

Tip Hovering your cursor over a non-sticky pop-up notification causes it to be sticky.

The system determines which messages it displays to users in pop-up notifications and the Message Center based on their licenses, domains, and access roles.

Message Types

The Message Center displays messages reporting system activities and status organized into three different tabs:

Deployments

This tab displays current status related to configuration deployment for each appliance in your system, grouped by domain. The Firepower System reports the following deployment status values on this tab. You can get additional detail about the deployment jobs by clicking **Show History**.

- Running (**Spinning**) — The configuration is in the process of deploying.
- **Success** — The configuration has successfully been deployed.
- **Warning** ( ) — Warning deployment statuses contribute to the message count displayed with the **Warning System Status icon**.
- **Failure** — The configuration has failed to deploy; see Out-of-Date Policies, on page 357. Failed deployments contribute to the message count displayed with the **Error System Status icon**.

Health

This tab displays current health status information for each appliance in your system, grouped by domain. Health status is generated by health modules as described in About Health Monitoring, on page 275. The Firepower System reports the following health status values on this tab:

- **Warning** ( ) — Indicates that warning limits have been exceeded for a health module on an appliance and the problem has not been corrected. The Health Monitoring page indicates these conditions with a **Yellow Triangle** ( ). Warning statuses contribute to the message count displayed with the **Warning System Status icon**.
- **Critical** ( ) — Indicates that critical limits have been exceeded for a health module on an appliance and the problem has not been corrected. The Health Monitoring page indicates these conditions

with a **Critical** () icon. Critical statuses contribute to the message count displayed with the **Error System Status icon**.

- **Error** () — Indicates that a health monitoring module has failed on an appliance and has not been successfully re-run since the failure occurred. The Health Monitoring page indicates these conditions with a **Error icon**. Error statuses contribute to the message count displayed with the **Error System Status icon**.

You can click on links in the Health tab to view related detailed information on the Health Monitoring page. If there are no current health status conditions, the Health tab displays no messages.

Tasks

In the Firepower System, you can perform certain tasks (such as configuration backups or update installation) that can require some time to complete. This tab displays the status of these long-running tasks, and can include tasks initiated by you or, if you have appropriate access, other users of the system. The tab presents messages in reverse chronological order based on the most recent update time for each message. Some task status messages include links to more detailed information about the task in question. The Firepower System reports the following task status values on this tab:

- **Waiting**() — Indicates a task that is waiting to run until another in-progress task is complete. This message type displays an updating progress bar.
- **Running** — Indicates a task that is in-progress. This message type displays an updating progress bar.
- **Retrying** — Indicates a task that is automatically retrying. Note that not all tasks are permitted to try again. This message type displays an updating progress bar.
- **Success** — Indicates a task that has completed successfully.
- **Failure** — Indicates a task that did not complete successfully. Failed tasks contribute to the message count displayed with the **Error System Status icon**.
- **Stopped or Suspended** — Indicates a task that was interrupted due to a system update. Stopped tasks cannot be resumed. After normal operations are restored, start the task again.
- Skipped — A process in progress prevented the task from starting. Try again to start the task.

New messages appear in this tab as new tasks are started. As tasks complete (status success, failure, or stopped), this tab continues to display messages with final status indicated until you remove them. Cisco recommends you remove messages to reduce clutter in the Tasks tab as well as the message database.

Message Management

From the Message Center you can:

- Configure pop-up notification behavior (choosing whether to display them).
- Display additional task status messages from the system database (if any are available that have not been removed).
- Remove individual task status messages. (This affects all users who can view the removed messages.)
- Remove task status messages in bulk. (This affects all users who can view the removed messages.)

Tip Cisco recommends that you periodically remove accumulated task status messages from the Task tab to reduce clutter in the display as well the database. When the number of messages in the database approaches 100,000, the system automatically deletes task status messages that you have removed.

View Basic System Information

The About page displays information about your appliance, including the model, serial number, and version information for various components of the Firepower System. It also includes Cisco copyright information.

Procedure

Step 1 Click **Help** in the toolbar at the top of the page.

Step 2 Choose **About**.

View Appliance Information

Procedure

Choose **System** > **Configuration**.

Managing System Messages

Procedure

Step 1 Click System Status to display the Message Center.

Step 2 You have the following choices:

- Click **Deployments** to view messages related to configuration deployments. See Viewing Deployment Messages, on page 315. You must be an Admin user or have the **Deploy Configuration to Devices** permission to view these messages.
- Click **Health** to view messages related to the health of your Firepower Management Center and the devices registered to it. See Viewing Health Messages, on page 316. You must be an Admin user or have the **Health** permission to view these messages.
- Click **Tasks** to view or manage messages related to long-running tasks. See Viewing Task Messages, on page 316 or Managing Task Messages, on page 317. Everyone can see their own tasks. To see the tasks of other users, you must be an Admin user or have the **View Other Users' Tasks** permission.

- Click **Cog** ( ) in the upper right corner of the Message Center to configure pop-up notification behavior. See Configuring Notification Behavior, on page 317.

Viewing Deployment Messages

You must be an Admin user or have the **Deploy Configuration to Devices** permission to view these messages.

Procedure

Step 1 Click System Status to display the Message Center.

Step 2 Click **Deployments**.

Step 3 You have the following choices:

- Click **total** to view all current deployment statuses.
- Click a status value to view only messages with that deployment status.
- Hover your cursor over the time elapsed indicator for a message (for example, **1m 5s**) to view the elapsed time, and start and stop times for the deployment.

Step 4 Click **show deployment history** to view more detailed information about the deployment jobs.

The Deployment History table lists the deployment jobs in the left column in reverse chronological order.

a) Select a deployment job.

The table in the right column shows each device that was included in the job, and the deployment status per device.

b) To view responses from the device, and commands sent to the device during deployment, click download in the **Transcript** column for the device.

The transcript includes the following sections:

- **Snort Apply**—If there are any failures or responses from Snort-related policies, messages appear in this section. Normally, the section is empty.
- **CLI Apply**—This section covers features that are configured using commands sent to the Lina process.
- **Infrastructure Messages**—This section shows the status of different deployment modules.

In the **CLI Apply** section, the deployment transcript includes commands sent to the device, and any responses returned from the device. These response can be informative messages or error messages. For failed deployments, look for messages that indicate errors with the commands. Examining these errors can be particularly helpful if you are using FlexConfig policies to configure customized features. These errors can help you correct the script in the FlexConfig object that is trying to configure the commands.

Note There is no distinction made in the transcript between commands sent for managed features and those generated from FlexConfig policies.

For example, the following sequence shows that the FMC sent commands to configure GigabitEthernet0/0 with the logical name outside. The device responded that it automatically set the security level to 0. Firepower Threat Defense does not use the security level for anything.

```
========= CLI APPLY =========

FMC >> interface GigabitEthernet0/0
FMC >>  nameif outside
FTDv 192.168.0.152 >> [info] : INFO: Security level for "outside" set to 0 by default.
```

Related Topics

Deploy Configuration Changes, on page 346

Viewing Health Messages

You must be an Admin user or have the **Health** permission to view these messages.

Procedure

Step 1 Click System Status to display the Message Center.

Step 2 Click **Health**.

Step 3 You have the following choices:

- Click **total** to view all current health statuses.
- Click on a status value to view only messages with that status.
- Hover your cursor over the relative time indicator for a message (for example, **3 day(s) ago**) to view the time of the most recent update for that message.
- To view detailed health status information for a particular message, click the message.
- To view complete health status on the Health Monitoring page, click **Health Monitor**.

Related Topics

About Health Monitoring, on page 275

Viewing Task Messages

Everyone can see their own tasks. To see the tasks of other users, you must be an Admin user or have the **View Other Users' Tasks** permission.

Procedure

Step 1 Click System Status to display the Message Center.

Step 2 Click **Tasks**.

Step 3 You have the following choices:

- Click **total** to view all current task statuses.
- Click a status value to view only messages for tasks with the that status.

Note Messages for stopped tasks appear only in the total list of task status messages. You cannot filter on stopped tasks.

- Hover your cursor over the relative time indicator for a message (e.g., **3 day(s) ago**) to view the time of the most recent update for that message.
- Click any link within a message to view more information about the task.
- If more task status messages are available for display, click **Fetch more messages** at the bottom of the message list to retrieve them.

Managing Task Messages

Everyone can see their own tasks. To see the tasks of other users, you must be an Admin user or have the **View Other Users' Tasks** permission.

Procedure

Step 1 Click System Status to display the Message Center.

Step 2 Click Tasks.

Step 3 You have the following choices:

- If more task status messages are available for display, click on **Fetch more messages** at the bottom of the message list to retrieve them.
- To remove a single message for a completed task (status stopped, success, or failure), click on **Remove** (✕) next to the message.
- To remove all messages for all tasks that have completed (status stopped, success, or failure), filter the messages on **total** and click on **Remove all completed tasks**.
- To remove all messages for all tasks that have completed successfully, filter the messages on **success**, and click on **Remove all successful tasks**.
- To remove all messages for all tasks that have failed, filter the messages on **failure**, and click on **Remove all failed tasks**.

Configuring Notification Behavior

Note This setting affects all pop-up notifications and persists between login sessions.

Procedure

Step 1 Click System Status to display the Message Center.

Step 2 Click **Cog** ( ) in the upper right corner of the Message Center.

Step 3 To enable or disable pop-up notification display, click the **Show notifications** slider.

Step 4 Click **Cog** ( ) again to hide the slider.

Step 5 Click System Status again to close the Message Center.

Memory Usage Thresholds for Health Monitor Alerts

The Memory Usage health module compares memory usage on an appliance to the limits configured for the module and alerts when usage exceeds the levels. The module monitors data from managed devices and from the FMC itself.

Two configurable thresholds for memory usage, Critical and Warning, can be set as a percentage of memory used. When these thresholds are exceeded, a health alarm is generated with the severity level specified. However, the health alarm system does not calculate these thresholds in an exact manner.

With high memory devices, certain processes are expected to use a larger percentage of total system memory than in a low memory footprint device. The design is to use as much of the physical memory as possible while leaving a small value of memory free for ancillary processes.

Compare two devices, one with 32 GB of memory and one with 4 GB of memory. In the device with 32 GB of memory, 5% of memory (1.6GB) is a much larger value of memory to leave for ancillary processes than in the device with 4 GB of memory (5% of 4GB = 200MB).

To account for the higher percentage use of system memory by certain processes, the FMC calculates the total memory to include both total physical memory and total swap memory. Thus the enforced memory threshold for the user-configured threshold input can result in a Health Event where the "Value" column of the event does not match the value that was entered to determine the exceeded threshold.

The following table shows examples of user-input thresholds and the enforced thresholds, depending on the installed system memory.

Note The values in this table are examples. You can use this information to extrapolate thresholds for devices that do not match the installed RAM shown here, or you can contact Cisco TAC for more precise threshold calculations.

Table 45: Memory Usage Thresholds Based On Installed RAM

User-input Threshold Value	Enforced Threshold Per Installed Memory (RAM)			
	4 GB	6 GB	32 GB	48 GB
10%	10%	34%	72%	81%

User-input Threshold Value	Enforced Threshold Per Installed Memory (RAM)			
	4 GB	6 GB	32 GB	48 GB
20%	20%	41%	75%	83%
30%	30%	48%	78%	85%
40%	40%	56%	81%	88%
50%	50%	63%	84%	90%
60%	60%	70%	88%	92%
70%	70%	78%	91%	94%
80%	80%	85%	94%	96%
90%	90%	93%	97%	98%
100%	100%	100%	100%	100%

Disk Usage and Drain of Events Health Monitor Alerts

The Disk Usage health module compares disk usage on a managed device's hard drive and malware storage pack to the limits configured for the module and alerts when usage exceeds the percentages configured for the module. This module also alerts when the system excessively deletes files in monitored disk usage categories, or when disk usage excluding those categories reaches excessive levels, based on module thresholds.

This topic describes the symptoms and troubleshooting guidelines for two health alerts generated by the Disk Usage health module:

- Frequent Drain of Events
- Drain of Unprocessed Events

The disk manager process manages the disk usage of a device. Each type of file monitored by the disk manager is assigned a silo. Based on the amount of disk space available on the system the disk manager computes a High Water Mark (HWM) and a Low Water Mark (LWM) for each silo.

To display detailed disk usage information for each part of the system, including silos, LWMs, and HWMs, use the **show disk-manager** command.

Examples

Following is an example of the disk manager information.

```
> show disk-manager
Silo                                Used        Minimum     Maximum
Temporary Files                     0 KB        499.197 MB  1.950 GB
Action Queue Results                0 KB        499.197 MB  1.950 GB
User Identity Events                0 KB        499.197 MB  1.950 GB
UI Caches                           4 KB        1.462 GB    2.925 GB
Backups                             0 KB        3.900 GB    9.750 GB
```

```
Updates                                 0 KB        5.850 GB    14.625 GB
Other Detection Engine                  0 KB        2.925 GB    5.850 GB
Performance Statistics                  33 KB       998.395 MB  11.700 GB
Other Events                            0 KB        1.950 GB    3.900 GB
IP Reputation & URL Filtering           0 KB        2.437 GB    4.875 GB
Archives & Cores & File Logs            0 KB        3.900 GB    19.500 GB
Unified Low Priority Events             1.329 MB    4.875 GB    24.375 GB
RNA Events                              0 KB        3.900 GB    15.600 GB
File Capture                            0 KB        9.750 GB    19.500 GB
Unified High Priority Events            0 KB        14.625 GB   34.125 GB
IPS Events                              0 KB        11.700 GB   29.250 GB
```

Health Alert Format

When the Health Monitor process on the FMC runs (once every 5 minutes or when a manual run is triggered) the Disk Usage module looks into the diskmanager.log file and, if the correct conditions are met, the respective health alert is triggered.

The structures of these health alerts are as follows:

- Frequent drain of <*SILO NAME*>
- Drain of unprocessed events from <*SILO NAME*>

For example,

- Frequent drain of Low Priority Events
- Drain of unprocessed events from Low Priority Events

It's possible for any silo to generate a *Frequent drain of <SILO NAME>* health alert. However, the most commonly seen are the alerts related to events. Among the event silos, the *Low Priority Events* are often seen because these type of events are generated by the device more frequently.

A *Frequent drain of <SILO NAME>* event has a **Warning** severity level when seen in relation to an event-related silo, because events will be queued to be sent to the FMC. For a non-event related silo, such as the *Backups* silo, the alert has a **Critical** severity level because this information is lost.



Important Only event silos generate a *Drain of unprocessed events from <SILO NAME>* health alert. This alert always has **Critical** severity level.

Additional symptoms besides the alerts can include:

- Slowness on the FMC user interface
- Loss of events

Common Troubleshoot Scenarios

A *Frequent drain of <SILO NAME>* event is caused by too much input into the silo for its size. In this case, the disk manager drains (purges) that file at least twice in the last 5-minute interval. In an event type silo, this is typically caused by excessive logging of that event type.

In the case of a *Drain of unprocessed events of <SILO NAME>* health alert, this can also be caused by a bottleneck in the event processing path.

There are three potential bottlenecks with respect to these Disk Usage alerts:

- Excessive logging — The EventHandler process on FTD is oversubscribed (it reads slower than what Snort writes).
- Sftunnel bottleneck — The Eventing interface is unstable or oversubscribed.
- SFDataCorrelator bottleneck — The data transmission channel between the FMC and the managed device is oversubscribed.

Excessive Logging

One of the most common causes for the health alerts of this type is excessive input. The difference between the Low Water Mark (LWM) and High Water Mark (HWM) gathered from the **show disk-manager** command shows how much space there is available to take on that silo to go from LWM (freshly drained) to the HWM value. If there are frequent drain of events (with or without unprocessed events) the first thing to review is the logging configuration.

- Check for double logging — Double logging scenarios can be identified if you look at the correlator *perfstats* on the FMC:

  ```
  admin@FMC:~$ sudo perfstats -Cq < /var/sf/rna/correlator-stats/now
  ```

- Check logging settings for the ACP — Review the logging settings of the Access Control Policy (ACP). If logging both "Beginning" and "End" of connection, log only the end as it will include everything included when the beginning is logged as well as reduce the amount of events.

  Ensure that you follow the best practices described in Best Practices for Connection Logging, on page 2216.

Communications Bottleneck – Sftunnel

Sftunnel is responsible for encrypted communications between the FMC and the managed device. Events are sent over the tunnel to the FMC. Connectivity issues and/or instability in the communication channel (sftunnel) between the managed device and the FMC can be due to:

- Sftunnel is down or is unstable (flaps).

  Ensure that the FMC and the managed device have reachability between their management interfaces on TCP port 8305.

  The sftunnel process should be stable and should not restart unexpectedly. Verify this by checking the **/var/log/message** file and search for messages that contain the ***sftunneld*** string.

- Sftunnel is oversubscribed.

  Review trend data from the Heath Monitor and look for signs of oversubscription of the FMC's management interface, which can be a spike in management traffic or a constant oversubscription.

  Use as a secondary management interface for Firepower-eventing. To use this interface, you must configure its IP address and other parameters at the FTD CLI using the **configure network management-interface** command.

Communications Bottleneck – SFDataCorrelator

The SFDataCorrelator manages data transmission between the FMC and the managed device; on the FMC, it analyzes binary files created by the system to generate events, connection data, and network maps. The first step is to review the **diskmanager.log** file for important information to be gathered, such as:

- The frequency of the drain.
- The number of files with Unprocessed Events drained.
- The occurrence of the drain with Unprocessed Events.

Each time the disk manager process runs it generates an entry for each of the different silos on its own log file, which is located under **[/ngfw]/var/log/diskmanager.log**. Information gathered from the diskmanager.log (in CSV format) can be used to help narrow the search for a cause.

Additional troubleshooting steps:

- The command **stats_unified.pl** can help you to determine if the managed device does have some data which needs to be sent to FMC. This condition can happen when the managed device and the FMC experience a connectivity issue. The managed device stores the log data onto a hard drive.

  ```
  admin@FMC:~$ sudo stats_unified.pl
  ```

- The **manage_proc.pl** command can reconfigure the correlator on the FMC side.

  ```
  root@FMC:~# manage_procs.pl
  ```

Before You Contact Cisco Technical Assistance Center (TAC)

It is highly recommended to collect these items before you contact Cisco TAC:

- Screenshots of the health alerts seen.
- Troubleshoot file generated from the FMC.
- Troubleshoot file generated from the affected managed device.

  Date and Time when the problem was first seen.

- Information about any recent changes done to the policies (if applicable).

  The output of the stats_unified.pl command as described in the Communications Bottleneck — SFDataCorrelator, on page 322.

Health Monitor Reports for Troubleshooting

In some cases, if you have a problem with your appliance, Support may ask you to supply troubleshooting files to help them diagnose the problem. The system can produce troubleshooting files with information targeted to specific functional areas, as well as advanced troubleshooting files you retrieve in cooperation with Support. You can select any of the options listed in the table below to customize the contents of a troubleshooting file for a specific function.

Note that some options overlap in terms of the data they report, but the troubleshooting files will not contain redundant copies, regardless of what options you select.

Table 46: Selectable Troubleshoot Options

This option...	Reports...
Snort Performance and Configuration	data and configuration settings related to Snort on the appliance
Hardware Performance and Logs	data and logs related to the performance of the appliance hardware
System Configuration, Policy, and Logs	configuration settings, data, and logs related to the current system configuration of the appliance
Detection Configuration, Policy, and Logs	configuration settings, data, and logs related to detection on the appliance
Interface and Network Related Data	configuration settings, data, and logs related to inline sets and network configuration of the appliance
Discovery, Awareness, VDB Data, and Logs	configuration settings, data, and logs related to the current discovery and awareness configuration on the appliance
Upgrade Data and Logs	data and logs related to prior upgrades of the appliance
All Database Data	all database-related data that is included in a troubleshoot report
All Log Data	all logs collected by the appliance database
Network Map Information	current network topology data

Producing Troubleshooting Files for Specific System Functions

You can generate and download customized troubleshooting files that you can send to Support.

In a multidomain deployment, you can generate and download troubleshooting files for devices in descendant domains.

Caution Generating troubleshooting files for lower-memory devices can trigger Automatic Application Bypass (AAB) when AAB is enabled, At a minimum, triggering AAB restarts the Snort process, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information. In some such cases, triggering AAB can render the device temporarily inoperable. If inoperability persists, contact Cisco Technical Assistance Center (TAC), who can propose a solution appropriate to your deployment. Susceptible devices include Firepower 7010, 7020, and 7030; ASA 5506-X, 5508-X, 5516-X, 5512-X , 5515-X, and 5525-X; NGIPSv.

Before you begin

You must be an Admin, Maintenance, Security Analyst, or Security Analyst (Read Only) user to perform this task.

Procedure

Step 1 Perform the steps in Viewing Appliance Health Monitors, on page 292.

Step 2 Click **Generate Troubleshooting Files.**

Step 3 Choose All Data to generate all possible troubleshooting data, or check individual boxes as described in Viewing Task Messages, on page 316.

Step 4 Click **OK.**

Step 5 View task messages in the Message Center; see Viewing Task Messages, on page 316.

Step 6 Find the task that corresponds to the troubleshooting files you generated.

Step 7 After the appliance generated the troubleshooting files and the task status changes to `Completed`, click **Click to retrieve generated files**.

Step 8 Follow your browser's prompts to download the file. (The troubleshooting files are downloaded in a single `.tar.gz` file.)

Step 9 Follow the directions from Support to send the troubleshooting files to Cisco.

Downloading Advanced Troubleshooting Files

In a multidomain deployment, you can generate and download troubleshooting files for devices in descendant domains. You can download files from the Firepower Management Center only from the Global domain.

Before you begin

You must be an Admin, Maintenance, Security Analyst, or Security Analyst (Read Only) user to perform this task.

Procedure

Step 1 View the health monitor for the appliance; see Viewing Appliance Health Monitors, on page 292.

Step 2 Click **Advanced Troubleshooting.**

Step 3 In **File Download**, enter the file name supplied by Support.

Step 4 Click **Download.**

Step 5 Follow your browser's prompts to download the file.

Note For managed devices, the system renames the file by prepending the device name to the file name.

Step 6 Follow the directions from Support to send the troubleshooting files to Cisco.

Advanced Troubleshooting for the Firepower Threat Defense Device

You can use Packet Tracer and Packet Capture features for performing an in-depth troubleshooting analysis on a Firepower Threat Defense device. Packet-tracer allows a firewall administrator to inject a virtual packet into the security appliance and track the flow from ingress to egress. Along the way, the packet is evaluated against flow and route lookups, ACLs, protocol inspection, NAT, and intrusion detection. The power of the utility comes from the ability to simulate real-world traffic by specifying source and destination addresses with protocol and port information. Packet capture is available with the trace option, which provides you with a verdict as to whether the packet is dropped or successful.

For more information about the troubleshooting files, see Downloading Advanced Troubleshooting Files, on page 324.

Using the Firepower Threat Defense CLI from the Web Interface

You can execute selected Firepower Threat Defense command line interface (CLI) commands from the Firepower Management Center web interface. These commands are **ping**, **traceroute**, and **show** (except for **show history** and **show banner**).

In a multidomain deployment, you can enter Firepower Threat Defense CLI commands through the Firepower Management Center web interface for managed devices in descendant domains.

Note In deployments using Firepower Management Center high availability, this feature is available only in the active Firepower Management Center.

For more information on the Firepower Threat Defense CLI, see the *Command Reference for Firepower Threat Defense*.

Before you begin

You must be an Admin, Maintenance, or Security Analyst user to use the CLI.

Procedure

Step 1 View the health monitor for the appliance; see Viewing Appliance Health Monitors, on page 292.

Step 2 Click **Advanced Troubleshooting.**

Step 3 Click **Threat Defense CLI**.

Step 4 From the **Command** drop-down list, select a command.

Step 5 Optionally, enter command parameters in the **Parameters** text box.

Step 6 Click **Execute** to view the command output.

Packet Tracer Overview

Using the packet tracer, you can test your policy configuration by modeling a packet based on source and destination addressing, and protocol characteristics. The trace does a policy lookup to test access rules, NAT, routing, access policies and rate limiting policies, to check if the packet would be permitted or denied. The packet flow is simulated based on interfaces, source address, destination address, ports and protocols. By testing packets this way, you can see the results of your policies and test whether the types of traffic you want to allow or deny are handled as desired. Besides verifying your configuration, you can use the tracer to debug unexpected behavior, such as packets being denied when they should be allowed. To simulate the packet fully, packet tracer traces the data path; slow-path and fast-path modules. Processing is transacted based on per-session and per-packet basis. Tracing packets and capture with trace log the tracing data on per packet basis when the Next-Generation Firewall (NGFW) processes packets per-session or per-packet.

Use the Packet Tracer

You can use packet tracer on Firepower Threat Defense devices. You must be an Admin or Maintenance user to use this tool.

Procedure

Step 1 On the Firepower Management Center, choose **Devices** > **Device Management**.

Step 2 Select a device.

Step 3 Click troubleshooting.
The **Health Monitor** page appears.

Step 4 Click **Advanced Troubleshooting**.

Step 5 Click **Packet Tracer**.

Step 6 Select the **Packet type** for the trace, and specify the protocol characteristics:

- **ICMP**—Enter the ICMP type, ICMP code (0-255), and optionally, the ICMP identifier.
- **TCP**/**UDP**/**SCTP**—Enter the source and destination port numbers.
- **IP**—Enter the protocol number, 0-255.

Step 7 Select the ingress **Interface** for the packet trace.

Step 8 Select the **Source** type for the packet trace, and enter the source IP address.

Source and destination types include IPv4, IPv6, and fully-qualified domain names (FQDN). You can specify IPv4 or IPv6 addresses and FQDN, if you use Cisco TrustSec.

Step 9 Select the **Source Port** for the packet trace.

Step 10 Select the **Destination** type for the packet trace, and enter the destination IP address.

Step 11 Select the **Destination Port** for the packet trace.

Step 12 Optionally, if you want to trace a packet where the Security Group Tag (SGT) value is embedded in the Layer 2 CMD header (TrustSec), enter a valid **SGT number**.

Step 13 If you want packet tracer to enter a parent interface, which is later redirected to a sub-interface, enter a **VLAN ID**.

This value is optional for non-sub-interfaces only, since all the interface types can be configured on a sub-interface.

Step 14 Specify a **Destination MAC Address** for the packet trace.

If the Firepower Threat Defense device is running in transparent firewall mode, and the ingress interface is VTEP, **Destination MAC Address** is required if you enter a value in **VLAN ID**. Whereas if the interface is a bridge group member, **Destination MAC Address** is optional if you enter a **VLAN ID** value, but required if you do not enter a **VLAN ID** value.

If the Firepower Threat Defense is running in routed firewall mode, **VLAN ID** and **Destination MAC Address** are optional if the input interface is a bridge group member.

Step 15 Select the **Output Format** for the packet logs.

Step 16 Click **Start**.

Packet Capture Overview

The packet capture feature with trace option allows real packets that are captured on the ingress interface to be traced through the system. The trace information is displayed at a later stage. These packets are not dropped on the egress interface, as they are real data-path traffic. Packet capture for Firepower Threat Defense devices supports troubleshooting and analysis of data packets.

Once the packet is acquired, Snort detects the tracing flag that is enabled in the packet. Snort writes tracer elements, through which the packet traverses. Snort verdict as a result of capturing packets can be one of .the following:

Table 47: Snort Verdicts

Verdict	Description
Pass	Allow analyzed packet.
Block	Packet not forwarded.
Replace	Packet modified.
White list	Flow passed without inspection.
Blacklist	Flow was blocked.
Ignore	Flow was blocked; occurs only for sessions with flows blocked on passive interfaces.
Retry	Flow is stalled, waiting on a enamelware or URL category/reputation query. In the event of a timeout, processing continues with an unknown result: in the case of enamelware, the file is allowed; in the case of URL category/reputation, AC rule lookup continues with an uncategorized and unknown reputation.

Based on the Snort verdict, the packets are dropped or allowed. For example, the packet is dropped if the Snort verdict is **Blacklist**, and the subsequent packets in the session are dropped before reaching Snort. When the Snort verdict is **Block** or **Blacklist**, the **Drop Reason** can be one of the following:

Table 48: Drop Reasons

Blocked or Flow Blocked by...	Cause
Snort	Snort is unable to process the packet, erg., snort can't decode packet since it is corrupted or has invalid format.
the App Id preprocessed	App Id module/preprocessed does not block packet by itself; but this may indicate that App Id detection causes other module (erg., firewall) to match a blocking rule.
the SSL preprocessed	There is a block/reset rule in SSL policy to match the traffic.
the firewall	There is a block/reset rule in firewall policy to match the traffic.
the captive portal preprocessed	There is a block/reset rule using the identity policy to match the traffic.
the safe search preprocessed	There is a block/reset rule using the safe-search feature in firewall policy to match the traffic.
the SI preprocessed	There is a block/reset rule a in Security Intelligence tab of AC Policy to block the traffic, erg., DNS or URL SI rule.
the filterer preprocessed	There is a block/reset rule in filterer tab of AC policy to match the traffic.
the stream preprocessed	There is an intrusion rule blocking/reset stream connection, erg., blocking when TCP normalization error.
the session preprocessed	This session was already blocked earlier by some other module, so session preprocessed is blocking further packets of the same session.
the fragmentation preprocessed	Blocking because earlier fragment of the data is blocked.
the snort response preprocessed	There is a react snort rule, erg., sending a response page on a particular HTTP traffic.
the snort response preprocessed	There is a snort rule to send custom response on packets matching conditions.
the reputation preprocessed	Packet matches a reputation rule, erg., blocking a given IP address.
the x-Link2State preprocessed	Blocking due to buffer overflow vulnerability detected in SMTP.

Blocked or Flow Blocked by...	Cause
back orifice preprocessed	Blocking due to detection of back orifice data.
the SMB preprocessed	There is a snort rule to block SMB traffic.
the file process preprocessed	There is file policy that blocks a file, erg., enamelware blocking.
the IPS preprocessed	There is a snort rule using IPS, erg., rate filtering.

The **file-size** option is used when you need to capture packets with the size limit more than 32 MB.

Use the Capture Trace

Packet capture data includes information from Snort and preprocessors about verdicts and actions the system takes while processing a packet. Multiple packet captures are possible at a time. You can configure the system to modify, delete, clear, and save captures.



Note Capturing packet data requires packet copy. This operation may cause delays while processing packets and may also degrade the packet throughput. Cisco recommends that you use packet filters to capture specific traffic data.

The saved traffic data can be downloaded in *pcap* or *ASCII* file formats.

Before you begin

You can use packet capture on Firepower Threat Defense devices. You must be an Admin or Maintenance user to use this tool.

Procedure

Step 1 On the Firepower Management Center, choose **Devices** > **Device Management**.

Step 2 Select a device.

Step 3 Click troubleshooting.
The **Health Monitor** page appears.

Step 4 Click **Advanced Troubleshooting**.

Step 5 Select **Capture w/Trace**.

Step 6 Click **Add Capture**.

Step 7 Enter the **Name** for capturing the trace.

Step 8 Select the **Interface** for the capturing the trace.

Step 9 Specify **Match Criteria** details:

a) Select the **Protocol**.
b) Enter the IP address for the **Source Host**.
c) Enter the IP address for the **Destination Host**.
d) (Optional) Check **SGT number** check box, and enter a Security Group Tag (SGT).

Step 10 Specify **Buffer** details:

a) (Optional) Enter a maximum **Packet Size**.
b) (Optional) Enter a minimum **Buffer Size**.
c) Select either **Continuous Capture** if you want the traffic captured without interruption, or **Stop when full** if you want the capture to stop when the maximum buffer size is reached.
d) Select **Trace** if you want to capture the details for each packet.
e) (Optional) Check **Trace Count** check box. Default value is 50. You can enter values in the range of 1-1000.

Step 11 Click **Save**.

Feature-Specific Troubleshooting

See the following table for feature-specific troubleshooting tips and techniques.

Table 49: Feature-Specific Troubleshooting Topics

Feature	Relevant Troubleshooting Information
Application control	Troubleshoot Application Control Rules, on page 383
LDAP external authentication	Troubleshooting LDAP Authentication Connections, on page 92
Licensing	Troubleshoot FTD Licensing, on page 123
FMC high availability	Troubleshooting Firepower Management Center High Availability, on page 204
7000 and 8000 Series device high-availability state sharing	Device High Availability State Sharing Statistics for Troubleshooting, on page 529
User rule conditions	Troubleshoot User Control, on page 387
User identity sources	Troubleshoot the User Agent Identity Source, on page 1894 Troubleshoot ISE or Cisco TrustSec Issues, on page 1900 Troubleshoot the TS Agent Identity Source, on page 1902 Troubleshoot the Captive Portal Identity Source, on page 1914 Troubleshooting LDAP Authentication Connections, on page 92
URL filtering	Troubleshoot URL Filtering, on page 1241
Realms and user data downloads	Troubleshoot Realms and User Downloads, on page 1945
Network discovery	Troubleshooting Your Network Discovery Strategy, on page 1938
Custom Security Group Tag (SGT) rule conditions	Troubleshooting Custom SGT Conditions, on page 390
SSL rules	Troubleshoot TLS/SSL Rules, on page 1393

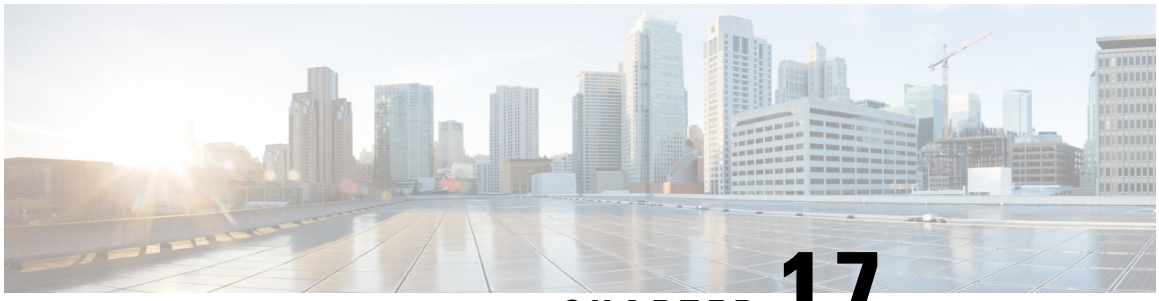
Feature	Relevant Troubleshooting Information
Firepower Threat Defense syslog	About Configuring Syslog, on page 962
Intrusion performance statistics	Intrusion Performance Statistic Logging Configuration, on page 1665
7000 and 8000 Series NGIPSv ASA with FirePOWER Services	generate-troubleshoot, on page 2488 (Command in the Command Line Interface (CLI))



PART IV

Deployment Management

- Domain Management, on page 335
- Policy Management, on page 343
- Rule Management: Common Characteristics, on page 361
- Reusable Objects, on page 393
- Firepower Threat Defense Certificate-Based Authentication, on page 481



CHAPTER 17

Domain Management

The following topics describe how to manage multitenancy using domains:

- Introduction to Multitenancy Using Domains, on page 335
- Requirements and Prerequisites for Domains, on page 338
- Managing Domains, on page 338
- Creating New Domains, on page 339
- Moving Data Between Domains, on page 340
- Moving Devices Between Domains, on page 341

Introduction to Multitenancy Using Domains

The Firepower System allows you to implement multitenancy using *domains*. Domains segment user access to managed devices, configurations, and events. You can create up to 50 subdomains under a top-level Global domain, in two or three levels.

When you log into the Firepower Management Center, you log into a single domain, called the *current domain*. Depending on your user account, you may be able to switch to other domains.

In addition to any restrictions imposed by your user role, your current domain level can also limit your ability to modify various Firepower System configurations. The system limits most management tasks, like system software updates, to the Global domain.

The system limits other tasks to *leaf domains*, which are domains with no subdomains. For example, you must associate each managed device with a leaf domain, and perform device management tasks from the context of that leaf domain.

Each leaf domain builds its own network map, based on the discovery data collected by that leaf domain's devices. Events reported by a managed device (connection, intrusion, malware, and so on) are also associated with the device's leaf domain.

One Domain Level: Global

If you do not configure multitenancy, all devices, configurations, and events belong to the Global domain, which in this scenario is also a leaf domain. Except for domain management, the system hides domain-specific configurations and analysis options until you add subdomains.

Two Domain Levels: Global and Second-Level

In a two-level multidomain deployment, the Global domain has direct descendant domains only. For example, a managed security service provider (MSSP) can use a single Firepower Management Center to manage network security for multiple customers:

- Administrators at the MSSP logging into the Global domain, cannot view or edit customers' deployments. They must log into respective second-level named subdomains to manage the customers' deployment.
- Administrators for each customer can log into second-level named subdomains to manage only the devices, configurations, and events applicable to their organizations. These local administrators cannot view or affect the deployments of other customers of the MSSP.

Three Domain Levels: Global, Second-Level, and Third-Level

In a three-level multidomain deployment, the Global domain has subdomains, at least one of which has its own subdomain. To extend the previous example, consider a scenario where an MSSP customer—already restricted to a subdomain—wants to further segment its deployment. This customer wants to separately manage two classes of device: devices placed on network edges and devices placed internally:

- Administrators for the customer logging into the second-level subdomain cannot view or edit the customer's edge network deployments. They must log into the respective leaf domain to manage the devices deployed on the network edge.
- Administrators for the customer's edge network can log into a third-level (leaf) domain to manage only the devices, configurations, and events applicable to devices deployed on the network edge. Similarly, administrators for the customer's internal network can log into a different third-level domain to manage internal devices, configurations, and events. Edge and internal administrators cannot view each other's deployment.



Note In an FMC that uses multi-tenancy, the SSO configuration can be applied only at the global domain level, and applies to the global domain and all subdomains.

Domains Terminology

This documentation uses the following terms when describing domains and multidomain deployments:

Global Domain

In a multidomain deployment, the top-level domain. If you do not configure multitenancy, all devices, configurations, and events belong to the Global domain. Administrators in the Global domain can manage the entire Firepower System deployment.

Subdomain

A second or third-level domain.

Second-level domain

A child of the Global domain. Second-level domains can be leaf domains, or they can have subdomains.

Third-level domain

A child of a second-level domain. Third-level domains are always leaf domains.

Leaf domain

A domain with no subdomains. Each device must belong to a leaf domain.

Descendant domain

A domain descending from the current domain in the hierarchy.

Child domain

A domain's direct descendant.

Ancestor domain

A domain from which the current domain descends.

Parent domain

A domain's direct ancestor.

Sibling domain

A domain with the same parent.

Current domain

The domain you are logged into now. The system displays the name of the current domain before your user name at the top right of the web interface. Unless your user role is restricted, you can edit configurations in the current domain.

Domain Properties

To modify a domain's properties, you must have Administrator access in that domain's parent domain.

Name and Description

Each domain must have a unique name within its hierarchy. A description is optional.

Parent Domain

Second- and third-level domains have a parent domain. You cannot change a domain's parent after you create the domain.

Devices

Only leaf domains may contain devices. In other words, a domain may contain subdomains or devices, but not both. You cannot save a deployment where a non-leaf domain directly controls a device.

In the domain editor, the web interface displays available and selected devices according to their current place in your domain hierarchy.

Host Limit

The number of hosts a FMC can monitor, and therefore store in network maps, depends on its model. In a multidomain deployment, leaf domains share the available pool of monitored hosts, but have separate network maps.

To ensure that each leaf domain can populate its network map, you can set host limits at each subdomain level. If you set a domain's host limit to `0`, the domain shares in the general pool.

Setting the host limit has a different effect at each domain level:

- Leaf — For a leaf domain, a host limit is a simple limit on the number of hosts the leaf domain can monitor.
- Second Level — For a second-level domain that manages third-level leaf domains, a host limit represents the total number of hosts that the leaf domains can monitor. The leaf domains share the pool of available hosts.
- Global — For the Global domain, the host limit is equal to the total number of hosts a FMC can monitor. You cannot change it

The sum of subdomains' host limits can add up to more than their parent domain's host limit. For example, if the Global domain host limit is 150,000, you can configure multiple subdomains each with a host limit of 100,000. Any of those domains, but not all, can monitor 100,000 hosts.

The network discovery policy controls what happens when you detect a new host after you reach the host limit; you can drop the new host, or replace the host that has been inactive for the longest time. Because each leaf domain has its own network discovery policy, each leaf domain governs its own behavior when the system discovers a new host.

If you reduce the host limit for a domain and its network map contains more hosts than the new limit, the system deletes the hosts that have been inactive the longest.

Related Topics

Firepower System Host Limit, on page 1832
Network Discovery Data Storage Settings, on page 1934

Requirements and Prerequisites for Domains

Model Support

Any.

Supported Domains

Any

User Roles

- Admin

Managing Domains

To modify a domain's properties, you must have Administrator access in that domain's parent domain.

Procedure

Step 1 Choose **System** > **Domains**.

Step 2 Manage your domains:

- Add — Click **Add Domain**, or click **Add Subdomain** next to the parent domain; see Creating New Domains, on page 339.
- Edit — Click **Edit** ( ) next to the domain you want to modify; see Domain Properties, on page 337.
- Delete — Click **Delete** ( ) next to the empty domain you want to delete, then confirm your choice. Move devices from domains you want to delete by editing their destination domain.

Step 3 When you are done making changes to the domain structure and all devices are associated with leaf domains, click **Save** to implement your changes.

Step 4 If prompted, make additional changes:

- If you changed a leaf domain to a parent domain, move or delete the old network map; see Moving Data Between Domains, on page 340.
- If you moved devices between domains and must assign new policies and security zones or interface groups, see Moving Devices Between Domains, on page 341.

What to do next

- Configure user roles and policies (access control, network discovery, and so on) for any new domains. Update device properties as needed.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Creating New Domains

You can create up to 50 subdomains under a top-level Global domain, in two or three levels.

You must assign all devices to a leaf domain before you can implement the domain configuration. When you add a subdomain to a leaf domain, the domain stops being a leaf domain and you must reassign its devices.

Procedure

Step 1 In a Global or a second-level domain, choose **System** > **Domains**.

Step 2 Click **Add Domain**, or click **Add Subdomain** next to the parent domain.

Step 3 Enter a **Name** and **Description**.

Step 4 Choose a **Parent Domain**.

Step 5 On **Devices**, choose the **Available Devices** to add to the domain, then click **Add to Domain** or drag and drop into the list of **Selected Devices**.

Step 6 Optionally, click **Advanced** to limit the number of hosts the new domain may monitor; see Domain Properties, on page 337.

Step 7 Click **Save** to return to the domain management page.

The system warns you if any devices are assigned to non-leaf domains. Click **Create New Domain** to create a new domain for those devices. Click **Keep Unassigned** if you plan to move the devices to existing domains.

Step 8 When you are done making changes to the domain structure and all devices are associated with leaf domains, click **Save** to implement your changes.

Step 9 If prompted, make additional changes:

- If you changed a leaf domain to a parent domain, move or delete the old network map; see Moving Data Between Domains, on page 340.
- If you moved devices between domains and must assign new policies and security zones or interface groups, see Moving Devices Between Domains, on page 341.

What to do next

- Configure user roles and policies (access control, network discovery, and so on) for any new domains. Update device properties as needed.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Moving Data Between Domains

Because events and network maps are associated with leaf domains, when you change a leaf domain to a parent domain, you have two choices:

- Move the network map and associated events to a new leaf domain.
- Delete the network map but retain the events. In this case, the events remain associated with the parent domain until the system prunes events as needed or as configured. Or, you can delete old events manually.

Before you begin

Implement a domain configuration where a former leaf domain is now a parent domain; see Managing Domains, on page 338.

Procedure

Step 1 For each former leaf domain that is now a parent domain:

- Choose a new **Leaf Domain** to inherit the **Parent Domain**'s events and network map.
- Choose **None** to delete the parent domain's network map, but retain old events.

Step 2 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Moving Devices Between Domains

You can move devices between domains when you are in the global domain or a second-level domain. Moving a device between domains can affect the configurations and policies applied to the device. The system automatically retains and updates what it can. It deletes what it cannot update, namely, object overrides, dynamic routing configuration, static routes, IP pool associated with the diagnostic interface,and DDNS.

If you assign a Remote Access VPN policy to a device, you cannot move the same device from one domain to another domain.

Specifically:

- If the health policy applied to a moved device is inaccessible in the new domain, you can choose a new health policy.
- If the access control policy assigned to a moved device is not valid or accessible in the new domain, choose a new policy. Every device must have an assigned access control policy.
- If the interfaces on the moved device belong to a security zone that is inaccessible in the new domain, you can choose a new zone.
- Interfaces are removed from:
  - Security zones that are inaccessible in the new domain and not used in an access control policy.
  - All interface groups.

If devices require a policy update but you do not need to move interfaces between zones, the system displays a message stating that zone configurations are up to date. For example, if a device's interfaces belong to a security zone configured in a common ancestor domain, you do not need to update zone configurations when you move devices from subdomain to subdomain.

Before you begin

- Implement a domain configuration where you moved a device from domain to domain and now must assign new policies and security zones; see Managing Domains, on page 338.

Procedure

Step 1 In the **Move Devices** dialog box, under **Select Device(s) to Configure**, check the device you want to configure.

Check multiple devices to assign the same health and access control policies.

Step 2 Choose an **Access Control Policy** to apply to the device, or choose **New Policy** to create a new policy.

Step 3 Choose a **Health Policy** to apply to the device, or choose **None** to leave the device without a health policy.

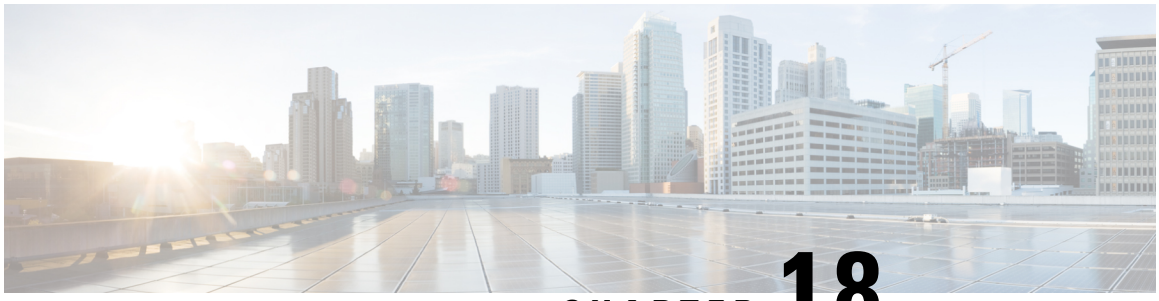
Step 4 If prompted to assign interfaces to new zones, choose a **New Security Zone** for each listed interface, or choose **None** to assign it later.

Step 5 After you configure all affected devices, click **Save** to save policy and zone assignments.

Step 6 Click **Save** to implement the domain configuration.

What to do next

- Update other configurations on the moved device that were affected by the move.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.



CHAPTER 18

Policy Management

The following topics describe how to manage various policies on the Firepower Management Center:

- Requirements and Prerequisites for Policy Management, on page 343
- Policy Deployment, on page 344
- Policy Comparison, on page 355
- Policy Reports, on page 356
- Out-of-Date Policies, on page 357
- Performance Considerations for Limited Deployments, on page 358

Requirements and Prerequisites for Policy Management

Model Support

Any.

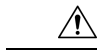
Supported Domains

Any

User Roles

- Admin
- Network Admin
- Security Approver

Policy Deployment



Caution Do NOT push the FMC deployments over a VPN tunnel that is terminating directly on the Firepower Threat Defense. Pushing the FMC deployments can potentially inactivate the tunnel and disconnect the FMC and the Firepower Threat Defense.

Recovering the device from this situation can be very disruptive and require executing the disaster recovery procedure. This procedure resets the Firepower Threat Defense configuration to factory defaults by changing manager from FMC to local and configuring the device from beginning. For more information, see Deploying the FMC Policy Configuration over VPN Tunnel, on page 344.

After you configure your deployment, and any time you change that configuration, you must deploy the changes to affected devices. You can view deployment status in the Message Center.

Deploying updates the following components:

- Device and interface configurations
- Device-related policies: NAT, VPN, QoS, platform settings
- Access control and related policies: DNS, file, identity, intrusion, network analysis, prefilter, SSL
- Network discovery policy
- Intrusion rule updates
- Configurations and objects associated with any of these elements

You can configure the system to deploy automatically by scheduling a deploy task or by setting the system to deploy when importing intrusion rule updates. Automating policy deployment is especially useful if you allow intrusion rule updates to modify system-provided base policies for intrusion and network analysis. Intrusion rule updates can also modify default values for the advanced preprocessing and performance options in your access control policies.

In a multidomain deployment, you can deploy changes for any domain where your user account belongs:

- Switch to an ancestor domain to deploy changes to all subdomains at the same time.
- Switch to a leaf domain to deploy changes to only that domain.

Best Practices for Deploying Configuration Changes

The following are guidelines for deploying configuration changes.

Deploying the FMC Policy Configuration over VPN Tunnel

You can deploy the FMC policy configuration over a VPN tunnel, only if the deployment is for a device that does not terminate the tunnel. The FMC to Firepower Threat Defense management traffic should be its own secure transport SF tunnel and does not need to be over S2S VPN tunnel for any connectivity.

For policy-based VPN tunnel, choose the protected networks on both side to exclude the FMC to Firepower Threat Defense management traffic. For route-based VPN tunnel, configure the routing to exclude the FMC to Firepower Threat Defense management traffic to the VTI interface.

When you push the FMC deployments over the VPN tunnel with the management traffic that is also passing through the tunnel, in the event of any VPN misconfiguration, it inactivates the tunnel and results in disconnecting the FMC and the Firepower Threat Defense.

To reinstantiate the tunnel configuration, you can either:

- Remove the sensor from the Firepower Threat Defense and the FMC (resulting in losing all of its configuration), and then add the sensor again to the FMC.

  Or

- Contact Cisco TAC.

Note Reinstantiating the tunnel configuration requires overhauling of the system.

Inline vs Passive Deployments

Do not apply inline configurations to devices deployed passively, and vice versa.

Time to Deploy and Memory Limitations

The time it takes to deploy depends on multiple factors, including (but not limited to):

- The configurations you send to the device. For example, if you dramatically increase the number of Security Intelligence entries you block, deploy can take longer.
- Device model and memory. On lower-memory devices, deploying can take longer. For example, it can take up to five minutes to deploy to a Firepower 7010, 7020, or 7030 device.

Do not exceed the capability of your devices. If you exceed the maximum number of rules or policies supported by a target device, the system displays a warning. The maximum depends on a number of factors—not only memory and the number of processors on the device, but also on policy and rule complexity. For information on optimizing policies and rules, see Best Practices for Access Control Rules, on page 1194.

Interruptions to Traffic Flow and Inspection During Deploy

When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 and Configurations that Restart the Snort Process When Deployed or Activated, on page 352.

Caution We *strongly* recommend you deploy in a maintenance window or at a time when interruptions will have the least impact.

Auto-Enabling Application Detectors

If you are performing application control but disable required detectors, the system automatically enables the appropriate system-provided detectors upon policy deploy. If none exist, the system enables the most recently modified user-defined detector for the application.

Asset Rediscovery with Network Discovery Policy Changes

When you deploy changes to a network discovery policy, the system deletes and then rediscovers MAC address, TTL, and hops information from the network map for the hosts in your monitored networks. Also, the affected managed devices discard any discovery data that has not yet been sent to the FMC.

Related Topics

Snort® Restart Scenarios, on page 348

Deploy Configuration Changes



Caution Do NOT push the FMC deployments over a VPN tunnel that is terminating directly on the Firepower Threat Defense. Pushing the FMC deployments can potentially inactivate the tunnel and disconnect the FMC and the Firepower Threat Defense.

Recovering the device from this situation can be very disruptive and require executing the disaster recovery procedure. This procedure resets the Firepower Threat Defense configuration to factory defaults by changing manager from FMC to local and configuring the device from beginning. For more information, see Deploying the FMC Policy Configuration over VPN Tunnel, on page 344.

After you change configurations, deploy them to the affected devices. We *strongly* recommend that you deploy in a maintenance window or at a time when any interruptions to traffic flow and inspection will have the least impact.



Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 and Configurations that Restart the Snort Process When Deployed or Activated, on page 352.

Before you begin

- Review the guidelines described in Best Practices for Deploying Configuration Changes, on page 344.
- Be sure all managed devices use the same revision of the Security Zones object. If you have edited security zone objects: Do not deploy configuration changes to any device until you edit the zone setting for interfaces on *all* devices you want to sync. You must deploy to all managed devices at the same time. See Synchronizing Security Zone Object Revisions, on page 501.

Note Policy deployment process fails if the sensor configuration is being read by the system during deployment. Executing commands such as `show running-config` from the sensor CLI disturbs the deployment, which results in deployment failure.

Procedure

Step 1 On the FMC menu bar, click **Deploy**.

The Deploy Policies dialog lists devices with out-of-date configurations. The **Version** at the top of the dialog specifies when you last made configuration changes. The **Current Version** column in the device table specifies when you last deployed changes to each device.

Step 2 Identify and choose the devices where you want to deploy configuration changes.

- Sort—Sort the device list by clicking a column heading.
- Expand—Click **Plus** to expand a device listing to view the configuration changes to be deployed. The system marks out-of-date policies with an **Index**.
- Filter—Filter the device list. Click the arrow in the upper-right corner of any column heading in the display, enter text in the **Filters** text box, and press Enter. Check or uncheck the check box to activate or deactivate the filter.
- Arrange—Place the mouse on a column heading to drag and drop the column in your preferred order.

Step 3 (Optional) Enable **Check for rule conflicts** to check for rule conflicts before deployment begins.

In policies that evaluate rules in a specific order, conflicts occur when a rule will never match traffic because an earlier rule always matches the traffic first. Although rule conflicts do not prevent deployment, properly creating and ordering rules is essential to building an effective deployment. The deployment process can check for rule conflicts in access control policies. However, in complex deployments, checking for rule conflicts can take some time.

Step 4 Click **Deploy**.

Step 5 If the system identifies errors or warnings in the changes to be deployed, it displays them in the **Errors and Warnings for Requested Deployment** window.

You have the following choices:

- Proceed—Continue deploying without resolving warning conditions. If you enabled rule conflict checks, warnings include detected rule conflicts. You cannot proceed if the system identifies errors.
- Cancel—Exit without deploying. Resolve the error and warning conditions, and attempt to deploy the configuration again.

What to do next

- (Optional) Monitor deployment status; see Viewing Deployment Messages, on page 315.
- If deploy fails, see Best Practices for Deploying Configuration Changes, on page 344.

Related Topics

Snort® Restart Scenarios, on page 348

Redeploy Existing Configurations to a Device

You can force-deploy existing (unchanged) configurations to a single managed device. We *strongly* recommend you deploy in a maintenance window or at a time when any interruptions to traffic flow and inspection will have the least impact.

Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 and Configurations that Restart the Snort Process When Deployed or Activated, on page 352.

Before you begin

Review the guidelines described in Best Practices for Deploying Configuration Changes, on page 344.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **Edit** ( ) next to the device where you want to force deployment.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Device**.

Step 4 Click **Edit** ( ) next to the **General** section heading.

Step 5 Click **Force Deploy** ( ).

Step 6 Click **Deploy**.

The system identifies any errors or warnings with the configurations you are deploying. You can click **Proceed** to continue without resolving warning conditions. However, you cannot proceed if the system identifies an error.

Related Topics

Snort® Restart Scenarios, on page 348

Snort® Restart Scenarios

When the traffic inspection engine referred to as *the Snort process* on a managed device restarts, inspection is interrupted until the process resumes. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350

for more information. Additionally, resource demands may result in a small number of packets dropping without inspection when you deploy, regardless of whether the Snort process restarts.

Any of the scenarios in the following table cause the Snort process to restart.

Table 50: Snort Restart Scenarios

Restart Scenario	More Information
Deploying a specific configuration that requires the Snort process to restart.	Configurations that Restart the Snort Process When Deployed or Activated, on page 352
Modifying a configuration that immediately restarts the Snort process.	Changes that Immediately Restart the Snort Process, on page 354
Traffic-activation of the currently deployed Automatic Application Bypass (AAB) configuration.	Configure Automatic Application Bypass, on page 241

Related Topics

Access Control Policy Advanced Settings, on page 1211

Configurations that Restart the Snort Process When Deployed or Activated, on page 352

Inspect Traffic During Policy Apply

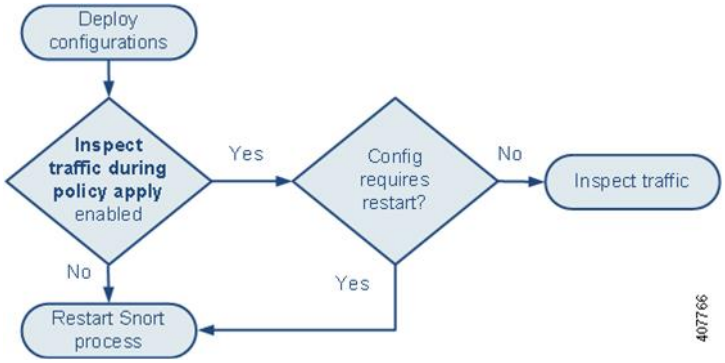
Inspect traffic during policy apply is an advanced access control policy general setting that allows managed devices to inspect traffic while deploying configuration changes; this is the case unless a configuration that you deploy requires the Snort process to restart. You can configure this option as follows:

- Enabled — Traffic is inspected during the deployment unless certain configurations require the Snort process to restart.

  When the configurations you deploy do not require a Snort restart, the system initially uses the currently deployed access control policy to inspect traffic, and switches during deployment to the access control policy you are deploying.

- Disabled — Traffic is not inspected during the deployment. The Snort process always restarts when you deploy.

The following graphic illustrates how Snort restarts can occur when you enable or disable **Inspect traffic during policy apply**.



Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 and Configurations that Restart the Snort Process When Deployed or Activated, on page 352.

Snort® Restart Traffic Behavior

The following tables explain how different devices handle traffic when the Snort process restarts.

Table 51: FTD and FTDv Restart Traffic Effects

Interface Configuration	Restart Traffic Behavior
inline: **Snort Fail Open: Down**: disabled	dropped
inline: **Snort Fail Open: Down**: enabled	passed without inspection Some packets can be delayed in buffer for several seconds before the system recognizes that Snort is down. This delay can vary depending upon the load distribution. However, the buffered packets are eventually passed.
routed, transparent (including EtherChannel, redundant, subinterface) when **preserve-connection** is enabled (**configure snort preserve-connection enable**; default) Note that both the FMC and Firepower Threat Defense must be running Version 6.2.0.2 or a subsequent 6.2.0.x patch. For more information, see Cisco Firepower Threat Defense Command Reference.	existing TCP/UDP flows: passed without inspection so long as at least one packet arrives while Snort is down new TCP/UDP flows and all non-TCP/UDP flows: dropped Note that the following traffic drops even when **preserve-connection** is enabled: • When Snort goes down before the system tags traffic that would have matched a **Do not decrypt** SSL rule or default policy action (tagging occurs after a few handshake packet exchanges) • plaintext, passthrough prefilter tunnel traffic that matches an **Analyze** rule action or an **Analyze all tunnel traffic** default policy action • connections that do not match an access control rule and are instead handled by the default action. • decrypted TLS/SSL traffic • a safe search flow • a captive portal flow

Interface Configuration	Restart Traffic Behavior
routed, transparent (including EtherChannel, redundant, subinterface) when either of the following occurs: • **preserve-connection** is disabled (**configure snort preserve-connection disable**) • the FMC or Firepower Threat Defense is not running version 6.2.0.2 or a subsequent 6.2.0.x patch	dropped
inline: tap mode	egress packet immediately, copy bypasses Snort
passive	uninterrupted, not inspected

Table 52: 7000 and 8000, NGIPSv Restart Traffic Effects

Interface Configuration	Restart Traffic Behavior
inline: **Failsafe** enabled or disabled	passed without inspection A few packets might drop if **Failsafe** is disabled and Snort is busy but not down.
inline: tap mode	egress packet immediately, copy bypasses Snort
passive	uninterrupted, not inspected
routed, switched (7000 and 8000 Series only)	dropped

Table 53: ASA FirePOWER Restart Traffic Effects

Interface Configuration	Restart Traffic Behavior
routed or transparent with fail-open	passed without inspection
routed or transparent with fail-close	dropped

Note In addition to traffic handling when the Snort process is down while it restarts, traffic can also pass without inspection or drop when the Snort process is busy, depending on the configuration of the Failsafe option (see Inline Sets, on page 509) or the Snort Fail Open **Busy** option (see Configure an Inline Set, on page 659). A device supports either the Failsafe option or the Snort Fail Open option, but not both.

Warning Do not reboot the system while the Snort Rule Update is in progress.

Snort-busy drops happen when snort is not able to process the packets fast enough. Lina does not know whether Snort is busy due to processing delay, or if is stuck or due to call blocking. When transmission queue is full, snort-busy drops occur. Based on Transmission queue utilization, Lina will try to access if the queue is being serviced smoothly.



Note When the Snort process is busy but not down during configuration deployment, some packets may drop on routed, switched, or transparent interfaces if the total CPU load exceeds 50 percent.

Configurations that Restart the Snort Process When Deployed or Activated

Deploying any of the following configurations except AAB restarts the Snort process as described. Deploying AAB does not cause a restart, but excessive packet latency activates the currently deployed AAB configuration, causing a partial restart of the Snort process.

Access Control Policy

- Add the first or remove the last URL category/reputation condition in an access control rule.
- Change the total number of active intrusion policies by adding an intrusion policy that is not currently used, or by removing the last instance of an intrusion policy. You can use an intrusion policy in an access control rule, as the default action, or as the default intrusion policy.

Access Control Policy Advanced Settings

- Deploy when **Inspect Traffic During Policy Apply** is disabled.
- Configure a non-default value under Files and Malware Settings.
- Add or remove an SSL policy.
- Enable or disable adaptive profiles.
- Enable or disable the **Log Session/Protocol Distribution** troubleshooting option.

Security Intelligence

Add or delete multiple Security Intelligence whitelist or blacklist networks or network objects. Changes can be to custom or system-provided lists, and whether the Snort process restarts can vary by device, depending on the memory available for inspection.

SSL Policy

- Add the first or remove the last category/reputation condition in an SSL rule.

File Policy

Deploy the first or last of any one of the following configurations; note that while otherwise deploying these file policy configurations does not cause a restart, deploying non-file-policy configurations can cause restarts.

- Enable or disable **Inspect Archives**.
- Take either of the following actions:

- Enable or disable **Inspect Archives** when the deployed access control policy includes at least one file policy.
- Add the first or remove the last file policy rule when **Inspect Archives** is enabled (note that at least one rule is required for **Inspect Archives** to be meaningful).

- Select **Detect Files** or **Block Files** in a file rule.
- Enable or disable **Store files** in a **Detect Files** or **Block Files** rule.
- Add the first or remove the last active file rule that combines the **Malware Cloud Lookup** or **Block Malware** rule action with an analysis option (**Spero Analysis or MSEXE**, **Dynamic Analysis**, or **Local Malware Analysis**) or a store files option (**Malware, Unknown**, **Clean**, or **Custom**).

Note that access control rules that deploy these file policy configurations to security zones or tunnel zones cause a restart only when your configuration meets the following conditions:

- Source or destination security zones in your access control rule must match the security zones associated with interfaces on the target devices.
- Unless the destination zone in you access control rule is *any*, a source tunnel zone in the rule must match a tunnel zone assigned to a tunnel rule in the prefilter policy.

Identity Policy

- When SSL decryption is disabled (that is, when the access control policy does not include an SSL policy), add the first or remove the last active authentication rule.

  An active authentication rule has either an **Active Authentication** rule action, or a **Passive Authentication** rule action with **Use active authentication if passive authentication cannot identify user** selected.

Network Analysis Policy

- Change the total number of network analysis policies by adding a network analysis policy that is not currently used, or by removing the last instance of a network analysis policy. You can use a network analysis policy with network analysis rules or as the default network analysis policy.
- Change the value for the IMAP, POP, or SMTP preprocessor **Base64 Decoding Depth**, **7-Bit/8-Bit/Binary Decoding Depth**, **Quoted-Printable Decoding Depth**, or **Unix-to-Unix Decoding Depth**.

Network Discovery

- Enable or disable non-authoritative, traffic-based user detection over the HTTP, FTP, or MDNS protocols, using the network discovery policy.

Device Management

- Routing: Add a routed interface pair or virtual router to a 7000 or 8000 Series device.
- VPN: Add or remove a VPN on a 7000 or 8000 Series device.

Caution The system does not warn you that the Snort process restarts when you add or remove a VPN on a 7000 or 8000 Series device.

- MTU: Change the highest MTU value among all non-management interfaces on a device.
- 7000/8000 series high availability: Change a high-availability state sharing option.
- Automatic Application Bypass (AAB): The currently deployed AAB configuration activates when a malfunction of the Snort process or a device misconfiguration causes a single packet to use an excessive amount of processing time. The result is a partial restart of the Snort process to alleviate extremely high latency or prevent a complete traffic stall. This partial restart causes a few packets to pass without inspection, or drop, depending on how the device handles traffic.

Updates

- System update: Deploy configurations the first time after a software update that includes a new version of the Snort binary or data acquisition library (DAQ).
- VDB: Deploy configurations the first time after installing a vulnerability database (VDB) .
- Intrusion rule update: Deploying configurations the first time after importing an intrusion rule update (also known as a *Snort Rule Update* or *SRU*).

Caution Intrusion rule updates are cumulative. Any shared object rule that is added or modified since your last update causes a restart when you deploy, even if the current update has no shared object rule changes.

Related Topics

Deploy Configuration Changes, on page 346
Snort® Restart Scenarios, on page 348

Changes that Immediately Restart the Snort Process

The following changes immediately restart the Snort process without going through the deploy process. How the restart affects traffic depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

- Take any of the following actions involving applications or application detectors:
  - Activate or deactivate a system or custom application detector.
  - Delete an activated custom detector.
  - **Save and Reactivate** an activated custom detector.
  - Create a user-defined application.

  The Snort process restarts on all managed devices.
- Create or break a Firepower Threat Defense high availability pair—Restarts the Snort process on the primary and secondary devices.

- Install a vulnerability database (VDB) update.
- Restart the Snort process in the 7000 or 8000 Series user interface (**System > Configuration > Process**)—The system prompts you for confirmation and allows you to cancel.

Policy Comparison

To review policy changes for compliance with your organization's standards or to optimize system performance, you can examine the differences between two policies or between a saved policy and the running configuration.

You can compare the following policy types:

- DNS
- File
- Health
- Identity
- Intrusion (Only Snort 2 policies)
- Network Analysis
- SSL

The comparison view displays both policies in a side-by-side format. Differences between the two policies are highlighted:

- Blue indicates that the highlighted setting is different in the two policies, and the difference is noted in red text.
- Green indicates that the highlighted setting appears in one policy but not the other.

Comparing Policies

You can compare policies only if you have access rights and any required licenses for the specific policy, and you are in the correct domain for configuring the policy.

Procedure

Step 1 Access the management page for the policy you want to compare:

- DNS—**Policies** > **Access Control** > **DNS**
- File—**Policies** > **Access Control** > **Malware & File**
- Health—**System** > **Health** > **Policy**
- Identity—**Policies** > **Access Control** > **Identity**
- Intrusion—**Policies** > **Access Control** > **Intrusion**

Note You can compare only Snort 2 policies.

- Network Analysis—**Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**

  Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

- SSL—**Policies** > **Access Control** > **SSL**

Step 2 Click **Compare Policies**.

Step 3 From the **Compare Against** drop-down list, choose the type of comparison you want to make:

- To compare two different policies, choose **Other Policy**.
- To compare two revisions of the same policy, choose **Other Revision**.
- To compare another policy to the currently active policy, choose **Running Configuration**.

Step 4 Depending on the comparison type you choose, you have the following choices:

- If you are comparing two different policies, choose the policies you want to compare from the **Policy A** and **Policy B** drop-down lists.
- If you are comparing the running configuration to another policy, choose the second policy from the **Policy B** drop-down list.

Step 5 Click **OK**.

Step 6 Review the comparison results:

- Comparison Viewer—To use the comparison viewer to navigate individually through policy differences, click **Previous** or **Next** above the title bar.
- Comparison Report—To generate a PDF report that lists the differences between the two policies, click **Comparison Report**.

Policy Reports

For most policies, you can generate two kinds of reports. A report on a single policy provides details on the policy's current saved configuration, while a comparison report lists only the differences between two policies. You can generate a single-policy report for all policy types except health.



Note Intrusion policy reports combine the settings in the base policy with the settings of the policy layers, and make no distinction between which settings originated in the base policy or policy layer.

Generating Current Policy Reports

You can generate policy reports only if you have access rights and any required licenses for the specific policy, and you are in the correct domain for configuring the policy.

Procedure

Step 1 Access the management page for the policy for which you want to generate a report:

- Access Control—**Policies** > **Access Control**
- DNS—**Policies** > **Access Control** > **DNS**
- File—**Policies** > **Access Control** > **Malware & File**
- Health—**System** > **Health** > **Policy**
- Identity—**Policies** > **Access Control** > **Identity**
- Intrusion—**Policies** > **Access Control** > **Intrusion**
- NAT for 7000 & 8000 Series devices—**Devices** > **NAT**
- Network Analysis—**Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**

  Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

- SSL—**Policies** > **Access Control** > **SSL**

Step 2 Click **Report** () next to the policy for which you want to generate a report.

Out-of-Date Policies

The Firepower System marks out-of-date policies with red status text that indicates how many of its targeted devices need a policy update. To clear this status, you must re-deploy the policy to the devices.

Configuration changes that require a policy re-deploy include:

- Modifying an access control policy: any changes to access control rules, the default action, policy targets, Security Intelligence filtering, advanced options including preprocessing, and so on.
- Modifying any of the policies that the access control policy invokes: the SSL policy, network analysis policies, intrusion policies, file policies, identity policies, or DNS policies.
- Changing any reusable object or configuration used in an access control policy or policies it invokes:
  - network, port, VLAN tag, URL, and geolocation objects
  - Security Intelligence lists and feeds
  - application filters or detectors
  - intrusion policy variable sets
  - file lists
  - decryption-related objects and security zones
- Updating the system software, intrusion rules, or the vulnerability database (VDB).

Keep in mind that you can change some of these configurations from multiple places in the web interface. For example, you can modify security zones using the object manager (**Objects** > **Object Management**), but modifying an interface type in a device's configuration (**Devices** > **Device Management**) can also change a zone and require a policy re-deploy.

Note that the following updates do **not** require policy re-deploy:

- automatic updates to Security Intelligence feeds and additions to the Security Intelligence global Block or Do Not Block list using the context menu
- automatic updates to URL filtering data
- scheduled geolocation database (GeoDB) updates

Performance Considerations for Limited Deployments

Host, application, and user discovery data allow the system to create a complete, up-to-the-minute profile of your network. The system can also act as an intrusion detection and prevention system (IPS), analyzing network traffic for intrusions and exploits and, optionally, dropping offending packets.

Combining discovery and IPS gives context to your network activity and allows you to take advantage of many features, including:

- impact flags and indications of compromise, which can tell you which of your hosts are vulnerable to a particular exploit, attack, or piece of malware
- adaptive profile updates and Firepower recommendations, which allow you to examine traffic differently depending on the destination host
- correlation, which allows you to respond to intrusions (and other events) differently depending on the affected host

However, if your organization is interested in performing only IPS, or only discovery, there are a few configurations that can optimize the performance of the system.

Discovery Without Intrusion Prevention

The *discovery* feature allows you to monitor network traffic and determine the number and types of hosts (including network devices) on your network, as well as the operating systems, active applications, and open ports on those hosts. You can also configure managed devices to monitor user activity on your network. You can use discovery data to perform traffic profiling, assess network compliance, and respond to policy violations.

In a basic deployment (discovery and simple, network-based access control only), you can improve a device's performance by following a few important guidelines when configuring its access control policy.



Note You must use an access control policy, even if it simply allows all traffic. The network discovery policy can **only** examine traffic that the access control policy allows to pass.

First, make sure your access control policy does not require complex processing and uses only simple, network-based criteria to handle network traffic. You must implement **all** of the following guidelines; misconfiguring any one of these options eliminates the performance benefit:

- Do **not** use the Security Intelligence feature. Remove any populated global Block or Do Not Block list from the policy's Security Intelligence configuration.
- Do **not** include access control rules with Monitor or Interactive Block actions. Use only Allow, Trust, and Block rules. Keep in mind that allowed traffic can be inspected by discovery; trusted and blocked traffic cannot.
- Do **not** include access control rules with application, user, URL, ISE attribute, or geolocation-based network conditions. Use only simple network-based conditions: zone, IP address, VLAN tag, and port.
- Do **not** include access control rules that perform file, malware, or intrusion inspection. In other words, do not associate a file policy or intrusion policy with any access control rule.
- In the Advanced settings for the access control policy, make sure that **Intrusion Policy used before Access Control rule is determined** is set to **No Rules Active**.
- Select **Network Discovery Only** as the policy's default action. Do **not** choose a default action for the policy that performs intrusion inspection.

In conjunction with the access control policy, you can configure and deploy the network discovery policy, which specifies the network segments, ports, and zones that the system examines for discovery data, as well as whether hosts, applications, and users are discovered on the segments, ports, and zones.

Related Topics

Inspection of Packets That Pass Before Traffic Is Identified, on page 1670

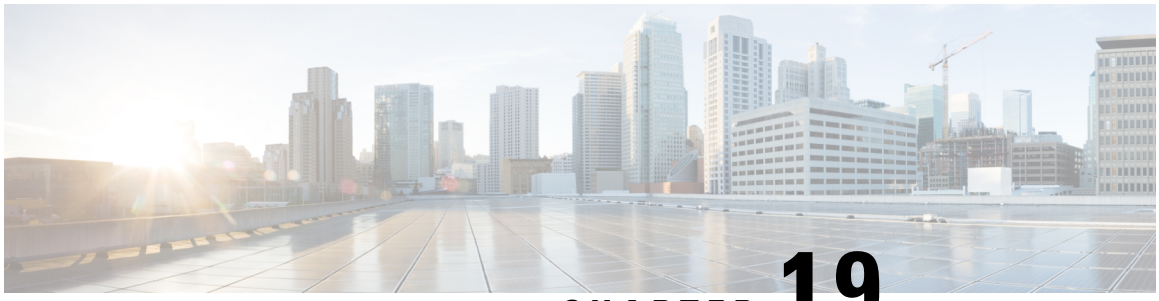
Intrusion Prevention Without Discovery

Disabling discovery if you don't need it (for example, in an IPS-only deployment) can improve performance. To disable discovery you must implement *all* of these changes:

- Delete *all* rules from your network discovery policy.
- Use *only* simple network-based conditions to perform access control: zone, IP address, VLAN tag, and port.

  Do *not* perform any kind of, application, user, URL, or geolocation control. Although you can disable storage of discovery data, the system still must collect and examine it to implement those features.

After you deploy, new discovery halts on target devices. The system gradually deletes information in the network map according to your timeout preferences. Or, you can purge all discovery data immediately.



CHAPTER 19

Rule Management: Common Characteristics

The following topics describe how to manage common characteristics of rules in various policies on the Firepower Management Center:

- Requirements and Prerequisites for Rule Management, on page 361
- Introduction to Rules, on page 361
- Rule Condition Types, on page 363
- Searching for Rules, on page 390
- Filtering Rules by Device, on page 391
- Rule and Other Policy Warnings, on page 392

Requirements and Prerequisites for Rule Management

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Introduction to Rules

Rules in various policies exert granular control over network traffic. The system evaluates traffic against rules in the order that you specify, using a first-match algorithm.

Although these rules may include other configurations that are not consistent across policies, they share many basic characteristics and configuration mechanics, including:

- Conditions: Rule conditions specify the traffic that each rule handles. You can configure each rule with multiple conditions. Traffic must match all conditions to match the rule.
- Action: A rule's action determines how the system handles matching traffic. Note that even if a rule does not have an **Action** list you can choose from, the rule still has an associated action. For example, a custom network analysis rule uses a network analysis policy as its "action." As another example, QoS rules do not have an explicit action because all QoS rules do the same thing: rate limit traffic.
- Position: A rule's position determines its evaluation order. When using a policy to evaluate traffic, the system matches traffic to rules in the order you specify. Usually, the system handles traffic according to the first rule where all the rule's conditions match the traffic. (Monitor rules, which are designed to track and log, are an exception.) Proper rule order reduces the resources required to process network traffic, and prevents rule preemption.
- Category: To organize some rule types, you can create custom rule categories in each parent policy.
- Logging: For many rules, logging settings govern whether and how the system logs connections handled by the rule. Some rules (such as identity and network analysis rules) do not include logging settings because the rules neither determine the final disposition of connections, nor are they specifically designed to log connections. As another example, QoS rules do not include logging settings; you cannot log a connection simply because it was rate limited.
- Comments: For some rule types, each time you save changes, you can add comments. For example, you might summarize the overall configuration for the benefit of other users, or note when you change a rule and the reason for the change.



Tip A right-click menu in many policy editors provides shortcuts to many rule management options, including editing, deleting, moving, enabling, and disabling.

Rules with Shared Characteristics

This chapter documents many common aspects of the following rules and configurations. For information on non-shared configurations, see:

- Access control rules: Access Control Rules, on page 1215
- Prefilter rules: Tunnel and Prefilter Rule Components, on page 1281
- Tunnel rules: Tunnel and Prefilter Rule Components, on page 1281
- SSL rules: Creating and Modifying TLS/SSL Rules, on page 1354
- DNS rules: Creating and Editing DNS Rules, on page 1266
- Identity rules: Create an Identity Rule, on page 1962
- Network analysis rules: Configuring Network Analysis Rules, on page 1674
- QoS rules: Configuring QoS Rules, on page 677
- Intelligent Application Bypass (IAB): Intelligent Application Bypass, on page 1289
- Application filters: Application Filters, on page 406

Rules without Shared Characteristics

Rules whose configurations are not documented in this chapter include:

- Intrusion rules: Tuning Intrusion Policies Using Rules, on page 1489
- File and malware rules: File Rules, on page 1424
- Correlation rules: Configuring Correlation Rules, on page 1989
- NAT rules (Classic): NAT for 7000 and 8000 Series Devices, on page 987
- NAT rules (Firepower Threat Defense): Network Address Translation (NAT) for Firepower Threat Defense, on page 1005
- 8000 Series fastpath rules: Configure Fastpath Rules (8000 Series), on page 243

Rule Condition Types

The following table describes the common rule conditions documented in this chapter, and lists the configurations where they are used.

Condition	Controls Traffic By...	Supported Rules/Configurations
Interface Conditions, on page 366	Source and destination interfaces, and where supported, tunnel zones	Access control rules Tunnel rules Prefilter rules SSL rules DNS rules Identity rules Network analysis rules QoS rules
Network Conditions, on page 368	Source and destination IP address, and where supported, geographical location or originating client	Access control rules Prefilter rules SSL rules DNS rules Identity rules Network analysis rules QoS rules
Tunnel Endpoint Conditions, on page 370	Source and destination tunnel endpoints for plaintext, passthrough tunnels	Tunnel rules

Condition	Controls Traffic By...	Supported Rules/Configurations
VLAN Conditions, on page 372	VLAN tag	Access control rules **Note** For FTD, VLAN tags in access rules only apply to inline sets; they cannot be used in access rules applied to firewall interfaces. Tunnel rules Prefilter rules SSL rules DNS rules Identity rules Network analysis rules
Port and ICMP Code Conditions, on page 373	Source and destination ports, protocols, and ICMP codes	Access control rules Prefilter rules SSL rules Identity rules QoS rules
Encapsulation Conditions, on page 375	Encapsulation protocol (nonencrypted)	Tunnel rules
Application Conditions (Application Control), on page 375	Application or application characteristic (type, risk, business relevance, category, and tags)	Access control rules SSL rules Identity rules QoS rules Application filters Intelligent Application Bypass (IAB)
URL Conditions (URL Filtering), on page 384	URL, and where supported, URL characteristic (category and reputation)	Access control rules SSL rules QoS rules
User, Realm, and ISE Attribute Conditions (User Control), on page 384	Logged-in authoritative user of a host, or that user's realm, group, or ISE attributes	Access control rules SSL rules (no ISE attributes) QoS rules

Condition	Controls Traffic By...	Supported Rules/Configurations
Custom SGT Conditions, on page 389	Custom Security Group Tag (SGT)	Access control rules

Rule Condition Mechanics

Rule conditions specify the traffic that each rule handles. You can configure each rule with multiple conditions, and traffic must match all conditions to match the rule. The available condition types depend on the rule type.

In rule editors, each condition type has its own tab page. Build conditions by choosing the traffic characteristics you want to match. In general, choose criteria from one or two lists of available items on the left, then add or combine those criteria into one or two lists of selected items on the right. For example, in URL conditions in access control rules, you can combine URL category and reputation criteria to create a single group of websites to block.

To help you build conditions, you can match traffic using various system-provided and custom configurations, including realms, ISE attributes, and various types of objects and object groups. Often, you can manually specify rule criteria.

Leave matching criteria empty whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against every combination of the contents of the criteria you specify.



Caution Failure to set up your access control rules properly can have unexpected results, including traffic being allowed that should be blocked. In general, application control rules should be lower in your access control list because it takes longer for those rules to match than rules based on IP address, for example.

Access control rules that use *specific* conditions (such as networks and IP addresses) should be ordered *before* rules that use general conditions (such as applications). If you're familiar with the Open Systems Interconnect (OSI) model, use similar numbering in concept. Rules with conditions for layers 1, 2, and 3 (physical, data link, and network) should be ordered first in your access control rules. Conditions for layers 5, 6, and 7 (session, presentation, and application) should be ordered later in your access control rules. For more information about the OSI model, see this Wikipedia article.

Source and Destination Criteria

Where a rule involves source and destination criteria (zones, networks, ports), usually you can use either or both criteria as constraints. If you use both, matching traffic must originate from one of the specified source zones, networks, or ports and leave through one of the destination zones, networks, or ports.

Items per Condition

You can add up to 50 items to each condition. For rules with source and destination criteria, you can use up to 50 of each. Traffic that matches any of the selected items matches the condition.

Simple Rule Mechanics

In rule editors, you have the following general choices. For detailed instructions on building conditions, see the topics for each condition type.

- Choose Item—Click an item or check its check box. Often you can use Ctrl or Shift to choose multiple items, or right-click to **Select All**.
- Search—Enter criteria in the search field. The list updates as you type. The system searches item names and, for objects and object groups, their values. Click **Reload** ( ) or **Clear** ( ) to clear the search.
- Add Predefined Item—After you choose one or more available items, click an **Add** button or drag and drop. The system prevents you from adding invalid items: duplicates, invalid combinations, and so on.
- Add Manual Item—Click the field under the **Selected** items list, enter a valid value, and click **Add**. When you add ports, you may also choose a protocol from the drop-down list.
- Create Object—Click **Add** ( ) to create a new, reusable object that you can immediately use in the condition you are building, then manage in the object manager. When using this method to add application filters on the fly, you cannot save a filter that includes another user-created filter.
- Delete—Click the **Delete** ( ) for an item, or choose one or more items and right-click to **Delete Selected**.

Interface Conditions

Interface rule conditions control traffic by its source and destination interfaces.

Depending on the rule type and the devices in your deployment, you can use predefined *interface objects* called *security zones* or *interface groups* to build interface conditions. Interface objects segment your network to help you manage and classify traffic flow by grouping interfaces across multiple devices; see Interface Objects: Interface Groups and Security Zones, on page 410.

Tip Constraining rules by interface is one of the best ways to improve system performance. If a rule excludes all of a device's interfaces, that rule does not affect that device's performance.

Just as all interfaces in an interface object must be of the same type (all inline, passive, switched, routed, or ASA FirePOWER), all interface objects used in an interface condition must be of the same type. Because devices deployed passively do not transmit traffic, in passive deployments you cannot constrain rules by destination interface.

Tunnel Zones vs Security Zones

In some configurations, you can use tunnel zones instead of security zones to constrain interface conditions. Tunnel zones allow you to use prefiltering to tailor subsequent traffic handling to certain types of encapsulated connections.

Note If a configuration supports tunnel zone constraints, a rezoned connection—a connection with an assigned tunnel zone—does **not** match security zone constraints. For more information, see Tunnel Zones and Prefiltering, on page 1282.

Rules with Interface Conditions

Rule Type	Supports Security Zones?	Supports Tunnel Zones?	Supports Interface Groups?
Access control	yes	yes	no
Tunnel and prefilter	yes	n/a; you assign tunnel zones in the prefilter policy	yes
SSL	yes	no	no
DNS (source only)	yes	no	no
Identity	yes	no	no
Network analysis	yes	no	no
QoS (routed only, required)	yes	no	yes

Example: Access Control Using Security Zones

Consider a deployment where you want hosts to have unrestricted access to the internet, but you nevertheless want to protect them by inspecting incoming traffic for intrusions and malware.

First, create two security zones: Internal and External. Then, assign interface pairs on one or more devices to those zones, with one interface in each pair in the Internal zone and one in the External zone. Hosts connected to the network on the Internal side represent your protected assets.



Note You are not required to group all internal (or external) interfaces into a single zone. Choose the grouping that makes sense for your deployment and security policies.

Then, configure an access control rule with a destination zone condition set to Internal. This simple rule matches traffic that leaves the device from any interface in the Internal zone. To inspect matching traffic for intrusions and malware, choose a rule action of **Allow**, then associate the rule with an intrusion and a file policy.

Configuring Interface Conditions

Before you begin

- (Access control only) If you want to constrain traffic by tunnel zones instead of security zones, make sure the associated prefilter policy assigns those zones; see Associating Other Policies with Access Control, on page 1213.

Procedure

Step 1 In the rule editor, click the following for interface conditions:

- Interface groups and security zones (tunnel, prefilter, QoS)—Click **Interface Objects**.
- Security zones (access control, SSL, DNS, identity, network analysis)—Click **Zones**.
- Tunnel zones (access control)—Click **Zones**.

Step 2 Find and choose the interfaces you want to add from the **Available Interface Objects** or **Available Zones** list.

(Access control only) To match connections in rezoned tunnels, choose tunnel zones instead of security zones. You cannot use tunnel and security zones in the same rule. For more information, see Tunnel Zones and Prefiltering, on page 1282.

Step 3 Click **Add to Source** or **Add to Destination**, or drag and drop.

Step 4 Save or continue editing the rule.

What to do next

- (Access control only) If you rezoned tunnels during prefiltering, configure additional rules if necessary to ensure complete coverage. Connections in rezoned tunnels do **not** match rules with security zone constraints. For more information, see Using Tunnel Zones, on page 1283.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Network Conditions

Network rule conditions control traffic by its source and destination IP address, using inner headers. Tunnel rules, which use outer headers, have tunnel endpoint conditions instead of network conditions.

You can use predefined objects to build network conditions, or manually specify individual IP addresses or address blocks.



Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Leave matching criteria empty whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against every combination of the contents of the criteria you specify.

Geolocation in Network Conditions

Some rules can match traffic using the geographical location of the source or destination. If a rule type supports geolocation, you can mix network and geolocation criteria. To ensure you are using up-to-date geolocation data to filter your traffic, Cisco strongly recommends you regularly update the geolocation database (GeoDB).

Original Client in Network Conditions (Filtering Proxied Traffic)

For some rules, you can handle proxied traffic based on the originating client. Use a source network condition to specify proxy servers, then add an original client constraint to specify original client IP addresses. The system uses a packet's X-Forwarded-For (XFF), True-Client-IP, or custom-defined HTTP header field to determine original client IP.

Traffic matches the rule if the proxy's IP address matches the rule's source network constraint, **and** the original client's IP address matches the rule's original client constraint. For example, to allow traffic from a specific original client address, but only if it uses a specific proxy, create three access control rules:

Access Control Rule 1: Blocks non-proxied traffic from a specific IP address (209.165.201.1)

Source Networks: 209.165.201.1
Original Client Networks: none/any
Action: Block

Access Control Rule 2: Allows proxied traffic from the same IP address, but only if the proxy server for that traffic is one you choose (209.165.200.225 or 209.165.200.238)

Source Networks: 209.165.200.225 and 209.165.200.238
Original Client Networks: 209.165.201.1
Action: Allow

Access Control Rule 3: Blocks proxied traffic from the same IP address if it uses any other proxy server.

Source Networks: any
Original Client Networks: 209.165.201.1
Action: Block

Rules with Network Conditions

Rule Type	Supports Geolocation Constrains?	Supports Original Client Constraints?
Access control	yes	yes
Prefilter	no	no
SSL	yes	no
DNS (source networks only)	no	no
Identity	yes	no
Network analysis	no	no
QoS	yes	no

Configuring Network Conditions

Procedure

Step 1 In the rule editor, click **Networks**.

Step 2 Find and choose the predefined networks you want to add from the **Available Networks** list.

If the rule supports geolocation, you can mix network and geolocation criteria in the same rule:

- Networks—Click **Networks** to choose networks.
- Geolocation—Click **Geolocation** to choose geolocation objects.

Step 3 (Optional) If the rule supports original client constraints, under **Source Networks**, configure the rule to handle proxied traffic based on its original client:

- Source/Proxy—Click **Source** to specify proxy servers.
- Original Client—Click **Original Client** to add a network as an original client constraint. In proxied connections, the original client's IP address must match one of these networks to match the rule.

Step 4 Click **Add to Source**, **Add to Original Client**, or **Add to Destination**, or drag and drop.

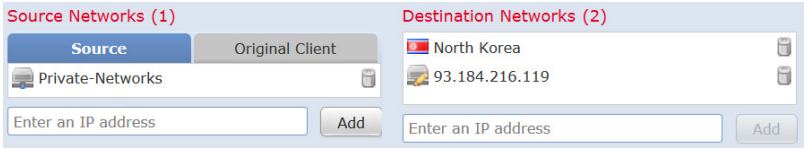
Step 5 Add networks that you want to specify manually. Enter a source, original client, or destination IP address or address block, then click **Add**.

Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Step 6 Save or continue editing the rule.

Example: Network Condition in an Access Control Rule

The following graphic shows the network condition for an access control rule that blocks connections originating from your internal network and attempting to access resources either in North Korea or on 93.184.216.119 (example.com).



In this example, a network object group called Private Networks (that comprises the IPv4 and IPv6 Private Networks network objects, not shown) represents your internal networks. The example also manually specifies the example.com IP address, and uses a system-provided North Korea geolocation object to represent North Korea IP addresses.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Tunnel Endpoint Conditions

Tunnel endpoint conditions are specific to tunnel rules. They are similar to the network conditions for other rule types.

Tunnel endpoint conditions control certain types of plaintext, passthrough tunnels (see Encapsulation Conditions, on page 375) by their source and destination IP address, using outer encapsulation headers. These are the IP addresses of the tunnel endpoints—the routed interfaces of the network devices on either side of the tunnel.

Tunnel rules are bidirectional by default, and handle all matching tunnels between any of the source endpoints and any of the destination endpoints. However, you can configure unidirectional tunnel rules that match source-to-destination traffic only; see Tunnel and Prefilter Rule Components, on page 1281.

You can use predefined network objects to build tunnel endpoint conditions, or manually specify individual IP addresses or address blocks.



Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Configuring Tunnel Endpoint Conditions

Tunnel endpoint conditions apply to Firepower Threat Defense devices only.

Procedure

Step 1 In the rule editor, click **Tunnel Endpoints**.

Step 2 Find and choose the predefined networks you want to add from the **Available Tunnel Endpoints** list.

Because tunnel endpoints are simply the IP addresses of the routed interfaces of the network devices on either side of the tunnel, you can use network objects to build tunnel endpoint conditions.

Step 3 Click **Add to Source** or **Add to Destination**, or drag and drop.

Tunnel rules are bidirectional by default so they can handle all traffic between the two endpoints. However, if you choose to **Match tunnels only from source**, the tunnel rule matches source-to-destination traffic only.

Step 4 Add endpoints that you want to specify manually. Enter a source or destination IP address or address block, then click **Add**.

Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Step 5 Save or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

VLAN Conditions

VLAN rule conditions control VLAN-tagged traffic, including Q-in-Q (stacked VLAN) traffic. The system uses the innermost VLAN tag to filter VLAN traffic, with the exception of the prefilter policy, which uses the outermost VLAN tag in its rules.

Note the following Q-in-Q support:

- NGIPSv, Firepower 7000, Firepower 8000—Supports Q-in-Q for all interface types.
- ASA FirePOWER module—Does not support Q-in-Q (supports only one VLAN tag).
- FTD on Firepower 4100/9300—Does not support Q-in-Q (supports only one VLAN tag).
- FTD on all other models:
  - Inline sets and passive interfaces—Supports Q-in-Q, up to 2 VLAN tags.
  - Firewall interfaces—Does not support Q-in-Q (supports only one VLAN tag).

You can use predefined objects to build VLAN conditions, or manually enter any VLAN tag from **1** to **4094**. Use a hyphen to specify a range of VLAN tags.

You can specify a maximum of 50 VLAN conditions.

Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal VLAN tags to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Rules with VLAN Conditions

The folllowing rule types support VLAN conditions:

- Access control

Note For FTD, VLAN tags in access rules only apply to inline sets; they cannot be used in access rules applied to firewall interfaces.

- Tunnel and prefilter (uses outermost VLAN tag)
- SSL
- DNS
- Identity
- Network analysis

Port and ICMP Code Conditions

Port conditions allow you to control traffic by its source and destination ports. Depending on the rule type, "port" can represent any of the following:

- TCP and UDP—You can control TCP and UDP traffic based on the transport layer protocol. The system represents this configuration using the protocol number in parentheses, plus an optional associated port or port range. For example: TCP(6)/22.
- ICMP—You can control ICMP and ICMPv6 (IPv6-ICMP) traffic based on its internet layer protocol plus an optional type and code. For example: ICMP(1):3:3.
- No port—You can control traffic using other protocols that do not use ports.

Leave matching criteria empty whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against every combination of the contents of the criteria you specify.

Best Practices for Port-Based Rules

Specifying ports is the traditional way to target applications. However, applications can be configured to use unique ports to bypass access control blocks. Thus, whenever possible, use application filtering criteria rather than port criteria to target traffic.

Application filtering is also recommended for applications, like FTP, that open separate channels dynamically for control vs. data flow. Using port-based access control rules can prevent these kinds of applications from performing correctly, and could result in blocking desirable connections.

Using Source and Destination Port Constraints

If you add both source and destination port constraints, you can only add ports that share a single transport protocol (TCP or UDP). For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).

If you add only source ports or only destination ports, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as source port conditions in a single access control rule.

Matching Non-TCP Traffic with Port Conditions

Although you can configure port conditions to match non-TCP traffic, there are some restrictions:

- Access control rules—For Classic devices, you can match GRE-encapsulated traffic with an access control rule by using the GRE (47) protocol as a destination port condition. To a GRE-constrained rule, you can add only network-based conditions: zone, IP address, port, and VLAN tag. Also, the system uses outer headers to match **all** traffic in access control policies with GRE-constrained rules. For Firepower Threat Defense devices, use tunnel rules in the prefilter policy to control GRE-encapsulated traffic.
- SSL rules—SSL rules support TCP port conditions only.
- Identity rules—The system cannot enforce active authentication on non-TCP traffic. If an identity rule action is Active Authentication or if you check the option to **Use active authentication if passive authentication cannot identify user**, use TCP ports constraints only. If the identity rule action is Passive Authentication or No Authentication, you can create port conditions based on non-TCP traffic.

Caution Adding the first or removing the last active authentication rule when SSL decryption is disabled (that is, when the access control policy does not include an SSL policy) restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Note that an active authentication rule has either an **Active Authentication** rule action, or a **Passive Authentication** rule action with **Use active authentication if passive authentication cannot identify user** selected.

- IMCP echo—A destination ICMP port with the type set to 0 or a destination ICMPv6 port with the type set to 129 only matches unsolicited echo replies. ICMP echo replies sent in response to ICMP echo requests are ignored. For a rule to match on any ICMP echo, use ICMP type 8 or ICMPv6 type 128.

Rules with Port Conditions

The following rules support port conditions:

- Access control
- Prefilter
- SSL (supports TCP traffic only)
- Identity (active authentication supports TCP traffic only)
- QoS

Configuring Port Conditions

Procedure

Step 1 In the rule editor, click **Ports**.

Step 2 Find and choose the predefined ports you want to add from the **Available Ports** list.

Step 3 Click **Add to Source** or **Add to Destination**, or drag and drop.

Step 4 Add any source or destination ports that you want to specify manually:

- Source—Choose a **Protocol**, enter a single **Port** from 0 to 65535, and click **Add**.
- Destination (non-ICMP)—Choose or enter a **Protocol**. If you do not want to specify a protocol, or if you choose **TCP** or **UDP**, enter a single **Port** from 0 to 65535. Click **Add**.
- Destination (ICMP)—Choose **ICMP** or **IPv6-ICMP** from the **Protocol** drop down list, then choose a **Type** and related **Code** in the pop-up window that appears. For more information on ICMP types and codes, see the Internet Assigned Numbers Authority (IANA) website.

Step 5 Save or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Encapsulation Conditions

Encapsulation conditions are specific to tunnel rules.

These conditions control certain types of plaintext, passthrough tunnels by their encapsulation protocol. You must choose at least one protocol to match before you can save the rule. You can choose:

- GRE (47)
- IP-in-IP (4)
- IPv6-in-IP (41)
- Teredo (UDP (17)/3455)

Application Conditions (Application Control)

When the system analyzes IP traffic, it can identify and classify the commonly used applications on your network. This discovery-based *application awareness* is the basis for *application control*—the ability to control application traffic.

System-provided *application filters* help you perform application control by organizing applications according to basic characteristics: type, risk, business relevance, category, and tags. You can create reuseable user-defined filters based on combinations of the system-provided filters, or on custom combinations of applications.

At least one detector must be enabled for each application rule condition in the policy. If no detector is enabled for an application, the system automatically enables all system-provided detectors for the application; if none exist, the system enables the most recently modified user-defined detector for the application. For more information about application detectors, see Application Detector Fundamentals, on page 1874.

You can use both application filters and individually specified applications to ensure complete coverage. However, understand the following note before you order your access control rules.

As part of application control, you can also use access control rules to enforce content restriction (such as Safe Search and YouTube EDU).



Caution Failure to set up your access control rules properly can have unexpected results, including traffic being allowed that should be blocked. In general, application control rules should be lower in your access control list because it takes longer for those rules to match than rules based on IP address, for example.

Access control rules that use *specific* conditions (such as networks and IP addresses) should be ordered *before* rules that use general conditions (such as applications). If you're familiar with the Open Systems Interconnect (OSI) model, use similar numbering in concept. Rules with conditions for layers 1, 2, and 3 (physical, data link, and network) should be ordered first in your access control rules. Conditions for layers 5, 6, and 7 (session, presentation, and application) should be ordered later in your access control rules. For more information about the OSI model, see this Wikipedia article.

Benefits of Application Filters

Application filters help you quickly configure application control. For example, you can easily use system-provided filters to create an access control rule that identifies and blocks all high risk, low business relevance applications. If a user attempts to use one of those applications, the system blocks the session.

Using application filters simplifies policy creation and administration. It assures you that the system controls application traffic as expected. Because Cisco frequently updates and adds application detectors via system and vulnerability database (VDB) updates, you can ensure that the system uses up-to-date detectors to monitor application traffic. You can also create your own detectors and assign characteristics to the applications they detect, automatically adding them to existing filters.

Configurations with Application Conditions

The configurations in the following table help you perform application control. The table also shows how you can constrain application control, depending on the configuration.

Configuration	Type, Risk, Relevance, Category	Tags	User-Defined Filters	Content Restriction
Access control rules	yes	yes	yes	yes
SSL rules	yes	no; automatically constrained to encrypted application traffic by the SSL Protocol tag	no	no
Identity rules (to exempt applications from active authentication)	yes	no; automatically constrained by the User-Agent Exclusion tag	no	no
QoS rules	yes	yes	yes	no
User-defined application filter in the object manager	yes	yes	no; you cannot nest user-defined filters	no
Intelligent Application Bypass (IAB)	yes	yes	yes	no

Related Topics

Overview: Application Detection, on page 1873

Configuring Application Conditions and Filters

To build an application condition or filter, choose the applications whose traffic you want to control from a list of available applications. Optionally (and recommended), constrain the available applications using filters. You can use filters and individually specified applications in the same condition.

Before you begin

- Adaptive profiling **must** be enabled (its default state) as described in Configuring Adaptive Profiles, on page 1814 for access control rules to perform application control.
- For Classic device models, you must have the Control license to configure these conditions.

Procedure

Step 1 Invoke the rule or configuration editor:

- Access control, SSL, QoS rule condition—In the rule editor, click **Applications**.
- Identity rule condition—In the rule editor, click **Realms & Settings** and enable active authentication; see Create an Identity Rule, on page 1962.
- Application filter—On the Application Filters page of the object manager, add or edit an application filter. Provide a unique **Name** for the filter.
- Intelligent Application Bypass (IAB)—In the access control policy editor, click **Advanced**, edit IAB settings, then click **Bypassable Applications and Filters**.

Step 2 (Optional) For an access control rule, enable content restriction features by clicking dimmed for **Safe search** () or **YouTube EDU** () and setting related options.

For additional configuration requirements, see Using Access Control Rules to Enforce Content Restriction, on page 1299.

In most cases, enabling content restriction populates the condition's **Selected Applications and Filters** list with the appropriate values. The system does not automatically populate the list if applications or filters related to content restriction are already present in the list when you enable content restriction.

Continue with the procedure to refine your application and filter selections, or skip to saving the rule.

Step 3 Find and choose the applications you want to add from the **Available Applications** list.

To constrain the applications displayed in **Available Applications**, choose one or more **Application Filters** or search for individual applications.

Tip Click **Information** () next to an application to display summary information and internet search links. **Unlock** marks applications that the system can identify only in decrypted traffic.

When you choose filters, singly or in combination, the Available Applications list updates to display only the applications that meet your criteria. You can choose system-provided filters in combination, but not user-defined filters.

- Multiple filters for the same characteristic (risk, business relevance, and so on)—Application traffic must match only one of the filters. For example, if you choose both the medium and high-risk filters, the Available Applications list displays all medium and high-risk applications.
- Filters for different application characteristics—Application traffic must match both filter types. For example, if you choose both the high-risk and low business relevance filters, the Available Applications list displays only applications that meet both criteria.

Step 4 Click **Add to Rule**, or drag and drop.

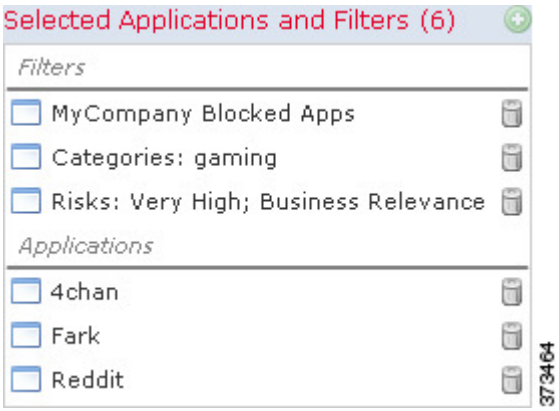
Tip Before you add more filters and applications, click **Clear Filters** to clear your current choices.

The web interface lists filters added to a condition above and separately from individually added applications.

Step 5 Save or continue editing the rule or configuration.

Example: Application Condition in an Access Control Rule

The following graphic shows the application condition for an access control rule that blocks a user-defined application filter for MyCompany, all applications with high risk and low business relevance, gaming applications, and some individually selected applications.



What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Application Characteristics

The system characterizes each application that it detects using the criteria described in the following table. Use these characteristics as application filters.

Table 54: Application Characteristics

Characteristic	Description	Example
Type	Application protocols represent communications between hosts. Clients represent software running on a host. Web applications represent the content or requested URL for HTTP traffic.	HTTP and SSH are application protocols. Web browsers and email clients are clients. MPEG video and Facebook are web applications.
Risk	The likelihood that the application is being used for purposes that might be against your organization's security policy.	Peer-to-peer applications tend to have a very high risk.

Characteristic	Description	Example
Business Relevance	The likelihood that the application is being used within the context of your organization's business operations, as opposed to recreationally.	Gaming applications tend to have a very low business relevance.
Category	A general classification for the application that describes its most essential function. Each application belongs to at least one category.	Facebook is in the social networking category.
Tag	Additional information about the application. Applications can have any number of tags, including none.	Video streaming web applications often are tagged `high bandwidth` and `displays ads`.

Best Practices for Application Control

Keep in mind the following guidelines and limitations for application control:

Ensuring that Adaptive Profiling is Enabled

If adaptive profiling is not enabled (its default state), access control rules cannot perform application control.

Automatically Enabling Application Detectors

If no detector is enabled for an application you want to detect, the system automatically enables all system-provided detectors for the application. If none exist, the system enables the most recently modified user-defined detector for the application.

Configure Your Policy to Examine the Packets That Must Pass Before an Application Is Identified

The system cannot perform application control, including Intelligent Application Bypass (IAB) and rate limiting, before *both* of the following occur:

- A monitored connection is established between a client and server
- The system identifies the application in the session

This identification should occur in 3 to 5 packets, or after the server certificate exchange in the SSL handshake if the traffic is encrypted.

Important! To ensure that your system examines these initial packets, see Specify a Policy to Handle Packets That Pass Before Traffic Identification, on page 1670.

If early traffic matches all other criteria but application identification is incomplete, the system allows the packet to pass and the connection to be established (or the SSL handshake to complete). After the system completes its identification, the system applies the appropriate action to the remaining session traffic.

Create Separate Rules for URL and Application Filtering

Create separate rules for URL and application filtering whenever possible, because combining application and URL criteria can lead to unexpected results, especially for encrypted traffic.

Rules that include both application and URL criteria should come after application-only or URL-only rules, unless the application+URL rule is acting as an exception to a more general application-only or URL-only rule.

URL Rules Before Application and Other Rules

For the most effective URL matching, place rules that include URL conditions before other rules, particularly if the URL rules are block rules and the other rules meet both of the following criteria:

- They include application conditions.
- The traffic to be inspected is encrypted.

Application Control for Encrypted and Decrypted Traffic

The system can identify and filter encrypted and decrypted traffic:

- Encrypted traffic—The system can detect application traffic encrypted with StartTLS, including SMTPS, POPS, FTPS, TelnetS, and IMAPS. In addition, it can identify certain encrypted applications based on the Server Name Indication in the TLS ClientHello message, or the subject distinguished name value from the server certificate. These applications are tagged `SSL Protocol`; in an SSL rule, you can choose only these applications. Applications without this tag can only be detected in unencrypted or decrypted traffic.
- Decrypted traffic—The system assigns the `decrypted traffic` tag to applications that the system can detect in decrypted traffic only, not encrypted or unencrypted.

Exempting Applications from Active Authorization

In an identity policy, you can exempt certain applications from active authentication, allowing traffic to continue to access control. These applications are tagged `User-Agent Exclusion`. In an identity rule, you can choose only these applications.

Handling Application Traffic Packets Without Payloads

When performing access control, the system applies the default policy action to packets that do not have a payload in a connection where an application is identified.

Handling Referred Application Traffic

To handle traffic referred by a web server, such as advertisement traffic, match the referred application rather than the referring application.

Controlling Application Traffic That Uses Multiple Protocols (Skype, Zoho)

Some applications use multiple protocols. To control their traffic, make sure your access control policy covers all relevant options. For example:

- Skype—To control Skype traffic, choose the **Skype** tag from the **Application Filters** list rather than selecting individual applications. This ensures that the system can detect and control all Skype traffic the same way.
- Zoho—To control Zoho mail, choose *both* **Zoho** and **Zoho mail** from the Available Application list.

Search Engines Supported for Content Restriction Features

The system supports Safe Search filtering for specific search engines only. The system assigns the `safesearch supported` tag to application traffic from these search engines.

Controlling Evasive Application Traffic

See Application-Specific Notes and Limitations, on page 382.

Additional Guidelines for Rule Ordering for Application Control

For guidelines about rule ordering for application control, see Best Practices for Configuring Application Control, on page 381.

Related Topics

Inspection of Packets That Pass Before Traffic Is Identified, on page 1670
Special Considerations for Application Detection, on page 1877

Best Practices for Configuring Application Control

We recommend controlling applications' access to the network as follows:

- To allow or block application access from a less secure network to a more secure network: Use **Port** (Selected Destination Port) conditions on the access control rule

  For example, allow ICMP traffic from the internet (less secure) to an internal network (more secure.)

- To allow or block applications being accessed by user groups: Use **Application** conditions on the access control rule

  For example, block Facebook from being accessed by members of the Contractors group



Caution Failure to set up your access control rules properly can have unexpected results, including traffic being allowed that should be blocked. In general, application control rules should be lower in your access control list because it takes longer for those rules to match than rules based on IP address, for example.

Access control rules that use *specific* conditions (such as networks and IP addresses) should be ordered *before* rules that use general conditions (such as applications). If you're familiar with the Open Systems Interconnect (OSI) model, use similar numbering in concept. Rules with conditions for layers 1, 2, and 3 (physical, data link, and network) should be ordered first in your access control rules. Conditions for layers 5, 6, and 7 (session, presentation, and application) should be ordered later in your access control rules. For more information about the OSI model, see this Wikipedia article.

The following table provides an example of how to set up your access control rules:

Type of control	Action	Zones, Networks, VLAN Tags	Users	Applications	Ports	URLs	SGT/ISE Attributes	Inspection, Logging, Comments
Application from more secure to less secure network when application uses a port (for example, SSH)	Your choice (**Allow** in this example)	Destination zones or networks using the outside interface	Any	Do not set	Available Ports : **SSH** Add to **Selected Destination Ports**	Any	Use only with ISE.	Any
Application from more secure to less secure network when application does not use a port (for example, ICMP)	Your choice (**Allow** in this example)	Destination zones or networks using the outside interface	Any	Do not set	Selected Destination Ports **Protocol**: **ICMP** **Type**: **Any**	Do not set	Use only with ISE.	Any
Application access by a user group	Your choice (**Block** in this example)	Your choice	Choose a user group (Contractors group in this example)	Choose the name of the application (**Facebook** in this example)	Do not set	Do not set	Use only with ISE.	Your choice

Application-Specific Notes and Limitations

- Office 365 Admin Portal:

  Limitation: If the access policy has logging enabled at the beginning as well as at the end, the first packet will be detected as Office 365 and the end of connection will be detected as Office 365 Admin Portal. This should not affect blocking.

- Skype:

  See Best Practices for Application Control, on page 379

- GoToMeeting

  In order to fully detect GoToMeeting, your rule must include all of the following applications:

  - GoToMeeting
  - Citrix Online
  - Citrix GoToMeeting Platform
  - LogMeIn
  - STUN

- Zoho:

  See Best Practices for Application Control, on page 379

- Evasive applications such as Bittorrent, Tor, Psiphon, and Ultrasurf:

  For evasive applications, only the highest-confidence scenarios are detected by default. If you need to take action on this traffic (such as block or implement QoS), it may be necessary to configure more aggressive detection with better effectiveness. To do this, contact TAC to review your configurations as these changes may result in false positives.

- WeChat:

  It is not possible to selectively block WeChat Media if you allow WeChat.

Troubleshoot Application Control Rules

If your application control rules don't function as you expect, use the guidelines discussed in this section.

We recommend controlling applications' access to the network as follows:

- To allow or block application access from a less secure network to a more secure network: Use **Port** (Selected Destination Port) conditions on the access control rule

  For example, allow ICMP traffic from the internet (less secure) to an internal network (more secure.)

- To allow or block applications being accessed by user groups: Use **Application** conditions on the access control rule

  For example, block Facebook from being accessed by members of the Contractors group



Caution Failure to set up your access control rules properly can have unexpected results, including traffic being allowed that should be blocked. In general, application control rules should be lower in your access control list because it takes longer for those rules to match than rules based on IP address, for example.

Access control rules that use *specific* conditions (such as networks and IP addresses) should be ordered *before* rules that use general conditions (such as applications). If you're familiar with the Open Systems Interconnect (OSI) model, use similar numbering in concept. Rules with conditions for layers 1, 2, and 3 (physical, data link, and network) should be ordered first in your access control rules. Conditions for layers 5, 6, and 7 (session, presentation, and application) should be ordered later in your access control rules. For more information about the OSI model, see this Wikipedia article.

The following table provides an example of how to set up your access control rules:

Type of control	Action	Zones, Networks, VLAN Tags	Users	Applications	Ports	URLs	SGT/ISE Attributes	Inspection, Logging, Comments
Application from more secure to less secure network when application uses a port (for example, SSH)	Your choice (**Allow** in this example)	Destination zones or networks using the outside interface	Any	Do not set	Available Ports : **SSH** Add to **Selected Destination Ports**	Any	Use only with ISE.	Any
Application from more secure to less secure network when application does not use a port (for example, ICMP)	Your choice (**Allow** in this example)	Destination zones or networks using the outside interface	Any	Do not set	Selected Destination Ports **Protocol**: **ICMP** **Type**: **Any**	Do not set	Use only with ISE.	Any
Application access by a user group	Your choice (**Block** in this example)	Your choice	Choose a user group (Contractors group in this example)	Choose the name of the application (**Facebook** in this example)	Do not set	Do not set	Use only with ISE.	Your choice

Initial Packets Are Passing Uninspected

See Inspection of Packets That Pass Before Traffic Is Identified, on page 1670 and subtopics.

Related Topics

Best Practices for Ordering Rules, on page 1194

URL Conditions (URL Filtering)

Use URL conditions to control the websites that users on your network can access.

For complete information, see URL Filtering, on page 1229.

User, Realm, and ISE Attribute Conditions (User Control)

You can perform *user control* with the *authoritative user identity data* collected by the Firepower System.

Identity sources monitor users as they log in and out, or as they authenticate using Microsoft Active Directory (AD) or LDAP credentials. You can then configure rules that use this collected identity data to handle traffic based on the logged-in authoritative user associated with a monitored host. A user remains associated with a

host until the user logs off (as reported by an identity source), a realm times out the session, or you delete the user data from the system's database.

For information on the authoritative user identity sources supported in your version of the Firepower System, see About User Identity Sources, on page 1891.

You can use the following rule conditions to perform user control:

- User and realm conditions—Match traffic based on the logged-in authoritative user of a host. You can control traffic based on realms, individual users, or the groups those users belong to.
- ISE attribute conditions—Match traffic based on a user's ISE-assigned Security Group Tag (SGT), Device Type (also referred to as Endpoint Profile), or Location IP (also referred to as Endpoint Location). Requires that you configure ISE as an identity source.



Note In some rules, custom SGT conditions can match traffic tagged with SGT attributes that were **not** assigned by ISE. This is not considered user control, and only works if you are not using ISE as an identity source; see Custom SGT Conditions, on page 389.

Rules with User Conditions

Rule Type	Supports User and Realm Conditions?	Supports ISE Attribute Conditions?
Access control	yes	yes
SSL	yes	no
QoS	yes	yes

Related Topics

The User Agent Identity Source, on page 1893
The ISE Identity Source, on page 1895
The Terminal Services (TS) Agent Identity Source, on page 1901
The Captive Portal Identity Source, on page 1903

User Control Prerequisites

Configure Identity Sources/Authentication Methods

Configure identity sources for the types of authentication you want to perform. For more information, see About User Identity Sources, on page 1891.

If you configure an ISE or user agent device to monitor a large number of user groups, or if you have a very large number of users mapped to hosts on your network, the system may drop user mappings based on groups, due to your Firepower Management Center user limit. As a result, rules with realm, user, or user group conditions may not match traffic as expected.

Configure Realms

Configure a realm for each AD or LDAP server you want to monitor, including your ISE, User Agent, and perform a user download. For more information, see Create a Realm, on page 1949.

Note If you are configuring an ISE SGT attribute rule condition, configuring a realm is optional. The ISE SGT attribute rule condition can be configured in policies with or without an associated identity policy (where realms are invoked).

When you configure a realm, you specify the users and user groups whose activity you want to monitor. Including a user group automatically includes all of that group's members, including members of any secondary groups. However, if you want to use the secondary group as a rule criterion, you must explicitly include the secondary group in the realm configuration.

For each realm, you can enable automatic download of user data to refresh authoritative data for users and user groups.

Create Identity Policies

Create an identity policy to associate the realm with an authentication method, and associate that policy with access control. For more information, see Create an Identity Policy, on page 1961.

Policies that perform user control on a device (access control, SSL, QoS) share an identity policy. That identity policy determines the realms, users, and groups that you can use in rules affecting traffic on those devices.

Before you configure a user condition in a QoS rule, you **must** make sure the devices targeted by the QoS policy are using the correct identity policy, as defined in the access control policy deployed to the devices. Because the QoS policy and access control policy deployed to the same device are not explicitly linked, the QoS rule editor can allow you to choose invalid realms, users, and groups. These invalid elements are those from identity policies that exist on the Firepower Management Center, but that are not applied to the QoS-targeted devices. If you use these elements, the system cannot determine that you made an invalid choice until deploy-time.

Configuring User and Realm Conditions

You can constrain a rule by realm, or by users and user groups within that realm.

Before you begin

- Fulfill the user control prerequsities described in User, Realm, and ISE Attribute Conditions (User Control), on page 384.
- For Classic device models, you must have the Control license to configure these conditions.

Procedure

Step 1 In the rule editor, click **Users**.

Step 2 (Optional) Find and choose the realm you want to use from the **Available Realms**.

Step 3 (Optional) Further constrain the rule by choosing users and groups from the **Available Users** list.

Step 4 Click **Add to Rule**, or drag and drop.

Step 5 Save or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configuring ISE Attribute Conditions

Before you begin

- Fulfill the user control prerequisites described in User, Realm, and ISE Attribute Conditions (User Control), on page 384.
- For Classic device models, you must have the Control license to configure these conditions.

Procedure

Step 1 In the rule editor, click the following for ISE attribute conditions:

- Access control—Click **SGT/ISE Attributes**.

You can use ISE-assigned Security Group Tags (SGTs) to constrain ISE attribute conditions. To use custom SGTs in access control rules, see Custom SGT Conditions, on page 389.

Step 2 Find and choose the ISE attributes you want to use from the **Available Attributes** list:

- Security Group Tag (SGT)
- Device Type (also referred to as Endpoint Profile)
- QoS—Click **ISE Attributes**.
- Location IP (also referred to as Endpoint Location)

Step 3 Further constrain the rule by choosing attribute metadata from the **Available Metadata** list. Or, keep the default: **any**.

Step 4 Click **Add to Rule**, or drag and drop.

Step 5 (Optional) Constrain the rule with an IP address in the **Add a Location IP Address** field, then click **Add**.

The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.

Step 6 Save or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Troubleshoot User Control

If you notice unexpected user rule behavior, consider tuning your rule, identity source, or realm configurations. For other related troubleshooting information, see:

- Troubleshoot the User Agent Identity Source, on page 1894
- Troubleshoot ISE or Cisco TrustSec Issues, on page 1900
- Troubleshoot the TS Agent Identity Source, on page 1902
- Troubleshoot the Captive Portal Identity Source, on page 1914
- Troubleshoot Realms and User Downloads, on page 1945

Rules targeting realms, users, or user groups are not matching traffic

If you configure a user agent or ISE device to monitor a large number of user groups, or if you have a very large number of users mapped to hosts on your network, the system may drop user records due to your Firepower Management Center user limit. As a result, rules with user conditions may not match traffic as expected.

Rules targeting user groups or users within user groups are not matching traffic as expected

If you configure a rule with a user group condition, your LDAP or Active Directory server must have user groups configured. The system cannot perform user group control if the server organizes the users in basic object hierarchy.

Rules targeting users in secondary groups are not matching traffic as expected

If you configure a rule with a user group condition that includes or excludes users who are members of a secondary group on your Active Directory server, your server may be limiting the number of users it reports.

By default, Active Directory servers limit the number of users they report from secondary groups. You must customize this limit so that all of the users in your secondary groups are reported to the Firepower Management Center and eligible for use in rules with user conditions.

Rules are not matching users when seen for the first time

After the system detects activity from a previously-unseen user, the system retrieves information about them from the server. Until the system successfully retrieves this information, activity seen by this user is *not* handled by matching rules. Instead, the user session is handled by the next rule it matches (or the policy's default action, if applicable).

For example, this might explain when:

- Users who are members of user groups are not matching rules with user group conditions.
- Users who were reported by a TS Agent, user agent, or ISE device are not matching rules, when the server used for user data retrieval is an Active Directory server.

Note that this might also cause the system to delay the display of user data in event views and analysis tools.

Rules are not matching all ISE users

This is expected behavior. You can perform user control on ISE users who were authenticated by an Active Directory domain controller. You cannot perform user control on ISE users who were authenticated by an LDAP, RADIUS, or RSA domain controller.

Custom SGT Conditions

If you do not configure Cisco ISE as an identity source, you can control traffic using Security Group Tags (SGTs) that were **not** assigned by ISE. SGTs specify the privileges of traffic sources within a trusted network.

Custom SGT rule conditions use manually created SGT objects to filter traffic, rather than ISE SGTs obtained from the system's connection to an ISE server. These manually created SGT objects correspond to the SGT attributes on the traffic you want to control. Controlling traffic using custom SGTs is not considered user control.

Rules with Custom SGT Conditions

Only access control rules support custom SGT conditions.

ISE SGT vs Custom SGT Rule Conditions

Some rules allow you to control traffic based on assigned SGT. Depending on the rule type and your identity source configuration, you can use either ISE-assigned SGTs or custom SGTs to match traffic with assigned SGT attributes.

Note If you use ISE SGTs to match traffic, even if a packet does not have an assigned SGT attribute, the packet still matches an ISE SGT rule if the SGT associated with the packet's source IP address is known in ISE.

Condition Type	Requires	SGTs Listed in Rule Editor
ISE SGT	ISE identity source	SGTs obtained by querying the ISE server, with automatically updated metadata
Custom SGT	No ISE identity source	Static SGT objects you create

Related Topics

User, Realm, and ISE Attribute Conditions (User Control), on page 384

Autotransition from Custom SGTs to ISE SGTs

If you create rules that match custom SGTs, then configure ISE as an identity source, the system:

- Disables **Security Group Tag** options in the object manager. Although the system retains existing SGT objects, you cannot modify them or add new ones.
- Retains existing rules with custom SGT conditions. However, these rules do not match traffic. You also cannot add additional custom SGT criteria to existing rules, or create new rules with custom SGT conditions.

If you configure ISE, Cisco recommends that you delete or disable existing rules with custom SGT conditions. Instead, use ISE attribute conditions to match traffic with SGT attributes.

Related Topics

Configure ISE for User Control, on page 1897

Configuring Custom SGT Conditions

The following procedure describes how to filter traffic tagged with SGT attributes that were **not** assigned by ISE. This is not considered user control, and only works if you are not using ISE as an identity source; see ISE SGT vs Custom SGT Rule Conditions, on page 389.

Before you begin

- Disable ISE connections. Custom SGT matching does not work if you use ISE as an identity source.
- Configure Security Group Tag objects that correspond with the SGTs you want to match; see Creating Security Group Tag Objects, on page 407.
- For Classic device models, you must have the Control license to configure these conditions.

Procedure

Step 1 In the rule editor, click **SGT/ISE Attributes**.

Step 2 Choose **Security Group Tag** from the **Available Attributes** list.

Step 3 In the **Available Metadata** list, find and choose a custom SGT object.

If you choose **Any**, the rule matches all traffic with an SGT attribute. For example, you might choose this value if you want an access control rule to block traffic from hosts that are not configured for TrustSec.

Step 4 Click **Add to Rule**, or drag and drop.

Step 5 Save or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Troubleshooting Custom SGT Conditions

If you notice unexpected rule behavior, consider tuning your custom SGT object configuration.

Security Group Tag objects unavailable

Custom SGT objects are only available if you do not configure ISE as an identity source. For more information, see Autotransition from Custom SGTs to ISE SGTs, on page 389.

Searching for Rules

In many policies, you can search for and within rules. The system matches your input to rule names and condition values, including objects and object groups.

You cannot search for values in a Security Intelligence or URL list or feed.

Procedure

Step 1 In the policy editor, click **Rules**.

Step 2 Click **Search Rules**, enter a complete or partial search string, then press Enter.
The matching value is highlighted for each matching rule. A status message displays the current match and the total number of matches.

Step 3 View the rules you are interested in.

To navigate between matching rules, click **Next-Match** or **Previous-Match** .

What to do next

- Before you begin a new search, click **Clear** ( ) to clear the search and any highlighting.

Filtering Rules by Device

Some policy editors allow you to filter your rule view by affected devices.

Filter by device only works for rules that use zones or interface groups. (Otherwise a rule applies to all devices.)

The system uses a rule's interface constraints to determine if the rule affects a device. If you constrain a rule by interface (security zone or interface group condition), the device where that interface is located is affected by that rule. Rules with no interface constraint apply to any interface, and therefore every device.

QoS rules are always constrained by interface.

Procedure

Step 1 In the policy editor, click **Rules**, then click **Filter by Device**.
A list of targeted devices and device groups appears.

Step 2 Check one or more check boxes to display only the rules that apply to those devices or groups. Or, check **All** to reset and display all of the rules.

Tip Hover your pointer over a rule criterion to see its value. If the criterion represents an object with device-specific overrides, the system displays the override value when you filter the rules list by only that device. If the criterion represents an object with domain-specific overrides, the system displays the override value when you filter the rules list by devices in that domain.

Step 3 Click **OK**.

Related Topics

Create and Edit Access Control Rules, on page 1220
Configure Prefiltering, on page 1278
Configuring QoS Rules, on page 677
Configure NAT for Threat Defense, on page 1019

Rule and Other Policy Warnings

Policy and rule editors use icons to mark configurations that could adversely affect traffic analysis and flow. Depending on the issue, the system may warn you when you deploy or prevent you from deploying entirely.



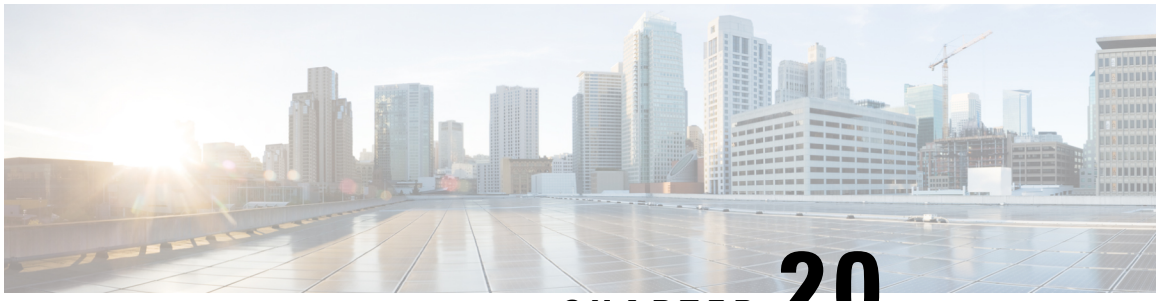
Tip Hover your pointer over an icon to read the warning, error, or informational text.

Table 55: Policy Error Icons

Icon	Description	Example
Errors () error	If a rule or configuration has an error, you cannot deploy until you correct the issue, even if you disable any affected rules.	A rule that performs category and reputation-based URL filtering is valid until you target a device that does not have a URL Filtering license. At that point, an error icon appears next to the rule, and you cannot deploy until you edit or delete the rule, retarget the policy, or enable the license.
Warning () warning	You can deploy a policy that displays rule or other warnings. However, misconfigurations marked with warnings have no effect. If you disable a rule with a warning, the warning icon disappears. It reappears if you enable the rule without correcting the underlying issue.	Preempted rules or rules that cannot match traffic due to misconfiguration have no effect. This includes conditions using empty object groups, application filters that match no applications, excluded LDAP users, invalid ports, and so on. However, if a warning icon marks a licensing error or model mismatch, you cannot deploy until you correct the issue.
Information () information	Information icons convey helpful information about configurations that may affect the flow of traffic. These issues do not prevent you from deploying.	With application control, the system might skip matching the first few packets of a connection against some rules, until the system identifies the application or web traffic in that connection. This allows connections to be established so that applications and HTTP requests can be identified.

Related Topics

Best Practices for Application Control, on page 379
Best Practices for URL Filtering, on page 1230



CHAPTER 20

Reusable Objects

The following topics describe how to manage reusable objects in the Firepower System:

- Introduction to Reusable Objects, on page 394
- The Object Manager, on page 396
- Network Objects, on page 402
- Port Objects, on page 404
- Tunnel Zones, on page 406
- Application Filters, on page 406
- VLAN Tag Objects, on page 406
- Security Group Tag Objects, on page 407
- URL Objects, on page 408
- Geolocation Objects, on page 409
- Interface Objects: Interface Groups and Security Zones, on page 410
- Variable Sets, on page 412
- Security Intelligence Lists and Feeds, on page 427
- Sinkhole Objects, on page 438
- File Lists, on page 438
- Cipher Suite Lists, on page 443
- Distinguished Name Objects, on page 444
- PKI Objects, on page 447
- SLA Monitor Objects, on page 461
- Prefix Lists, on page 463
- Route Maps, on page 464
- Access List, on page 467
- AS Path Objects, on page 470
- Community Lists, on page 470
- Policy Lists, on page 472
- VPN Objects, on page 473
- Address Pools, on page 478
- FlexConfig Objects, on page 479

Introduction to Reusable Objects

For increased flexibility and web interface ease-of-use, the Firepower System uses named *objects*, which are reusable configurations that associate a name with a value. When you want to use that value, use the named object instead. The system supports object use in various places in the web interface, including many policies and rules, event searches, reports, dashboards, and so on. The system provides many predefined objects that represent frequently used configurations.

Use the object manager to create and manage objects. Many configurations that use objects also allow you to create objects on the fly, as needed. You can also use the object manager to:

- View the policies, settings, and other objects where a network, port, VLAN, or URL object is used; see Viewing Objects and Their Usage, on page 397.
- Group objects to reference multiple objects with a single configuration; see Object Groups, on page 398.
- Override object values for selected devices or, in a multidomain deployment, selected domains; see Object Overrides, on page 399.

After you edit an object used in an active policy, you must redeploy the changed configuration for your changes to take effect. You cannot delete an object that is in use by an active policy.



Note An object is configured on a managed device if, and only if, the object is used in a policy that is assigned to that device. If you remove an object from all policies assigned to a given device, the object is also removed from the device configuration on the next deployment, and subsequent changes to the object are not reflected in the device configuration.

Object Types

The following table lists the objects you can create in the Firepower System, and indicates whether each object type can be grouped or configured to allow overrides.

Object Type	Groupable?	Allows Overrides?
Network	yes	yes
Port	yes	yes
Interface: • Security Zone • Interface Group	no	no
Tunnel Zone	no	no
Application Filter	no	no
VLAN Tag	yes	yes
Security Group Tag (SGT)	no	no

Object Type	Groupable?	Allows Overrides?
URL	yes	yes
Geolocation	no	no
Variable Set	no	no
Security Intelligence: Network, DNS, and URL lists and feeds	no	no
Sinkhole	no	no
File List	no	no
Cipher Suite List	no	no
Distinguished Name	yes	no
Public Key Infrastructure (PKI): • Internal and Trusted CA • Internal and External Certs	yes	no
SLA Monitor	no	no
Prefix List: IPv4 and IPv6	no	yes
Route Map	no	yes
Access List: Standard and Extended	no	yes
AS Path	no	yes
Community List	no	yes
Policy List	no	yes
FlexConfig: Text and FlexConfig objects	no	yes

Objects and Multitenancy

In a multidomain deployment, you can create objects in Global and descendant domains with the exception of Security Group Tag (SGT) objects, which you can create only in the Global domain. The system displays objects created in the current domain, which you can edit. It also displays objects created in ancestor domains, which you cannot edit, with the exception of security zones and interface groups.



Note Because security zones and interface groups are tied to device interfaces, which you configure at the leaf level, administrators in descendant domains can view and edit zones and groups created in ancestor domains. Subdomain users can add and delete interfaces from ancestor zones and groups, but cannot delete or rename the zones/groups.

Object names must be unique within the domain hierarchy. The system may identify a conflict with the name of an object you cannot view in your current domain.

For objects that support grouping, you can group objects in the current domain with objects inherited from ancestor domains.

Object overrides allow you to define device-specific or domain-specific values for certain types of object, including network, port, VLAN tag, and URL. In a multidomain deployment, you can define a default value for an object in an ancestor domain, but allow administrators in descendant domains to add override values for that object.

The Object Manager

You can use the object manager to create and manage objects and object groups.

The object manager displays 20 objects or groups per page. If you have more than 20 of any type of object or group, use the navigation links at the bottom of the page to view additional pages. You can also go to a specific page or click **Refresh** ( ) to refresh your view.

By default, the page lists objects and groups alphabetically by name. You can filter the objects on the page by name or value.

Editing Objects

In a multidomain deployment, the system displays objects created in the current domain, which you can edit. It also displays objects created in ancestor domains, which in most cases you cannot edit. To view and edit objects in a descendant domain, switch to that domain.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose an object type from the list; see Introduction to Reusable Objects, on page 394.

Step 3 Click **Edit** ( ) next to the object you want to edit.

If **View** ( ) appears instead, the object belongs to an ancestor domain and has been configured not to allow overrides, or you do not have permission to modify the object.

Step 4 Modify the object settings as desired.

Step 5 If you are editing a variable set, manage the variables in the set; see Managing Variables, on page 424.

Step 6 For objects that can be configured to allow overrides:

- If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 401. You can change this setting only for objects that belong to the current domain.
- If you want to add override values to this object, expand the Override section and click **Add**; see Adding Object Overrides, on page 401.

Step 7 Click **Save**.

Step 8 If you are editing a variable set, and that set is in use by an access control policy, click **Yes** to confirm that you want to save your changes.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Viewing Objects and Their Usage

You can view usage details of objects on the Object Management page. The network, port, VLAN, and URL object types are the only object types that provide this functionality.

Note In a multidomain deployment, you can view objects from any other domain. However, to find usage of objects in a descendant domain, switch to that domain.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose one of the following object types:

- Network
- Port
- VLAN Tag
- URL

Step 3 Click **Find Usage** ( ) next to the object.

The Object Usage window displays a list of all the policies, objects, and other settings where the object is in use. Click any of the listed items to know more about the object usage. For policies and some other settings where the object is used, you can click the corresponding links to visit the respective UI pages.

Filtering Objects or Object Groups

In a multidomain deployment, the system displays objects created in the current and ancestor domains, which you can filter.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Enter your filter criteria in the **Filter** field.

The page updates as you type to display matching items.

You can use the following wildcards:

- The asterisk (*) matches zero or more occurrences of a character.
- The caret (^) matches content at the beginning of a string.
- The dollar sign ($) matches content at the end of a string.

Object Groups

Grouping objects allows you to reference multiple objects with a single configuration. The system allows you to use objects and object groups interchangeably in the web interface. For example, anywhere you would use a port object, you can also use a port object group.

You can group network, port, VLAN tag, URL, and PKI objects. Network object groups can be nested, that is, you can add a network object group to another network object group up to 10 levels.

Objects and object groups of the same type cannot have the same name. In a multidomain deployment, the names of object groups must be unique within the domain hierarchy. Note that the system may identify a conflict with the name of an object group you cannot view in your current domain.

When you edit an object group used in a policy (for example, a network object group used in an access control policy), you must re-deploy the changed configuration for your changes to take effect.

Deleting a group does not delete the objects in the group, just their association with each other. Additionally, you cannot delete a group that is in use in an active policy. For example, you cannot delete a VLAN tag group that you are using in a VLAN condition in a saved access control policy.

Grouping Reusable Objects

In a multidomain deployment, the system displays objects created in the current domain, which you can edit. It also displays objects created in ancestor domains, which in most cases you cannot edit. To view and edit objects in a descendant domain, switch to that domain.

You can group objects in the current domain with objects inherited from ancestor domains.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 If the object type you want to group is **Network**, **Port**, **URL**, or **VLAN Tag**:

a) Choose the object type from the list of object types.
b) Choose **Add Group** from the **Add [Object Type]** drop-down list.

Step 3 If the object type you want to group is **Distinguished Name**:

a) Expand the **Distinguished Name** node.
b) Choose **Object Groups**.
c) Click **Add Distinguished Name Group**.

Step 4 If the object type you want to group is **PKI**:

a) Expand the **PKI** node.

b) Choose one of the following:

- **Internal CA Groups**
- **Trusted CA Groups**
- **Internal Cert Groups**
- **External Cert Groups**

c) Click **Add [Object Type] Group**.

Step 5 Enter a unique **Name**.

Step 6 Choose one or more objects from the list, and click **Add**.

You can also:

- Use the filter field **Search** ( ) to search for existing objects to include, which updates as you type to display matching items. Click **Reload** ( ) above the search field or click **Clear** ( ) in the search field to clear the search string.
- Click **Add** ( ) to create objects on the fly if no existing objects meet your needs.

Step 7 Optionally for **Network**, **Port**, **URL**, and **VLAN Tag** groups:

- Enter a **Description**.
- Check the **Allow Overrides** check box to allow overrides for this object group; see Allowing Object Overrides, on page 401.

Step 8 Click **Save**.

What to do next

- If an active policy references your object group, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Object Overrides

An object override allows you to define an alternate value for an object, which the system uses for the devices you specify.

You can create an object whose definition works for most devices, and then use overrides to specify modifications to the object for the few devices that need different definitions. You can also create an object that needs to be overridden for all devices, but its use allows you to create a single policy for all devices. Object overrides allow you to create a smaller set of shared policies for use across devices without giving up the ability to alter policies when needed for individual devices.

For example, you might want to deny ICMP traffic to the different departments in your company, each of which is connected to a different network. You can do this by defining an access control policy with a rule

that includes a network object called Departmental Network. By allowing overrides for this object, you can then create overrides on each relevant device that specifies the actual network where that device is connected.

In a multidomain deployment, you can define a default value for an object in an ancestor domain and allow administrators in descendant domains to add override values for that object. For example, a managed security service provider (MSSP) might use a single Firepower Management Center to manage network security for multiple customers. Administrators at the MSSP can define an object in the Global domain for use in all customers' deployments. Administrators for each customer can log into descendant domains to override that object for their organizations. These local administrators cannot view or affect the override values of other customers of the MSSP.

You can target an object override to a specific domain. In this case, the system uses the object override value for all devices in the targeted domain unless you override it at the device level.

From the object manager, you can choose an object that can be overridden and define a list of device-level or domain-level overrides for that object.

You can use object overrides with the following object types only:

- Network
- Port
- VLAN tag
- URL
- SLA Monitor
- Prefix List
- Route Map
- Access List
- AS Path
- Community List
- Policy List
- PKI Enrollment

If you can override an object, the **Override** column appears for the object type in the object manager. Possible values for this column include:

- Green checkmark — indicates that you can create overrides for the object and no overrides have been added yet
- Red X — indicates that you cannot create overrides for the object
- Number — represents a count of the overrides that have been added to that object (for example, "2" indicates two overrides have been added)

Managing Object Overrides

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose from the list of object types; see Introduction to Reusable Objects, on page 394.

Step 3 Click **Edit** ( ) next to the object you want to edit.

If **View** ( ) appears instead, the object belongs to an ancestor domain and has been configured not to allow overrides, or you do not have permission to modify the object.

Step 4 Manage the object overrides:

- Add—Add object overrides; see Adding Object Overrides, on page 401.
- Allow—Allow object overrides; see Allowing Object Overrides, on page 401.
- Delete—In the object editor, click **Delete** ( ) next to the override you want to remove.
- Edit—Edit object overrides; see Editing Object Overrides, on page 402.

Allowing Object Overrides

Procedure

Step 1 In the object editor, check the **Allow Overrides** check box.

Step 2 Click **Save**.

What to do next

Add object override values; see Adding Object Overrides, on page 401.

Adding Object Overrides

Before you begin

Allow object overrides; see Allowing Object Overrides, on page 401.

Procedure

Step 1 In the object editor, expand the **Override** section.

Step 2 Click **Add**.

Step 3 On **Targets**, choose domains or devices in the **Available Devices and Domains** list and click **Add**.

Step 4 On the Override tab, enter a **Name**.

Step 5 Optionally, enter a **Description**.

Step 6 Enter an override value.

Example:

For a network object, enter a network value.

Step 7 Click **Add**.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Editing Object Overrides

You can modify the description and the value of an existing override, but you cannot modify the existing target list. Instead, you must add a new override with new targets, which replaces the existing override.

Procedure

Step 1 In the object editor, expand the **Override** section.

Step 2 Click **Edit** ( ) next to the override you want to modify.

Step 3 Optionally, modify the **Description**.

Step 4 Modify the override value.

Step 5 Click **Save** to save the override.

Step 6 Click **Save** to save the object.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Network Objects

A network object represents one or more IP addresses. You can use network objects and groups in various places in the system's web interface, including access control policies, network variables, identity rules, network discovery rules, event searches, reports, and so on.

When you configure an option that requires a network object, the list is automatically filtered to show only those objects that are valid for the option. For example, some options require host objects, while other options require subnets.

A network object can be one of the following types:

Host

A single IP address.

IPv4 example:

`209.165.200.225`

IPv6 example:

`2001:DB8::0DB8:800:200C:417A` or `2001:DB8:0:0:0DB8:800:200C:417A`

Range

A range of IP addresses.

IPv4 example:

`209.165.200.225-209.165.200.250`

IPv6 example:

`2001:db8:0:cd30::1-2001:db8:0:cd30::1000`

Network

An address block, also known as a subnet.

IPv4 example:

`209.165.200.224/27`

IPv6 example:

`2001:DB8:0:CD30::/60`



Note Security Intelligence ignores IP address blocks using a `/0` netmask.

Group

A group of network objects or other network object groups.

For example:

```
209.165.200.225
209.165.201.1
209.165.202.129
```

You can create nested groups by adding one network object group to another network object group. You can nest up to 10 levels of groups.

Creating Network Objects

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Network** from the list of object types.

Step 3 Choose **Add Object** from the **Add Network** drop-down menu.

Step 4 Enter a **Name**.

In a multidomain deployment, object names must be unique within the domain hierarchy. The system may identify a conflict with the name of an object you cannot view in your current domain.

Step 5 Optionally, enter a **Description**.

Step 6 In the **Network** field, select the required option and enter an appropriate value; see Network Objects, on page 402.

Step 7 Manage overrides for the object:

- If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 401.
- If you want to add override values to this object, expand the Override section and click **Add**; see Adding Object Overrides, on page 401.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Port Objects

Port objects represent different protocols in slightly different ways:

TCP and UDP

A port object represents the transport layer protocol, with the protocol number in parentheses, plus an optional associated port or port range. For example: `TCP(6)/22`.

ICMP and ICMPv6 (IPv6-ICMP)

A port object represents the Internet layer protocol plus an optional type and code. For example: `ICMP(1):3:3`.

You can restrict an ICMP or IPV6-ICMP port object by type and, if applicable, code. For more information on ICMP types and codes, see:

- http://www.iana.org/assignments/icmp-parameters/icmp-parameters.xml
- http://www.iana.org/assignments/icmpv6-parameters/icmpv6-parameters.xml

Other

A port object can represent other protocols that do not use ports.

The Firepower System provides default port objects for well-known ports. You cannot modify or delete these default objects. You can create custom port objects in addition to the default objects.

You can use port objects and groups in various places in the system's web interface, including access control policies, identity rules, network discovery rules, port variables, and event searches. For example, if your

organization uses a custom client that uses a specific range of ports and causes the system to generate excessive and misleading events, you can configure your network discovery policy to exclude monitoring those ports.

When using port objects, observe the following guidelines:

- You cannot add any protocol other than TCP or UDP for source port conditions in access control rules. Also, you cannot mix transport protocols when setting both source and destination port conditions in a rule.
- If you add an unsupported protocol to a port object group used in a source port condition, the rule where it is used does not take affect on the managed device when the configuration is deployed.
- If you create a port object containing both TCP and UDP ports, then add it as a source port condition in a rule, you cannot add a destination port, and vice versa.

Creating Port Objects

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Port** from the list of object types.

Step 3 Choose **Add Object** from the **Add Port** drop-down list.

Step 4 Enter a **Name**.

Step 5 Choose a **Protocol**.

Step 6 Depending on the protocol you chose, constrain by **Port**, or choose an ICMP **Type** and **Code**.

You can enter ports from `1` to `65535`. Use a hyphen to specify a port range. You must constrain the object by port if you chose to match **All** protocols, using the **Other** drop-down list.

Step 7 Manage overrides for the object:

- If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 401.
- If you want to add override values to this object, expand the Override section and click **Add**; see Adding Object Overrides, on page 401.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Tunnel Zones

A *tunnel zone* represents certain types of plaintext, passthrough tunnels that you explicitly tag for special analysis. A tunnel zone is not an interface object, even though you can use it as an interface constraint in some configurations.

For detailed information, see Tunnel Zones and Prefiltering, on page 1282.

Application Filters

System-provided application filters help you perform application control by organizing applications according to basic characteristics: type, risk, business relevance, category, and tags. In the object manager, you can create and manage reuseable user-defined application filters based on combinations of the system-provided filters, or on custom combinations of applications. For detailed information, see Application Conditions (Application Control), on page 375.

VLAN Tag Objects

Each VLAN tag object you configure represents a VLAN tag or range of tags.

You can group VLAN tag objects. Groups represent multiple objects; using a range of VLAN tags in a single object is not considered a group in this sense.

You can use VLAN tag objects and groups in various places in the system's web interface, including rules and event searches. For example, you could write an access control rule that applies only to a specific VLAN.

Creating VLAN Tag Objects

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **VLAN Tag** from the list of object types.

Step 3 Choose **Add Object** from the **Add VLAN Tag** drop-down list.

Step 4 Enter a **Name**.

Step 5 Enter a **Description**.

Step 6 Enter a value in the **VLAN Tag** field. Use a hyphen to specify a range of VLAN tags.

Step 7 Manage overrides for the object:

- If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 401.
- If you want to add override values to this object, expand the Override section and click **Add**; see Adding Object Overrides, on page 401.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Security Group Tag Objects

A Security Group Tag (SGT) object specifies a single SGT value. You can use SGT objects in rules to control traffic with SGT attributes that were **not** assigned by Cisco ISE. You cannot group or override SGT objects.

Related Topics

Autotransition from Custom SGTs to ISE SGTs, on page 389
Custom SGT Conditions, on page 389
ISE SGT vs Custom SGT Rule Conditions, on page 389

Creating Security Group Tag Objects

You can create these objects in the global domain only. To use the object on Classic devices, you must have the Control license. For Smart Licensed devices, any license will do.

Before you begin

- Disable ISE connections. You cannot create custom SGT objects if you use ISE as an identity source.

Procedure

Step 1 Click **Objects** > **Object Management**.
Step 2 Click **Security Group Tag** from the list of object types.
Step 3 Click **Add Security Group Tag**.
Step 4 Enter a **Name**.
Step 5 Optionally, enter a **Description**.
Step 6 In the **Tag** field, enter a single SGT.
Step 7 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Autotransition from Custom SGTs to ISE SGTs, on page 389

Custom SGT Conditions, on page 389
ISE SGT vs Custom SGT Rule Conditions, on page 389

URL Objects



Important For best practices for using this and similar options in Security Intelligence configurations and for URL rules in access control and QoS policies, see Manual URL Filtering Options, on page 1239.

A URL object defines a single URL or IP address, whereas a URL group object can define more than one URL or address. You can use URL objects and groups in various places in the system's web interface, including access control policies and event searches.

When creating URL objects, keep the following points in mind:

- If you do not include a path (that is, there is no / character in the URL), the match is based on the server's hostname only. The hostname is considered a match if it comes after the :// separator, or after any dot in the hostname. For example, ign.com matches ign.com and www.ign.com, but it does not match verisign.com.
- If you include one or more / character, the entire URL string is used for a substring match, including the server name, path, and any query parameters. However, we recommend that you do not use manual URL filtering to block or allow individual web pages or parts of sites, as servers can be reorganized and pages moved to new paths. Substring matching can also lead to unexpected matches, where the string you include in the URL object also matches paths on unintended servers or strings within query parameters.
- The system disregards the encryption protocol (HTTP vs HTTPS). In other words, if you block a website, both HTTP and HTTPS traffic to that website is blocked, unless you use an application condition to target a specific protocol. When creating a URL object, you do not need to specify the protocol when creating an object. For example, use example.com rather than http://example.com.
- If you plan to use a URL object to match HTTPS traffic in an access control rule, create the object using the subject common name in the public key certificate used to encrypt the traffic. Also, the system disregards subdomains within the subject common name, so do not include subdomain information. For example, use example.com rather than www.example.com.

  However, please understand that the subject common name in the certificate might be completely unrelated to a web site's domain name. For example, the subject common name in the certificate for youtube.com is *.google.com (this of course might change at any time). You will get more consistent results if you use the SSL Decryption policy to decrypt HTTPS traffic so that URL filtering rules work on decrypted traffic.



Note URL objects will not match HTTPS traffic if the browser resumes a TLS session because the certificate information is no longer available. Thus, even if you carefully configure the URL object, you might get inconsistent results for HTTPS connections.

Creating URL Objects

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **URL** from the list of object types.

Step 3 Choose **Add Object** from the **Add URL** drop-down list.

Step 4 Enter a **Name**.

In a multidomain deployment, object names must be unique within the domain hierarchy. The system may identify a conflict with the name of an object you cannot view in your current domain.

Step 5 Optionally, enter a **Description**.

Step 6 Enter the **URL** or IP address.

Step 7 Manage overrides for the object:

- If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 401.
- If you want to add override values to this object, expand the Override section and click **Add**; see Adding Object Overrides, on page 401.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Geolocation Objects

Each geolocation object you configure represents one or more countries or continents that the system has identified as the source or destination of traffic on your monitored network. You can use geolocation objects in various places in the system's web interface, including access control policies, SSL policies, and event searches. For example, you could write an access control rule that blocks traffic to or from certain countries.

To ensure that you are using up-to-date information to filter your network traffic, Cisco strongly recommends that you regularly update your Geolocation Database (GeoDB).

Creating Geolocation Objects

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Geolocation** from the list of object types.

Step 3 Click **Add Geolocation**.

Step 4 Enter a **Name**.

In a multidomain deployment, object names must be unique within the domain hierarchy. The system may identify a conflict with the name of an object you cannot view in your current domain.

Step 5 Check the check boxes for the countries and continents you want to include in your geolocation object. Checking a continent chooses all countries within that continent, as well as any countries that GeoDB updates may add under that continent in the future. Unchecking any country under a continent unchecks the continent. You can choose any combination of countries and continents.

Step 6 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Interface Objects: Interface Groups and Security Zones

Interface objects segment your network to help you manage and classify traffic flow. An interface object simply groups interfaces. These groups may span multiple devices; you can also configure multiple interface objects on a single device.

There are two types of interface objects:

- Security zones—An interface can belong to only one security zone.
- Interface groups—An interface can belong to multiple interface groups (and to one security zone).

  You can use interface groups in Firepower Threat Defense NAT policies, prefilter policies, and QoS policies.

Although tunnel zones are not interface objects, you can use them in place of security zones in certain configurations; see Tunnel Zones and Prefiltering, on page 1282.

All interfaces in an interface object must be of the same type: all inline, passive, switched, routed, or ASA FirePOWER. After you create an interface object, you cannot change the type of interfaces it contains.

The Interface Objects page of the object manager lists the security zones and interface groups configured on your managed devices. The page also displays the type of interfaces in each interface object, and you can expand each interface object to view which interfaces on which devices belong to each object.



Note Create inline sets before you add security zones for the interfaces in the inline set; otherwise security zones are removed and you must add them again.

Model-Specific Notes and Warnings

During initial configuration of a 7000 or 8000 Series device, the system creates security zones based on the detection mode you selected for the device. For example, the system creates a Passive zone in passive

deployments, while in inline deployments the system creates External and Internal zones. When you register the device to the Firepower Management Center, those security zones are added to the FMC.

Interface Objects and Multitenancy

In a multidomain deployment, you can create interface objects at any level. An interface object created in an ancestor domain can contain interfaces that reside on devices in different domains. In this situation, subdomain users viewing the ancestor interface object configuration in the object manager can see only the interfaces in their domain.

Unless restricted by role, subdomain users can view **and** edit interface objects created in ancestor domains. Subdomain users can add and delete interfaces from these interface objects. They cannot, however, delete or rename the interface objects. You can neither view nor edit interface objects created in descendant domains.

Creating Security Zone and Interface Group Objects



Tip You can create empty interface objects and add interfaces to them later. To add an interface, the interface must have a name. You can also create security zones (but not interface groups) while configuring interfaces in **Devices** > **Device Management**.

Before you begin

- Understand the usage requirements and restrictions for each type of interface object. See Interface Objects: Interface Groups and Security Zones, on page 410.
- Carefully determine the interface objects you need. You cannot change an existing security zone to an interface group or vice-versa; instead you must create a new interface object.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Interface** from the list of object types.

Step 3 Click **Add > Security Zone** or **Add > Interface Group**.

Step 4 Enter a **Name**.

Step 5 Choose an **Interface Type**.

Step 6 From the **Device** > **Interfaces** drop-down list, choose a device that contains interfaces you want to add.

Step 7 Choose one or more interfaces.

Step 8 Click **Add** to add the interfaces you chose, grouped by device.

Step 9 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Variable Sets

Variables represent values commonly used in intrusion rules to identify source and destination IP addresses and ports. You can also use variables in intrusion policies to represent IP addresses in rule suppressions, adaptive profile updates, and dynamic rule states.

Tip Preprocessor rules can trigger events regardless of the hosts defined by network variables used in intrusion rules.

You use variable sets to manage, customize, and group your variables. You can use the default variable set provided by the system or create your own custom sets. Within any set you can modify predefined default variables and add and modify user-defined variables.

Most of the shared object rules and standard text rules that the Firepower System provides use predefined default variables to define networks and port numbers. For example, the majority of the rules use the variable `$HOME_NET` to specify the protected network and the variable `$EXTERNAL_NET` to specify the unprotected (or outside) network. In addition, specialized rules often use other predefined variables. For example, rules that detect exploits against web servers use the `$HTTP_SERVERS` and `$HTTP_PORTS` variables.

Rules are more effective when variables more accurately reflect your network environment. At a minimum, you should modify default variables in the default set. By ensuring that a variable such as `$HOME_NET` correctly defines your network and `$HTTP_SERVERS` includes all web servers on your network, processing is optimized and all relevant systems are monitored for suspicious activity.

To use your variables, you link variable sets to intrusion policies associated with access control rules or with the default action of an access control policy. By default, the default variable set is linked to all intrusion policies used by access control policies.

Adding a variable to any set adds it to all sets; that is, each variable set is a collection of all variables currently configured on your system. Within any variable set, you can add user-defined variables and customize the value of any variable.

Initially, the Firepower System provides a single, default variable set comprised of predefined default values. Each variable in the default set is initially set to its default value, which for a predefined variable is the value set by the Cisco Talos Intelligence Group (Talos) and provided in rule updates.

Although you can leave predefined default variables configured to their default values, Cisco recommends that you modify a subset of predefined variables.

You could work with variables only in the default set, but in many cases you can benefit most by adding one or more custom sets, configuring different variable values in different sets, and perhaps even adding new variables.

When using multiple sets, it is important to remember that the *current value* of any variable in the default set determines the *default value* of the variable in all other sets.

When you select **Variable Sets** on the Object Manager page, the object manager lists the default variable set and any custom sets you created.

On a freshly installed system, the default variable set is comprised only of the default variables predefined by Cisco.

Each variable set includes the default variables provided by the system and all custom variables you have added from any variable set. Note that you can edit the default set, but you cannot rename or delete the default set.

In a multidomain deployment, the system generates a default variable set for each subdomain.

Caution Importing an access control or an intrusion policy overwrites existing default variables in the default variable set with the imported default variables. If your existing default variable set contains a custom variable not present in the imported default variable set, the unique variable is preserved.

Related Topics

Managing Variables, on page 424
Managing Variable Sets, on page 422

Variable Sets in Intrusion Policies

By default, the Firepower System links the default variable set to all intrusion policies used in an access control policy. When you deploy an access control policy that uses an intrusion policy, intrusion rules that you have enabled in the intrusion policy use the variable values in the linked variable set.

When you modify a custom variable set used by an intrusion policy in an access control policy, the system reflects the status for that policy as out-of-date on the Access Control Policy page. You must re-deploy the access control policy to implement changes in your variable set. When you modify the default set, the system reflects the status of all access control policies that use intrusion policies as out-of-date, and you must re-deploy all access control policies to implement your changes.

Variables

Variables belong to one of the following categories:

Default Variables

Variables provided by the Firepower System. You cannot rename or delete a default variable, and you cannot change its default value. However, you can create a customized version of a default variable.

Customized Variables

Variables you create. These variables can include:

- *customized default variables*

  When you edit the value for a default variable, the system moves the variable from the Default Variables area to the Customized Variables area. Because variable values in the default set determine the default values of variables in custom sets, customizing a default variable in the default set modifies the default value of the variable in all other sets.

- *user-defined variables*

  You can add and delete your own variables, customize their values within different variable sets, and reset customized variables to their default values. When you reset a user-defined variable, it remains in the Customized Variables area.

  User-defined variables can be one of the following types:

- *network* variables specify the IP addresses of hosts in your network traffic.
- *port* variables specify TCP or UDP ports in network traffic, including the value `any` for either type.

For example, if you create custom standard text rules, you might also want to add your own user-defined variables to more accurately reflect your traffic or as shortcuts to simplify the rule creation process. Alternatively, if you create a rule that you want to inspect traffic in the "demilitarized zone" (or DMZ) only, you can create a variable named `$DMZ` whose value lists the server IP addresses that are exposed. You can then use the `$DMZ` variable in any rule written for this zone.

Advanced Variables

Variables provided by the Firepower System under specific conditions. These variables have a very limited deployment.

Predefined Default Variables

By default, the Firepower System provides a single default variable set, which is comprised of predefined default variables. The Cisco Talos Intelligence Group (Talos) uses rule updates to provide new and updated intrusion rules and other intrusion policy elements, including default variables.

Because many intrusion rules provided by the system use predefined default variables, you should set appropriate values for these variables. Depending on how you use variable sets to identify traffic on your network, you can modify the values for these default variables in any or all variable sets.



Caution Importing an access control or an intrusion policy overwrites existing default variables in the default variable set with the imported default variables. If your existing default variable set contains a custom variable not present in the imported default variable set, the unique variable is preserved.

The following table describes the variables provided by the system and indicates which variables you typically would modify. For assistance determining how to tailor variables to your network, contact Professional Services or Support.

Table 56: System-Provided Variables

Variable Name	Description	Modify?
`$AIM_SERVERS`	Defines known AOL Instant Messenger (AIM) servers, and is used in chat-based rules and rules that look for AIM exploits.	Not required.
`$DNS_SERVERS`	Defines Domain Name Service (DNS) servers. If you create a rule that affects DNS servers specifically, you can use the `$DNS_SERVERS` variable as a destination or source IP address.	Not required in current rule set.
`$EXTERNAL_NET`	Defines the network that the Firepower System views as the unprotected network, and is used in many rules to define the external network.	Yes, you should adequately define `$HOME_NET` and then exclude `$HOME_NET` as the value for `$EXTERNAL_NET`.
`$FILE_DATA_PORTS`	Defines non-encrypted ports used in intrusion rules that detect files in a network stream.	Not required.

Variable Name	Description	Modify?
$FTP_PORTS	Defines the ports of FTP servers on your network, and is used for FTP server exploit rules.	Yes, if your FTP servers use ports other than the default ports (you can view the default ports in the web interface).
$GTP_PORTS	Defines the data channel ports where the packet decoder extracts the payload inside a GTP (General Packet Radio Service [GPRS] Tunneling Protocol) PDU.	Not required.
$HOME_NET	Defines the network that the associated intrusion policy monitors, and is used in many rules to define the internal network.	Yes, to include the IP addresses for your internal network.
$HTTP_PORTS	Defines the ports of web servers on your network, and is used for web server exploit rules.	Yes, if your web servers use ports other than the default ports (you can view the default ports in the web interface).
$HTTP_SERVERS	Defines the web servers on your network. Used in web server exploit rules.	Yes, if you run HTTP servers.
$ORACLE_PORTS	Defines Oracle database server ports on your network, and is used in rules that scan for attacks on Oracle databases.	Yes, if you run Oracle servers.
$SHELLCODE_PORTS	Defines the ports you want the system to scan for shell code exploits, and is used in rules that detect exploits that use shell code.	Not required.
$SIP_PORTS	Defines the ports of SIP servers on your network, and is used for SIP exploit rules.	Not required.
$SIP_SERVERS	Defines SIP servers on your network, and is used in rules that address SIP-targeted exploits.	Yes, if you run SIP servers, you should adequately define $HOME_NET and then include $HOME_NET as the value for $SIP_SERVERS.
$SMTP_SERVERS	Defines SMTP servers on your network, and is used in rules that address exploits that target mail servers.	Yes, if you run SMTP servers.
$SNMP_SERVERS	Defines SNMP servers on your network, and is used in rules that scan for attacks on SNMP servers.	Yes, if you run SNMP servers.
$SNORT_BPF	Identifies a legacy advanced variable that appears only when it existed on your system in a Firepower System software release before Version 5.3.0 that you subsequently upgraded to Version 5.3.0 or greater.	No, you can only view or delete this variable. You cannot edit it or recover it after deleting it.
$SQL_SERVERS	Defines database servers on your network, and is used in rules that address database-targeted exploits.	Yes, if you run SQL servers.
$SSH_PORTS	Defines the ports of SSH servers on your network, and is used for SSH server exploit rules.	Yes, if your SSH servers use ports other than the default port (you can view the default ports in the web interface).

Variable Name	Description	Modify?
`$SSH_SERVERS`	Defines SSH servers on your network, and is used in rules that address SSH-targeted exploits.	Yes, if you run SSH servers, you should adequately define `$HOME_NET` and then include `$HOME_NET` as the value for `$SSH_SERVERS`.
`$TELNET_SERVERS`	Defines known Telnet servers on your network, and is used in rules that address Telnet server-targeted exploits.	Yes, if you run Telnet servers.
`$USER_CONF`	Provides a general tool that allows you to configure one or more features not otherwise available via the web interface. Conflicting or duplicate `$USER_CONF` configurations will halt the system.	No, only as instructed in a feature description or with the guidance of Support.

Network Variables

Network variables represent IP addresses you can use in intrusion rules that you enable in an intrusion policy and in intrusion policy rule suppressions, dynamic rule states, and adaptive profile updates. Network variables differ from network objects and network object groups in that network variables are specific to intrusion policies and intrusion rules, whereas you can use network objects and groups to represent IP addresses in various places in the system's web interface, including access control policies, network variables, intrusion rules, network discovery rules, event searches, reports, and so on.

You can use network variables in the following configurations to specify the IP addresses of hosts on your network:

- intrusion rules—Intrusion rule **Source IPs** and **Destination IPs** header fields allow you to restrict packet inspection to the packets originating from or destined to specific IP addresses.
- suppressions—The **Network** field in source or destination intrusion rule suppressions allows you to suppress intrusion event notifications when a specific IP address or range of IP addresses triggers an intrusion rule or preprocessor.
- dynamic rule states—The **Network** field in source or destination dynamic rule states allows you to detect when too many matches for an intrusion rule or preprocessor rule occur in a given time period.
- adaptive profile updates—When you enable adaptive profile updates, the adaptive profiles **Networks** field identifies hosts where you want to improve reassembly of packet fragments and TCP streams in passive deployments.

When you use variables in the fields identified in this section, the variable set you link to an intrusion policy determines the variable values in the network traffic handled by an access control policy that uses the intrusion policy.

You can add any combination of the following network configurations to a variable:

- any combination of network variables, network objects, and network object groups that you select from the list of available networks
- individual network objects that you add from the New Variable or Edit Variable page, and can then add to your variable and to other existing and future variables
- literal, single IP addresses or address blocks

You can list multiple literal IP addresses and address blocks by adding each individually. You can list IPv4 and IPv6 addresses and address blocks alone or in any combination. When specifying IPv6 addresses, you can use any addressing convention defined in RFC 4291.

The default value for included networks in any variable you add is the word `any`, which indicates any IPv4 or IPv6 address. The default value for excluded networks is `none`, which indicates no network. You can also specify the address `::` in a literal value to indicate any IPv6 address in the list of included networks, or no IPv6 addresses in the list of exclusions.

Adding networks to the excluded list negates the specified addresses and address blocks. That is, you can match any IP address with the exception of the excluded IP address or address blocks.

For example, excluding the literal address `192.168.1.1` specifies any IP address other than 192.168.1.1, and excluding `2001:db8:ca2e::fa4c` specifies any IP address other than 2001:db8:ca2e::fa4c.

You can exclude any combination of networks using literal or available networks. For example, excluding the literal values `192.168.1.1` and `192.168.1.5` *includes* any IP address other than 192.168.1.1 or 192.168.1.5. That is, the system interprets this as "**not** 192.168.1.1 **and not** 192.168.1.5," which matches any IP address other than those listed between brackets.

Note the following points when adding or editing network variables:

- You cannot logically exclude the value `any` which, if excluded, would indicate no address. For example, you cannot add a variable with the value `any` to the list of excluded networks.
- Network variables identify traffic for the specified intrusion rule and intrusion policy features. Note that preprocessor rules can trigger events regardless of the hosts defined by network variables used in intrusion rules.
- Excluded values must resolve to a subset of included values. For example, you cannot include the address block 192.168.5.0/24 and exclude 192.168.6.0/24.

Port Variables

Port variables represent TCP and UDP ports you can use in the **Source Port** and **Destination Port** header fields in intrusion rules that you enable in an intrusion policy. Port variables differ from port objects and port object groups in that port variables are specific to intrusion rules. You can create port objects for protocols other than TCP and UDP, and you can use port objects in various places in the system's web interface, including port variables, access control policies, network discovery rules, and event searches.

You can use port variables in the intrusion rule **Source Port** and **Destination Port** header fields to restrict packet inspection to packets originating from or destined to specific TCP or UDP ports.

When you use variables in these fields, the variable set you link to the intrusion policy associated with an access control rule or policy determines the values for these variables in the network traffic where you deploy the access control policy.

You can add any combination of the following port configurations to a variable:

- any combination of port variables and port objects that you select from the list of available ports

  Note that the list of available ports does not display port object groups, and you cannot add these to variables.

- individual port objects that you add from the New Variable or Edit Variable page, and can then add to your variable and to other existing and future variables

Only TCP and UDP ports, including the value `any` for either type, are valid variable values. If you use the new or edit variables page to add a valid port object that is not a valid variable value, the object is added to the system but is not displayed in the list of available objects. When you use the object manager to edit a port object that is used in a variable, you can only change its value to a valid variable value.

- single, literal port values and port ranges

  You must separate port ranges with a dash (-). Port ranges indicated with a colon (:) are supported for backward compatibility, but you cannot use a colon in port variables that you create.

  You can list multiple literal port values and ranges by adding each individually in any combination.

Note the following points when adding or editing port variables:

- The default value for included ports in any variable you add is the word `any`, which indicates any port or port range. The default value for excluded ports is `none`, which indicates no ports.

Tip To create a variable with the value `any`, name and save the variable without adding a specific value.

- You cannot logically exclude the value `any` which, if excluded, would indicate no ports. For example, you cannot save a variable set when you add a variable with the value `any` to the list of excluded ports.
- Adding ports to the excluded list negates the specified ports and port ranges. That is, you can match any port with the exception of the excluded ports or port ranges.
- Excluded values must resolve to a subset of included values. For example, you cannot include the port range 10-50 and exclude port 60.

Advanced Variables

Advanced variables allow you to configure features that you cannot otherwise configure via the web interface. The Firepower System currently provides only only one advanced variable, the USER_CONF variable.

USER_CONF

USER_CONF provides a general tool that allows you to configure one or more features not otherwise available via the web interface.

Caution Do **not** use the advanced variable USER_CONF to configure an intrusion policy feature unless you are instructed to do so in the feature description or by Support. Conflicting or duplicate configurations will halt the system.

When editing USER_CONF, you can type up to 4096 total characters on a single line; the line wraps automatically. You can include any number of valid instructions or lines until you reach the 8192 maximum character length for a variable or a physical limit such as disk space. Use the backslash (\) line continuation character after any complete argument in a command directive.

Resetting USER_CONF empties it.

Variable Reset

You can reset a variable to its default value on the variable set new or edit variables page. The following table summarizes the basic principles of resetting variables.

Table 57: Variable Reset Values

Resetting this variable type...	In this set type...	Resets it to...
default	default	the rule update value
user-defined	default	`any`
default or user-defined	custom	the current default set value (modified or unmodified)

Resetting a variable in a custom set simply resets it to the current value for that variable in the default set.

Conversely, resetting or modifying the value of a variable in the default set always updates the default value of that variable in all custom sets. When the reset icon is grayed out, indicating that you cannot reset the variable, this means that the variable has no customized value in that set. Unless you have customized the value for a variable in a custom set, a change to the variable in the default set updates the value used in any intrusion policy where you have linked the variable set.

Note It is good practice when you modify a variable in the default set to assess how the change affects any intrusion policy that uses the variable in a linked custom set, especially when you have not customized the variable value in the custom set.

You can hover your pointer over the **Reset icon** in a variable set to see the reset value. When the customized value and the reset value are the same, this indicates one of the following:

- you are in the custom or default set where you added the variable with the value `any`
- you are in the custom set where you added the variable with an explicit value and elected to use the configured value as the default value

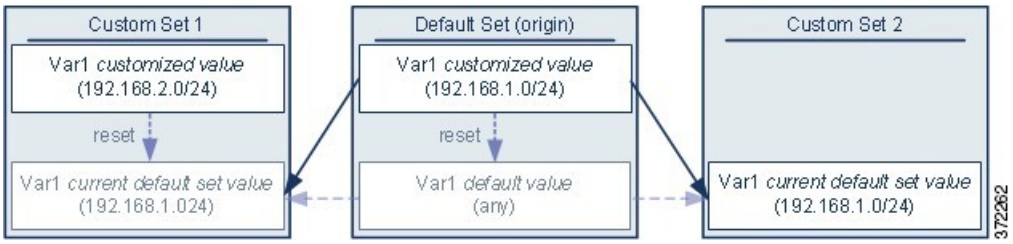
Adding Variables to Sets

Adding a variable to a variable set adds it to all other sets. When you add a variable from a custom set, you must choose whether to use the configured value as the customized value in the default set:

- **If you use the configured value** (for example, 192.168.0.0/16), the variable is added to the default set using the configured value as a customized value with a default value of `any`. Because the current value in the default set determines the default value in other sets, the initial, default value in other custom sets is the configured value (which in the example is 192.168.0.0/16).
- **If you do not use the configured value**, the variable is added to the default set using only the default value `any` and, consequently, the initial, default value in other custom sets is `any`.

Example: Adding User-Defined Variables to Default Sets

The following diagram illustrates set interactions when you add the user-defined variable `Var1` to the default set with the value `192.168.1.0/24`.



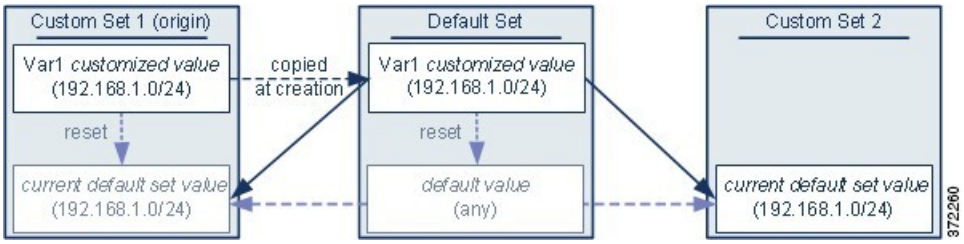
You can customize the value of `Var1` in any set. In Custom Set 2 where `Var1` has not been customized, its value is `192.168.1.0/24`. In Custom Set 1 the customized value `192.168.2.0/24` of `Var1` overrides the default value. Resetting a user-defined variable in the default set resets its default value to `any` in all sets.

It is important to note in this example that, if you do not update `Var1` in Custom Set 2, further customizing or resetting `Var1` in the default set consequently updates the current, default value of `Var1` in Custom Set 2, thereby affecting any intrusion policy linked to the variable set.

Although not shown in the example, note that interactions between sets are the same for user-defined variables and default variables except that resetting a default variable in the default set resets it to the value configured by Cisco in the current rule update.

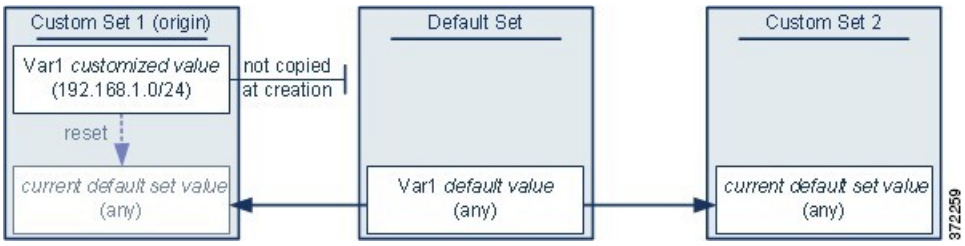
Example: Adding User-Defined Variables to Custom Sets

The next two examples illustrate variable set interactions when you add a user-defined variable to a custom set. When you save the new variable, you are prompted whether to use the configured value as the default value for other sets. In the following example, you elect **to use** the configured value.



Note that, except for the origin of `Var1` from Custom Set 1, this example is identical to the example above where you added `Var1` to the default set. Adding the customized value `192.168.1.0/24` for `Var1` to Custom Set 1 copies the value to the default set as a customized value with a default value of `any`. Thereafter, `Var1` values and interactions are the same as if you had added `Var1` to the default set. As with the previous example, keep in mind that further customizing or resetting `Var1` in the default set consequently updates the current, default value of `Var1` in Custom Set 2, thereby affecting any intrusion policy linked to the variable set.

In the next example, you add `Var1` with the value 192.168.1.0/24 to Custom Set 1 as in the previous example, but you elect **not to use** the configured value of `Var1` as the default value in other sets.



This approach adds `Var1` to all sets with a default value of `any`. After adding `Var1`, you can customize its value in any set. An advantage of this approach is that, by not initially customizing `Var1` in the default set, you decrease your risk of customizing the value in the default set and thus inadvertently changing the current value in a set such as Custom Set 2 where you have not customized `Var1`.

Nesting Variables

You can nest variables so long as the nesting is not circular. Nested, negated variables are not supported.

Valid Nested Variables

In this example, SMTP_SERVERS, HTTP_SERVERS, and OTHER_SERVERS are valid nested variables.

Variable	Type	Included Networks	Excluded Networks
SMTP_SERVERS	customized default	10.1.1.1	—
HTTP_SERVERS	customized default	10.1.1.2	—
OTHER_SERVERS	user-defined	10.2.2.0/24	—
HOME_NET	customized default	10.1.1.0/24 OTHER_SERVERS	SMTP_SERVERS HTTP_SERVERS

An Invalid Nested Variable

In this example, HOME_NET is an invalid nested variable because the nesting of HOME_NET is circular; that is, the definition of OTHER_SERVERS includes HOME_NET, so you would be nesting HOME_NET in itself.

Variable	Type	Included Networks	Excluded Networks
SMTP_SERVERS	customized default	10.1.1.1	—
HTTP_SERVERS	customized default	10.1.1.2	—
OTHER_SERVERS	user-defined	10.2.2.0/24 HOME_NET	—

Variable	Type	Included Networks	Excluded Networks
HOME_NET	customized default	10.1.1.0/24 OTHER_SERVERS	SMTP_SERVERS HTTP_SERVERS

An Unsupported Nested, Negated Variable

Because nested, negated variables are not supported, you cannot use the variable NONCORE_NET as shown in this example to represent IP addresses that are outside of your protected networks.

Variable	Type	Included Networks	Excluded Networks
HOME_NET	customized default	10.1.0.0/16 10.2.0.0/16 10.3.0.0/16	—
EXTERNAL_NET	customized default	—	HOME_NET
DMZ_NET	user-defined	10.4.0.0/16	—
NOT_DMZ_NET	user-defined	—	DMZ_NET
NONCORE_NET	user-defined	EXTERNAL_NET NOT_DMZ_NET	—

Alternative to an Unsupported Nested, Negated Variable

As an alternative to the example above, you could represent IP addresses that are outside of your protected networks by creating the variable NONCORE_NET as shown in this example.

Variable	Type	Included Networks	Excluded Networks
HOME_NET	customized default	10.1.0.0/16 10.2.0.0/16 10.3.0.0/16	—
DMZ_NET	user-defined	10.4.0.0/16	—
NONCORE_NET	user-defined	—	HOME_NET DMZ_NET

Managing Variable Sets

To use variable sets, you must have the Threat license (for FTD devices) or the Protection license (all other device types).

In a multidomain deployment, the system displays objects created in the current domain, which you can edit. It also displays objects created in ancestor domains, which in most cases you cannot edit. To view and edit objects in a descendant domain, switch to that domain.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Variable Set** from the list of object types.

Step 3 Manage your variable sets:

- Add — If you want to add a custom variable set, click **Add Variable Set**; see Creating Variable Sets, on page 423.
- Delete — If you want to delete a custom variable set, click **Delete** ( ) next to the variable set, then click **Yes**. You cannot delete the default variable set or variable sets belonging to ancestor domains.

  Note Variables created in a variable set you delete are not deleted or otherwise affected in other sets.

- Edit — If you want to edit a variable set, click **Edit** ( ) next to the variable set you want to modify; see Editing Objects, on page 396.
- Filter — If you want to filter variable sets by name, begin entering a name; as you type, the page refreshes to display matching names. If you want to clear name filtering, click **Clear** ( ) in the filter field.
- Manage Variables — To manage the variables included in variable sets, see Managing Variables, on page 424.

Creating Variable Sets

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Variable Set** from the list of object types.

Step 3 Click **Add Variable Set**.

Step 4 Enter a **Name**.

In a multidomain deployment, object names must be unique within the domain hierarchy. The system may identify a conflict with the name of an object you cannot view in your current domain.

Step 5 Optionally, enter a **Description**.

Step 6 Manage the variables in the set; see Managing Variables, on page 424.

Step 7 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Managing Variables

You must have the Threat license (for FTD devices) or the Protection license (all other device types).

In a multidomain deployment, the system displays objects created in the current domain, which you can edit. It also displays objects created in ancestor domains, which in most cases you cannot edit. To view and edit objects in a descendant domain, switch to that domain.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Variable Set** from the list of object types.

Step 3 Click **Edit** ( ) next to the variable set you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Manage your variables:

- Display — If you want to display the complete value for a variable, hover your pointer over the value in the **Value** column next to the variable.
- Add — If you want to add a variable, click **Add**; see Adding Variables, on page 425.
- Delete — Click **Delete** ( ) next to the variable. If you have saved the variable set since adding the variable, click **Yes** to confirm that you want to delete the variable.

  You *cannot* delete the following:

  - default variables
  - user-defined variables that are used by intrusion rules or other variables
  - variables belonging to ancestor domains

- Edit — Click **Edit** ( ) next to the variable you want to edit; see Editing Variables, on page 426.
- Reset — If you want to reset a modified variable to its default value, click **Reset** next to a modified variable. If reset is dimmed, one of the following is true:

  - The current value is already the default value.
  - The configuration belongs to an ancestor domain.

  Tip Hover your pointer over an active reset to display the default value.

Step 5 Click **Save** to save the variable set. If the variable set is in use by an access control policy, click **Yes** to confirm that you want to save your changes.

Because the current value in the default set determines the default value in all other sets, modifying or resetting a variable in the default set changes the current value in other sets where you have not customized the default value.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Adding Variables

You must have the Threat license (for FTD devices) or the Protection license (all other device types).

Procedure

Step 1 In the variable set editor, click **Add**.

Step 2 Enter a unique variable **Name**.

Step 3 From the **Type** drop-down list, choose either **Network** or **Port**.

Step 4 Specify values for the variable:

- If you want to move items from the list of available networks or ports to the list of included or excluded items, you can choose one or more items and then drag and drop, or click **Include** or **Exclude**.

  Tip If addresses or ports in the included and excluded lists for a network or port variable overlap, excluded addresses or ports take precedence.

- Enter a single literal value, then click **Add**. For network variables, you can enter a single IP address or address block. For port variables you can add a single port or port range, separating the upper and lower values with a hyphen (-). Repeat this step as needed to enter multiple literal values.
- If you want to remove an item from the included or excluded lists, click **Delete** ( ) next to the item.

Note The list of items to include or exclude can be comprised of any combination of literal strings and existing variables, objects, and network object groups in the case of network variables.

Step 5 Click **Save** to save the variable. If you are adding a new variable from a custom set, you have the following options:

- Click **Yes** to add the variable using the configured value as the customized value in the default set and, consequently, the default value in other custom sets.
- Click **No** to add the variable as the default value of `any` in the default set and, consequently, in other custom sets.

Step 6 Click **Save** to save the variable set. Your changes are saved, and any access control policy the variable set is linked to displays an out-of-date status.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Editing Variables

You must have the Threat license (for FTD devices) or the Protection license (all other device types).

In a multidomain deployment, the system displays objects created in the current domain, which you can edit. It also displays objects created in ancestor domains, which in most cases you cannot edit. To view and edit objects in a descendant domain, switch to that domain.

You can edit both custom and default variables.

You cannot change the **Name** or **Type** values in an existing variable.

Procedure

Step 1 In the variable set editor, click **Edit** ( ) next to the variable you want to modify.

If **View** ( ) appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 2 Modify the variable:

- If you want to move items from the list of available networks or ports to the list of included or excluded items, you can select one or more items and then drag and drop, or click **Include** or **Exclude**.

  Tip If addresses or ports in the included and excluded lists for a network or port variable overlap, excluded addresses or ports take precedence.

- Enter a single literal value, then click **Add**. For network variables, you can enter a single IP address or address block. For port variables you can add a single port or port range, separating the upper and lower values with a hyphen (-). Repeat this step as needed to enter multiple literal values.
- If you want to remove an item from the included or excluded lists, click **Delete** ( ) next to the item.

Note The list of items to include or exclude can be comprised of any combination of literal strings and existing variables, objects, and network object groups in the case of network variables.

Step 3 Click **Save** to save the variable.

Step 4 Click **Save** to save the variable set. If the variable set is in use by an access control policy, click **Yes** to confirm that you want to save your changes. Your changes are saved, and any access control policy the variable set is linked to displays an out-of-date status.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Security Intelligence Lists and Feeds

Security Intelligence functionality requires the Threat license (for FTD devices) or the Protection license (all other device types).

Security Intelligence *lists* and *feeds* are collections of IP addresses, domain names, and URLs that you can use to quickly filter traffic that matches an entry on a list or feed.

- A list is a static collection that you manage manually.
- A feed is a dynamic collection that updates on an interval over HTTP or HTTPS.

Security Intelligence lists/feeds are grouped into:

- DNS (Domain names)
- Network (IP addresses)
- URLs

System-Provided Feeds

Cisco provides the following feeds as Security Intelligence objects:

- Security Intelligence feeds updated regularly with the latest threat intelligence from Talos:
  - Cisco-DNS-and-URL-Intelligence-Feed (under DNS Lists and Feeds)
  - Cisco-Intelligence-Feed (for IP addresses, under Network Lists and Feeds)

  You cannot delete the system-provided feeds, but you can change the frequency of (or disable) their updates.

Predefined Lists: Global Block Lists and Global Do Not Block Lists

The system ships with predefined global Block lists and Do Not Block lists for domains (DNS), IP addresses (Networks), and URLs.

These lists are empty until you populate them. To build these lists, see Global and Domain Security Intelligence Lists, on page 428.

By default, access control and DNS policies use these lists as part of Security Intelligence.

Custom Feeds

You can use third-party feeds, or use a custom internal feed to easily maintain an enterprise-wide Block list in a large deployment with multiple Firepower Management Center appliances.

See Custom Security Intelligence Feeds, on page 434.

Custom Lists

Custom lists can augment and fine-tune feeds and the Global lists.

See Custom Security Intelligence Lists, on page 436.

Where Security Intelligence Lists and Feeds Are Used

- IP address and address blocks—Use Block and Do Not Block lists in access control policies, as part of Security Intelligence.
- Domain Names—Use Block and Do Not Block lists in DNS policies, as part of Security Intelligence.
- URLs—Use Block and Do Not Block lists in access control and QoS rules, whose analysis and traffic handling phases occur after Security Intelligence.

How to Modify Security Intelligence Objects

To add or delete entries on a Block list, Do Not Block list, feed, or sinkhole object:

Object Type	Edit Capabilities	Requires Redeploy After Edit?
Custom Block and Do Not Block lists	Upload new and replacement lists using the object manager.	Yes
Default (but custom-populated) Block lists and Do Not Block lists: Global, descendant, and domain-specific	Add entries using the context menu or delete entries using the object manager.	No
System-provided Intelligence Feeds	Disable or change update frequency using the object manager.	No
Custom feeds	Fully modify using the object manager.	No
Sinkhole	Fully modify using the object manager.	Yes

Global and Domain Security Intelligence Lists

Firepower Management Center ships with empty Global Block and Do-Not-Block lists to which you can instantly add URLs, domains, and IP addresses from events on your network at any time. These lists allow you to use Security Intelligence to always block particular connections, or to exempt particular connections from blocking by Security Intelligence, allowing them to be evaluated by other threat detection processes that you have configured.

For example, if you notice a set of routable IP addresses in intrusion events associated with exploit attempts, you can immediately block those IP addresses. Although it may take a few minutes for your changes to propagate, you do not have to redeploy.

By default, Access control and DNS policies use these Global lists, which apply to all security zones. You can opt not to use these lists on a per-policy basis.



Note These options apply to Security Intelligence only. Security Intelligence cannot block traffic that has already been fastpathed. Similarly, adding an item to a Security Intelligence Do Not Block list does not automatically trust or fastpath matching traffic. For more information, see About Security Intelligence, on page 1249.

In a multidomain deployment, you can choose the Firepower System domains where you want to enforce blocking, or exempting from Security Intelligence blocking, by adding items to Domain lists as well as the Global lists; see Security Intelligence Lists and Multitenancy, on page 429.

Security Intelligence Lists and Multitenancy

In a multidomain deployment, the Global domain owns the Global Block lists and Do Not Block lists. Only Global administrators can add to or remove items from the Global lists. So that subdomain users can add networks, domain names, and URLs to Block and Do Not Block lists, multitenancy adds:

- Domain lists—Block or Do Not Block lists whose contents apply to a particular subdomain only. The Global lists are Domain lists for the Global domain.
- Descendant Domain lists—Block or Do Not Block lists that aggregate the Domain lists of the current domain's descendants.

Domain Lists

In addition to being able to access (but not edit) the Global lists, each subdomain has its own named lists, the contents of which apply only to that subdomain. For example, a subdomain named Company A owns:

- Domain Block list - Company A and Domain Do Not Block list - Company A
- Domain Block list for DNS - Company A, Domain Do Not Block list for DNS - Company A
- Domain Block list for URL - Company A, Domain Do Not Block list for URL - Company A

Any administrator at or above the current domain can populate these lists. You can use the context menu to add an item to the Block or Do Not Block list in the current and all descendant domains. However, only an administrator in the associated domain can remove an item from a Domain list.

For example, a Global administrator could choose to add the same IP address to the Block list in the Global domain and Company A's domain, but not add it to the Block list in Company B's domain. This action would add the same IP address to:

- Global Block list (where it can be removed only by Global administrators)
- Domain Block list - Company A (where it can be removed only by Company A administrators)

The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.

Descendant Domain Lists

A Descendant Domain list is a Do Not Block list or Block list that aggregates the Domain lists of the current domain's descendants. Leaf domains do not have Descendant Domain lists.

Descendant Domain lists are useful because a higher-level domain administrator can enforce general Security Intelligence settings, while still allowing subdomain users to add items to a Block or Do Not Block list in their own deployment.

For example, the Global domain has the following Descendant Domain lists:

- Descendant Block lists - Global, Descendant Do Not Block lists - Global
- Descendant Block lists for DNS - Global, Descendant Do Not Block lists for DNS - Global
- Descendant Block lists for URL - Global, Descendant Do Not Block lists for URL - Global



Note Descendant Domain lists do not appear in the object manager because they are symbolic aggregations, not hand-populated lists. They appear where you can use them: in access control and DNS policies.

Add Entries to Global Security Intelligence Lists

When reviewing events and dashboards, you can instantly block future traffic involving IP addresses, domains, and URLs that appear in those events by adding them to a predefined Block list.

Similarly, if Security Intelligence is blocking traffic that you want evaluated by threat detection processes subsequent to Security Intelligence blocking, you can add IP addresses, domains, and URLs from events to a predefined Do Not Block list.

Traffic is evaluated against entries on these lists during the Security Intelligence phase of threat detection.

For more information about these lists, see Global and Domain Security Intelligence Lists, on page 428.

Before you begin

Because adding an entry to a Security Intelligence list affects access control, you must have one of the following user roles:

- Administrator
- A combination of roles: Network Admin or Access Admin, plus Security Analyst and Security Approver
- A custom role with both Modify Access Control Policy and Deploy Configuration to Devices permissions

If appropriate, verify that these lists are used in the policies in which you expect them to be used.

Procedure

Step 1 Navigate to an event that includes an IP address, domain, or URL that you want to always block using Security Intelligence, or exempt from Security Intelligence blocking.

Step 2 Right-click the IP address, domain, or URL and choose the appropriate option:

Target Item	Context Menu Option	Affected Global Lists
An IP address	Blacklist Now	Global Block List
	Whitelist Now	Global Whitelist

Target Item	Context Menu Option	Affected Global Lists
A URL	Blacklist HTTP/S Connections to URL Now Whitelist HTTP/S Connections to URL Now	Global Block List for URL Global Whitelist for URL
An entire domain	Blacklist HTTP/S Connections to Domain Now Whitelist HTTP/S Connections to Domain Now	Global Block List for URL Global Whitelist for URL
DNS requests for an entire domain	Blacklist DNS Requests to Domain Now Whitelist DNS Requests to Domain Now	Global Block List for DNS Global Whitelist for DNS

What to do next

You do NOT need to redeploy for these changes to take effect.

If you want to delete an item from a list, see Delete Entries from Global Security Intelligence Lists, on page 431.

Delete Entries from Global Security Intelligence Lists

Note

- In multi-domain deployments, the names of these lists may not be "Global." For more information, see Security Intelligence Lists and Multitenancy, on page 429.
- To add entries to these lists, see Add Entries to Global Security Intelligence Lists, on page 430.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Click **Security Intelligence**.

Step 3 Click the appropriate option:

- **Network Lists and Feeds** (for IP addresses)
- **DNS Lists and Feeds** (for domain names)
- **URL Lists and Feeds**

Step 4 Click the pencil beside the Global Block or Global Do-Not-Block list.

Step 5 Click the trash button beside the entry to delete.

List and Feed Updates for Security Intelligence

List and feed updates replace the existing list or feed file with the contents of the new file. Contents of existing and new files are not merged.

If the system downloads a corrupt feed or a feed with no recognizable entries, the system continues using the old feed data (unless it is the first download). However, if the system can recognize even one entry in the feed, it uses the entries it can recognize.

By default, each feed updates the Management Center every two hours; you can modify this frequency. Any updates the Management Center receives are passed immediately to managed devices. In addition, managed devices poll the FMC every 30 minutes for changes. You cannot modify this frequency.

In a multidomain deployment, the system-provided feeds belong to the Global domain and can be modified only by an administrator in that domain. You can modify the update frequency for custom feeds belonging to your domain.

To modify feed update intervals, see Changing the Update Frequency for Security Intelligence Feeds, on page 432.

Changing the Update Frequency for Security Intelligence Feeds

You can specify the intervals at which the Firepower Management Center updates Security Intelligence Feeds.

For details about feed updates, see List and Feed Updates for Security Intelligence, on page 432.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **Security Intelligence** node, then choose the feed type whose frequency you want to change.

The system-provided URL feed is combined with the domain feed under **DNS Lists and Feeds**.

Step 3 Next to the feed you want to update, click **Edit** ( ).

If **View** ( ) appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 4 Edit the **Update Frequency**.

Step 5 Click **Save**.

Custom Security Intelligence Lists and Feeds

Custom Lists and Feeds: Requirements

List and Feed Formatting

Each list or feed must be a simple text file no larger than 500MB. List files must have the .txt extension. Include one entry or comment per line: one IP address, one URL, one domain name.

Tip The number of entries you can include is limited by the maximum size of the file. For example, a URL list with no comments and an average URL length of 100 characters (including Punycode or percent Unicode representations and newlines) can contain more than 5.24 million entries.

In a DNS list entry, you can specify an asterisk (`*`) wildcard character for a domain label. All labels match the wildcard. For example, an entry of `www.example.*` matches both `www.example.com` and `www.example.co`.

If you add comment lines within the source file, they must start with the pound (`#`) character. If you upload a source file with comments, the system removes your comments during upload. Source files you download contain all your entries without your comments.

Feed Requirements

When you configure a feed, you specify its location using a URL; the URL cannot be Punycode-encoded.

If you use an MD5 checksum, the checksum must be stored in a simple text file with only the checksum. Comments are not supported.

URL Lists and Feeds: URL Syntax and Matching Criteria

Security Intelligence URL lists and feeds, including custom lists and feeds and entries in the global Block list and Do Not Block list, can include the following, which have the matching behavior as described:

- Hostnames

  For example, **`www.example.com`**.

- URLs

  `example.com` matches **`example.com`** and all subdomains, including **`www.example.com`**, **`eu.example.com`**, **`example.com/abc`**, and **`www.example.com/def`** -- but NOT **`example.co.uk`** or **`examplexyz.com`** or **`example.com.malicious-site.com`**

  You can also include an entire URL path, such as **`https://www.cisco.com/c/en/us/products/security/firewalls/index.html`**

- A slash at the end of a URL to specify an exact match

  `example.com/` matches ONLY **`example.com`**; it does NOT match **`www.example.com`** or any other URL.

- A wildcard (*) to represent any domain in a URL

  An asterisk can represent a complete domain string separated by dots, but not a partial domain string, and not any part of the URL following the first slash.

  Valid examples:

  - **`*.example.com`**
  - **`www.*.com`**
  - **`example.*`**

    (This will match **`example.com`** and **`example.org`** and **`example.de`**, for example, but NOT **`example.co.uk`**)

- `*.example.*`
- `example.*/`

Invalid examples:

- `example*.com`
- `example.com/*`

- IP addresses (IPv4)

  For IPv6 addresses, or to use ranges or CIDR notation, use the Security Intelligence Network object.

  You can include one or more wildcards representing an octet, for example 10.10.10.* or 10.10.*.*.

See also Custom Security Intelligence Lists, on page 436.

Custom Security Intelligence Feeds

Custom or third-party Security Intelligence feeds allow you to augment the system-provided Intelligence Feeds with other regularly-updated reputable Block lists and Do Not Block lists on the Internet. You can also set up an internal feed, which is useful if you want to update multiple Firepower Management Center appliances in your deployment using one source list.

Note You cannot add address blocks to Block or Do Not Block lists using a `/0` netmask in a Security Intelligence feed. If you want to monitor or block all traffic targeted by a policy, use an access control rule with the **Monitor** or **Block** rule action, respectively, and a default value of `any` for the **Source Networks** and **Destination Networks**.

You also can configure the system to use an MD5 checksum to determine whether to download an updated feed. If the checksum has not changed since the last time the system downloaded the feed, the system does not need to re-download it. You may want to use MD5 checksums for internal feeds, especially if they are large.

Note The system does **not** perform peer SSL certificate verification when downloading custom feeds, nor does the system support the use of certificate bundles or self-signed certificates to verify the remote peer.

If you want strict control over when the system updates a feed from the Internet, you can disable automatic updates for that feed. However, automatic updates ensure the most up-to-date, relevant data.

Manually updating Security Intelligence feeds updates all feeds, including the Intelligence Feeds.

See complete requirements at Custom Lists and Feeds: Requirements, on page 432.

Creating Security Intelligence Feeds

You must have the Threat license (for FTD devices) or the Protection license (all other device types).

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **Security Intelligence** node, then choose a feed type you want to add.

Step 3 Click the option appropriate to the feed type you chose above:

- **Add Network Lists and Feeds** (for IP addresses)
- **Add DNS Lists and Feeds**
- **Add URL Lists and Feeds**

Step 4 Enter a **Name** for the feed.

In a multidomain deployment, object names must be unique within the domain hierarchy. The system may identify a conflict with the name of an object you cannot view in your current domain.

Step 5 Choose **Feed** from the **Type** drop-down list.

Step 6 Enter a **Feed URL**.

Step 7 (Optional) Enter an **MD5 URL**.

This is used to determine whether the feed contents have changed since the last update, so the system does not download unchanged feeds.

Step 8 Choose an **Update Frequency**.

Step 9 Click **Save**.
Unless you disabled feed updates, the system attempts to download and verify the feed.

Manually Updating Security Intelligence Feeds

You must have the Threat license (for FTD devices) or the Protection license (all other device types).

Before you begin

At least one device must already be added to the management center.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **Security Intelligence** node, then choose a feed type.

Step 3 Click **Update Feeds**, then confirm.

Step 4 Click **OK**.

After the Firepower Management Center downloads and verifies the feed updates, it communicates any changes to its managed devices. Your deployment begins filtering traffic using the updated feeds.

Custom Security Intelligence Lists

Security Intelligence lists are simple static lists of IP addresses and address blocks, URLs, or domain names that you manually upload to the system. Custom lists are useful if you want to augment and fine-tune feeds or one of the global lists, for a single Firepower Management Center's managed devices.

For example, if a reputable feed improperly blocks your access to vital resources but is overall useful to your organization, you can create a custom Do Not Block list that contains only the improperly classified IP addresses, rather than removing the IP address feed object from the access control policy's Block list.



Note You cannot add address blocks to a Block or Do Not Block list using a `/0` netmask in a Security Intelligence list. If you want to monitor or block all traffic targeted by a policy, use an access control rule with the **Monitor** or **Block** rule action, respectively, and a default value of `any` for the **Source Networks** and **Destination Networks**.

Regarding list entry formatting, note the following:

- Netmasks for address blocks can be integers from `0` to `32` or `0` to `128`, for IPv4 and IPv6, respectively.
- Unicode in domain names must be encoded in Punycode format, and are case insensitive.
- Characters in domain names are case-insensitive.
- Unicode in URLs should be encoded in percent-encoding format.
- Characters in URL subdirectories are case-sensitive.
- List entries that start with the pound sign (`#`) are treated as comments.
- See additional formatting requirements at Custom Lists and Feeds: Requirements, on page 432.

Regarding matching list entries, note the following:

- The system matches sub-level domains if a higher-level domain exists in a URL or DNS list. For example, if you add `example.com` to a DNS list, the system matches both `www.example.com` and `test.example.com`.
- The system does not perform DNS lookups (forward or reverse) on DNS or URL list entries. For example, if you add `http://192.168.0.2` to a URL list, and it resolves to `http://www.example.com`, the system only matches `http://192.168.0.2`, not `http://www.example.com`.

Uploading New Security Intelligence Lists to the Firepower Management Center

To modify a Security Intelligence list, you must make your changes to the source file and upload a new copy. You cannot modify the file's contents using the web interface. If you do not have access to the source file, download a copy from the system.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **Security Intelligence** node, then choose a list type.

Step 3 Click the option appropriate to the list you chose above:

- **Add Network Lists and Feeds** (for IP addresses)
- **Add DNS Lists and Feeds**
- **Add URL Lists and Feeds**

Step 4 Enter a **Name**.

In a multidomain deployment, object names must be unique within the domain hierarchy. The system may identify a conflict with the name of an object you cannot view in your current domain.

Step 5 From the **Type** drop-down list, choose **List**.

Step 6 Click **Browse** to browse to the list `.txt` file, then click **Upload**.

Step 7 Click **Save**.

What to do next

If an active policy references your object, deploy configuration changes, see Deploy Configuration Changes, on page 346

Updating Security Intelligence Lists

In a multidomain deployment, the system displays objects created in the current domain, which you can edit. It also displays objects created in ancestor domains, which in most cases you cannot edit. To view and edit objects in a descendant domain, switch to that domain.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **Security Intelligence** node, then choose a list type.

Step 3 Next to the list you want to update, click **Edit** ( ).

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 If you need a copy of the list to edit, click **Download**, then follow your browser's prompts to save the list as a text file.

Step 5 Make changes to the list as necessary.

Step 6 On the Security Intelligence pop-up window, click **Browse** to browse to the modified list, then click **Upload**.

Step 7 Click **Save**.

What to do next

If an active policy references your object, deploy configuration changes, see Deploy Configuration Changes, on page 346.

Sinkhole Objects

A sinkhole object represents either a DNS server that gives non-routeable addresses for all domain names within the sinkhole, or an IP address that does not resolve to a server. You can reference the sinkhole object within a DNS policy rule to redirect matching traffic to the sinkhole. You must assign the object both an IPv4 address and an IPv6 address.

Creating Sinkhole Objects

You must have the Threat license (for FTD devices) or the Protection license (all other device types).

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Sinkhole** from the list of object types.

Step 3 Click **Add Sinkhole**.

Step 4 Enter a **Name**.

In a multidomain deployment, object names must be unique within the domain hierarchy. The system may identify a conflict with the name of an object you cannot view in your current domain.

Step 5 Enter the **IPv4 Address** and **IPv6 Address** of your sinkhole.

Step 6 You have the following options:

- If you want to redirect traffic to a sinkhole server, choose **Log Connections to Sinkhole**.
- If you want to redirect traffic to a non-resolving IP address, choose **Block and Log Connections to Sinkhole**.

Step 7 If you want to assign an Indication of Compromise (IoC) type to your sinkhole, choose one from the **Type** drop-down.

Step 8 Click **Save**.

File Lists

If you use AMP for Networks, and the AMP cloud incorrectly identifies a file's disposition, you can add the file to a *file list* to better detect the file in the future. These files are specified using SHA-256 hash values. Each file list can contain up to 10000 unique SHA-256 values.

There are two predefined categories of file lists:

Clean List

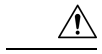
If you add a file to this list, the system treats it as if the AMP cloud assigned a clean disposition.

Custom Detection List

If you add a file to this list, the system treats it as if the AMP cloud assigned a malware disposition.

In a multidomain deployment, a clean list and custom detection list is present for each domain. In lower-level domains, you can view but not modify ancestor's lists.

Because you manually specify the blocking behavior for the files included in these lists, the system does not query the AMP cloud for these files' dispositions. You must configure a rule in the file policy with either a **Malware Cloud Lookup** or **Block Malware** action and a matching file type to calculate a file's SHA value.



Caution Do **not** include malware on the clean list. The clean list overrides both the AMP cloud and the custom detection list.

Source Files for File Lists

You can add multiple SHA-256 values to a file list by uploading a comma-separated value (CSV) source file containing a list of SHA-256 values and descriptions. The Firepower Management Center validates the contents and populates the file list with valid SHA-256 values.

The source file must be a simple text file with a .csv file name extension. Any header must start with a pound sign (`#`); it is treated as a comment and not uploaded. Each entry should contain a single SHA-256 value followed by a description and end with either the `LF` or `CR+LF` Newline character. The system ignores any additional information in the entry.

Note the following:

- Deleting a source file from the file list also removes all associated SHA-256 hashes from the file list.
- You cannot upload multiple files to a file list if the successful source file upload results in the file list containing more than 10000 distinct SHA-256 values.
- The system truncates descriptions exceeding 256 characters to the first 256 characters on upload. If the description contains commas, you must use an escape character (`\,`). If no description is included, the source file name is used instead.
- All non-duplicate SHA-256 values are added to the file list. If a file list contains a SHA-256 value, and you upload a source file containing that value, the newly uploaded value does not modify the existing SHA-256 value. When viewing captured files, file events, or malware events related to the SHA-256 value, any threat name or description is derived from the individual SHA-256 value.
- The system does not upload invalid SHA-256 values in a source file.
- If multiple uploaded source files contain an entry for the same SHA-256 value, the system uses the most recent value.
- If a source file contains multiple entries for the same SHA-256 value, the system uses the last one.
- You cannot directly edit a source file within the object manager. To make changes, you must first modify your source file directly, delete the copy on the system, then upload the modified source file.
- The number of entries associated with a source file refers to the number of distinct SHA-256 values. If you delete a source file from a file list, the total number of SHA-256 entries the file list contains decreases by the number of valid entries in the source file.

Adding Individual SHA-256 Values to File Lists

You must have the Malware license for this procedure.

You can submit a file's SHA-256 value to add it to a file list. You cannot add duplicate SHA-256 values.

In a multidomain deployment, the system displays objects created in the current domain, which you can edit. It also displays objects created in ancestor domains, which in most cases you cannot edit. To view and edit objects in a descendant domain, switch to that domain.

Before you begin

- Right-click a file or malware event from the event view, choose **Show Full Text** in the context menu, and copy the full SHA-256 value for pasting into the file list.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **File List** from the list of object types.

Step 3 Click **Edit** ( ) next to the clean list or custom detection list where you want to add a file.

If **View** ( ) appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 4 Choose `Enter SHA Value` from the **Add by** drop-down list.

Step 5 Enter a description of the source file in the **Description** field.

Step 6 Enter or paste the file's entire value in the **SHA-256** field. The system does not support matching partial values.

Step 7 Click **Add**.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Note After configuration changes are deployed, the system no longer queries the AMP cloud for files on the list.

Uploading Individual Files to File Lists

You must have the Malware license for this procedure.

If you have a copy of the file you want to add to a file list, you can upload the file to the Firepower Management Center for analysis; the system calculates the file's SHA-256 value and adds the file to the list. The system does not enforce a limit on the size of files for SHA-256 calculation.

In a multidomain deployment, the system displays objects created in the current domain, which you can edit. It also displays objects created in ancestor domains, which in most cases you cannot edit. To view and edit objects in a descendant domain, switch to that domain.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **File List** from the list of object types.

Step 3 Click **Edit** ( ) next to the clean list or custom detection list where you want to add a file.

If **View** ( ) appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 4 From the **Add by** drop-down list, choose **Calculate SHA**.

Step 5 Optionally, enter a description of the file in the **Description** field. If you do not enter a description, the file name is used for the description on upload.

Step 6 Click **Browse**, and choose a file to upload.

Step 7 Click **Calculate and Add SHA**.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Note After you deploy configuration changes, the system no longer queries the AMP cloud for files on the list.

Uploading Source Files to File Lists

You must have the Malware license for this procedure.

In a multidomain deployment, the system displays objects created in the current domain, which you can edit. It also displays objects created in ancestor domains, which in most cases you cannot edit. To view and edit objects in a descendant domain, switch to that domain.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Click **File List**.

Step 3 Click **Edit** ( ) next to the file list where you want to add values from a source file.

If **View** ( ) appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 4 In the **Add by** drop-down list, choose `List of SHAs`.

Step 5 Optionally, enter a description of the source file in the **Description** field. If you do not enter a description, the system uses the file name.

Step 6 Click **Browse** to browse to the source file, then click **Upload and Add List**.

Step 7 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Note After you deploy the policies, the system no longer queries the AMP cloud for files on the list.

Editing SHA-256 Values in File Lists

You must have the Malware license for this procedure.

You can edit or delete individual SHA-256 values on a file list. Note that you cannot directly edit a source file within the object manager. To make changes, you must first modify your source file directly, delete the copy on the system, then upload the modified source file.

In a multidomain deployment, the system displays objects created in the current domain, which you can edit. It also displays objects created in ancestor domains, which in most cases you cannot edit. To view and edit objects in a descendant domain, switch to that domain.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Click **File List**.

Step 3 Click **Edit** ( ) next to the clean list or custom detection list where you want to modify a file.

If **View** ( ) appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 4 You can:

- Click **Edit** ( ) next to the SHA-256 value you want to change, and modify the **SHA-256** or **Description** values as desired.
- Click **Delete** ( ) next to the SHA-256 value you want to delete.

Step 5 Click **Save** to update the file entry in the list.

Step 6 Click **Save** to save the file list.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Note After configuration changes are deployed, the system no longer queries the AMP cloud for files on the list.

Downloading Source Files from File Lists

You must have the Malware license for this procedure.

In a multidomain deployment, the system displays objects created in the current domain, which you can edit. It also displays objects created in ancestor domains, which in most cases you cannot edit. To view and edit objects in a descendant domain, switch to that domain.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **File List** from the list of object types.

Step 3 Click **Edit** ( ) next to the clean list or custom detection list where you want to download a source file.

If **View** ( ) appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 4 Next to the source file you want to download, click **View** ( ).

Step 5 Click **Download SHA List** and follow the prompts to save the source file.

Step 6 Click **Close**.

Cipher Suite Lists

A cipher suite list is an object comprised of several cipher suites. Each predefined cipher suite value represents a cipher suite used to negotiate an SSL- or TLS-encrypted session. You can use cipher suites and cipher suite lists in SSL rules to control encrypted traffic based on whether the client and server negotiated the SSL session using that cipher suite. If you add a cipher suite list to an SSL rule, SSL sessions negotiated with any of the cipher suites in the list match the rule.

Note Although you can use cipher suites in the web interface in the same places as cipher suite lists, you cannot add, modify, or delete cipher suites.

Creating Cipher Suite Lists

You can use these objects with any device type except NGIPSv.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Cipher Suite List** from the list of object types.

Step 3 Click **Add Cipher Suites**.

Step 4 Enter a **Name**.

In a multidomain deployment, object names must be unique within the domain hierarchy. The system may identify a conflict with the name of an object you cannot view in your current domain.

Step 5 Choose one or more cipher suites from the **Available Ciphers** list.

Step 6 Click **Add**.

Step 7 Optionally, click **Delete** ( ) next to any cipher suites in the **Selected Ciphers** list that you want to remove.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Distinguished Name Objects

Each distinguished name object represents the distinguished name for a public key certificate's subject or issuer. You can use distinguished name objects and groups in TLS/SSL rules to control encrypted traffic based on whether the client and server negotiated the TLS/SSL session using a server certificate with the distinguished name as subject or issuer.

(A *distinguished name group* is a named collection of existing distinguished name objects.)

The distinguished name can consist of country code, common name, organization, and organizational unit, but typically consists of a common name only. For example, the common name in the certificate for `https://www.cisco.com` is `cisco.com`. The certificate can contain multiple Subject Alternative Names (SANs) you can use as DNs in a rule condition. For detailed information about SANs, see RFC 528, section 4.2.1.6.

The format of a distginguished name object that references a common name is `CN=name`. If you add a DN rule condition without `CN=`, the system prepends `CN=` before saving the object.

The Firepower System uses Server Name Indication (SNI) to match the DN in the TLS/SSL rule whenever possible.

You can also add a distinguished name with one of each of the attributes listed in the following table, separated by commas.

Table 58: Distinguished name attributes

Attribute	Description	Allowed Values
C	Country Code	two alphabetic characters
CN	Common Name	up to 64 alphanumeric, backslash (/), hyphen (-), quotation ("), or asterisk (*) characters, or spaces
O	Organization	up to 64 alphanumeric, backslash (/), hyphen (-), quotation ("), or asterisk (*) characters, or spaces
OU	Organizational Unit	up to 64 alphanumeric, backslash (/), hyphen (-), quotation ("), or asterisk (*) characters, or spaces

Important notes about DN rule conditions

- The first time the system detects an encrypted session to a new server, DN data is not available for ClientHello processing, which *might* result in an undecrypted first session.
- You *cannot* configure a distinguished name condition if you also choose the **Decrypt - Known Key** action. Because that action requires you to choose a server certificate to decrypt traffic, the certificate already matches the traffic.

Wildcard examples

You can define one or more asterisks (*) as wildcards in an attribute. In a common name attribute, you can define one or more asterisks per domain name label. wildcards match only in that label, but you can define multiple labels with wildcards. See the following table for examples.

Table 59: Common Name attribute wildcard examples

Attribute	Matches	Does Not Match
CN=*ample.com	example.com	mail.example.com example.text.com ampleexam.com
CN=exam*.com	example.com	mail.example.com example.text.com ampleexam.com
CN=*xamp*.com	example.com	mail.example.com example.text.com ampleexam.com

Attribute	Matches	Does Not Match
CN=*.example.com	mail.example.com	www.myhost.example.com example.com example.text.com ampleexam.com



Note The DN object `CN=amp.cisco.com` would *not* match a CN like `CN=auth.amp.cisco.com`, which is why we recommend wildcards in these cases.

Creating Distinguished Name Objects

You can use these objects with any device type except NGIPSv.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **Distinguished Name** node, and choose **Individual Objects**.

Step 3 Click **Add Distinguished Name**.

Step 4 Enter a **Name**.

In a multidomain deployment, object names must be unique within the domain hierarchy. The system may identify a conflict with the name of an object you cannot view in your current domain.

Step 5 In the **DN** field, enter a value for the distinguished name or common name. You have the following options:

- If you add a distinguished name, you can include one of each attribute listed in Distinguished Name Objects, on page 444 separated by commas.
- If you add a common name, you can include multiple labels and wild cards.

Step 6 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

PKI Objects

PKI Objects for SSL Application

PKI objects represent the public key certificates and paired private keys required to support your deployment. Internal and trusted CA objects consist of certificate authority (CA) certificates; internal CA objects also contain the private key paired with the certificate. Internal and external certificate objects consist of server certificates; internal certificate objects also contain the private key paired with the certificate.

If you use trusted certificate authority objects and internal certificate objects to configure a connection to ISE, you can use ISE as an identity source.

If you use internal certificate objects to configure captive portal, the system can authenticate the identity of your captive portal device when connecting to users' web browsers.

If you use trusted certificate authority objects to configure realms, you can configure secure connections to LDAP or AD servers.

If you use PKI objects in SSL rules, you can match traffic encrypted with:

- the certificate in an external certificate object
- a certificate either signed by the CA in a trusted CA object, or within the CA's chain of trust

If you use PKI objects in SSL rules, you can decrypt:

- outgoing traffic by re-signing the server certificate with an internal CA object
- incoming traffic using the known private key in an internal certificate object

You can manually input certificate and key information, upload a file containing that information, or in some cases, generate a new CA certificate and private key.

When you view a list of PKI objects in the object manager, the system displays the certificate's Subject distinguished name as the object value. Hover your pointer over the value to view the full certificate Subject distinguished name. To view other certificate details, edit the PKI object.



Note The Firepower Management Center and managed devices encrypt all private keys stored in internal CA objects and internal certificate objects with a randomly generated key before saving them. If you upload private keys that are password protected, the appliance decrypts the key using the user-supplied password, then reencrypts it with the randomly generated key before saving it.

PKI Objects for Certificate Enrollment

A certificate enrollment object contains the Certification Authority (CA) server information and enrollment parameters that are required for creating Certificate Signing Requests (CSRs) and obtaining Identity Certificates from the specified CA. These activities occur in your Private Key Infrastructure (PKI).

The certificate enrollment object may also includes certificate revocation information. For more information on PKI, digital certificates, and certificate enrollment see PKI Infrastructure and Digital Certificates , on page 845.

Internal Certificate Authority Objects

Each internal certificate authority (CA) object you configure represents the CA public key certificate of a CA your organization controls. The object consists of the object name, CA certificate, and paired private key. You can use internal CA objects and groups in SSL rules to decrypt outgoing encrypted traffic by re-signing the server certificate with the internal CA.



Note If you reference an internal CA object in a **Decrypt - Resign** SSL rule and the rule matches an encrypted session, the user's browser may warn that the certificate is not trusted while negotiating the SSL handshake. To avoid this, add the internal CA object certificate to either the client or domain list of trusted root certificates.

You can create an internal CA object in the following ways:

- import an existing RSA-based or elliptic curve-based CA certificate and private key
- generate a new self-signed RSA-based CA certificate and private key
- generate an unsigned RSA-based CA certificate and private key. You must submit a certificate signing request (CSR) to another CA to sign the certificate before using the internal CA object.

After you create an internal CA object containing a signed certificate, you can download the CA certificate and private key. The system encrypts downloaded certificates and private keys with a user-provided password.

Whether system-generated or user-created, you can modify the internal CA object name, but cannot modify other object properties.

You cannot delete an internal CA object that is in use. Additionally, after you edit an internal CA object used in an SSL policy, the associated access control policy goes out-of-date. You must re-deploy the access control policy for your changes to take effect.

CA Certificate and Private Key Import

You can configure an internal CA object by importing an X.509 v3 CA certificate and private key. You can upload files encoded in one of the following supported formats:

- Distinguished Encoding Rules (DER)
- Privacy-enhanced Electronic Mail (PEM)

If the private key file is password-protected, you can supply the decryption password. If the certificate and key are encoded in the PEM format, you can also copy and paste the information.

You can upload only files that contain proper certificate or key information, and that are paired with each other. The system validates the pair before saving the object.



Note If you configure a rule with the **Decrypt - Resign** action, the rule matches traffic based on the referenced internal CA certificate's encryption algorithm type, in addition to any configured rule conditions. You must upload an elliptic curve-based CA certificate to decrypt outgoing traffic encrypted with an elliptic curve-based algorithm, for example.

Importing a CA Certificate and Private Key

You can use these objects with any device type except NGIPSv.

In a multidomain deployment, the system displays objects created in the current domain, which you can edit. It also displays objects created in ancestor domains, which in most cases you cannot edit. To view and edit objects in a descendant domain, switch to that domain.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Internal CAs**.

Step 3 Click **Import CA**.

Step 4 Enter a **Name**.

In a multidomain deployment, object names must be unique within the domain hierarchy. The system may identify a conflict with the name of an object you cannot view in your current domain.

Step 5 Above the **Certificate Data** field, click **Browse** to upload a DER or PEM-encoded X.509 v3 CA certificate file.

Step 6 Above the **Key** field, click **Browse** to upload a DER or PEM-encoded paired private key file.

Step 7 If the uploaded file is password-protected, check the **Encrypted, and the password is:** check box, and enter the password.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Generating a New CA Certificate and Private Key

You can use these objects with any device type except NGIPSv.

You can configure an internal CA object by providing identification information to generate a self-signed RSA-based CA certificate and private key.

The generated CA certificate is valid for ten years. The Valid From date is a week before generation.

In a multidomain deployment, the system displays objects created in the current domain, which you can edit. It also displays objects created in ancestor domains, which in most cases you cannot edit. To view and edit objects in a descendant domain, switch to that domain.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Internal CAs**.

Step 3 Click **Generate CA**.

Step 4 Enter a **Name**.

In a multidomain deployment, object names must be unique within the domain hierarchy. The system may identify a conflict with the name of an object you cannot view in your current domain.

Step 5 Enter the identification attributes.

Step 6 Click **Generate self-signed CA**.

New Signed Certificates

You can configure an internal CA object by obtaining a signed certificate from a CA. This involves two steps:

- Provide identification information to configure the internal CA object. This generates an unsigned certificate and paired private key, and creates a certificate signing request (CSR) to a CA you specify.
- After the CA issues the signed certificate, upload it to the internal CA object, replacing the unsigned certificate.

You can only reference an internal CA object in an SSL rule if it contains a signed certificate.

Creating an Unsigned CA Certificate and CSR

You can use these objects with any device type except NGIPSv.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Internal CAs**.

Step 3 Click **Generate CA**.

Step 4 Enter a **Name**.

In a multidomain deployment, object names must be unique within the domain hierarchy. The system may identify a conflict with the name of an object you cannot view in your current domain.

Step 5 Enter the identification attributes.

Step 6 Click **Generate CSR**.

Step 7 Copy the CSR to submit to a CA.

Step 8 Click **OK**.

What to do next

- You must upload a signed certificate issued by a CA as described in Uploading a Signed Certificate Issued in Response to a CSR, on page 450

Uploading a Signed Certificate Issued in Response to a CSR

You can use these objects with any device type except NGIPSv.

In a multidomain deployment, the system displays objects created in the current domain, which you can edit. It also displays objects created in ancestor domains, which in most cases you cannot edit. To view and edit objects in a descendant domain, switch to that domain.

Once uploaded, the signed certificate can be referenced in SSL rules.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Internal CAs**.

Step 3 Click **Edit** ( ) next to the CA object containing the unsigned certificate awaiting the CSR.

Step 4 Click **Install Certificate**.

Step 5 Click **Browse** to upload a DER or PEM-encoded X.509 v3 CA certificate file.

Step 6 If the uploaded file is password protected, check the **Encrypted, and the password is:** check box, and enter the password.

Step 7 Click **Save** to upload a signed certificate to the CA object.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

CA Certificate and Private Key Downloads

You can back up or transfer a CA certificate and paired private key by downloading a file containing the certificate and key information from an internal CA object.

Caution Always store downloaded key information in a secure location.

The system encrypts the private key stored in an internal CA object with a randomly generated key before saving it to disk. If you download a certificate and private key from an internal CA object, the system first decrypts the information before creating a file containing the certificate and private key information. You must then provide a password the system uses to encrypt the downloaded file.

Caution Private keys downloaded as part of a system backup are decrypted, then stored in the unencrypted backup file.

Downloading a CA Certificate and Private Key

You can use these objects with any device type except NGIPSv.

In a multidomain deployment, the system displays objects created in the current domain, which you can edit. It also displays objects created in ancestor domains, which in most cases you cannot edit. To view and edit objects in a descendant domain, switch to that domain.

You can download CA certificates for both the current domain and ancestor domains.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Internal CAs**.

Step 3 Next to the internal CA object whose certificate and private key you want to download, click **Edit** ( ).

In a multidomain deployment, click **View** ( ) to download the certificate and private key for an object in an ancestor domain.

Step 4 Click **Download**.

Step 5 Enter an encryption password in the **Password** and **Confirm Password** fields.

Step 6 Click **OK**.

Trusted Certificate Authority Objects

Each trusted certificate authority (CA) object you configure represents a CA public key certificate belonging to a trusted CA. The object consists of the object name and CA public key certificate. You can use external CA objects and groups in:

- your SSL policy to control traffic encrypted with a certificate signed either by the trusted CA, or any CA within the chain of trust.
- your realm configurations to establish secure connections to LDAP or AD servers.
- your ISE connection. Select trusted certificate authority objects for the **pxGrid Server CA** and **MNT Server CA** fields.

After you create the trusted CA object, you can modify the name and add certificate revocation lists (CRL), but cannot modify other object properties. There is no limit on the number of CRLs you can add to an object. If you want to modify a CRL you have uploaded to an object, you must delete the object and recreate it.

Note Adding a CRL to an object has no effect when the object is used in your ISE integration configuration.

You cannot delete a trusted CA object that is in use. Additionally, after you edit a trusted CA object that is in use, the associated access control policy goes out-of-date. You must re-deploy the access control policy for your changes to take effect.

Trusted CA Object

You can configure an external CA object by uploading an X.509 v3 CA certificate. You can upload a file encoded in one of the following supported formats:

- Distinguished Encoding Rules (DER)
- Privacy-enhanced Electronic Mail (PEM)

If the file is password-protected, you must supply the decryption password. If the certificate is encoded in the PEM format, you can also copy and paste the information.

You can upload a CA certificate only if the file contains proper certificate information; the system validates the certificate before saving the object.

Adding a Trusted CA Object

You can use these objects with any device type except NGIPSv.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Trusted CAs**.

Step 3 Click **Add Trusted CAs**.

Step 4 Enter a **Name**.

In a multidomain deployment, object names must be unique within the domain hierarchy. The system may identify a conflict with the name of an object you cannot view in your current domain.

Step 5 Click **Browse** to upload a DER or PEM-encoded X.509 v3 CA certificate file.

Step 6 If the file is password-protected, check the **Encrypted, and the password is:** check box, and enter the password.

Step 7 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Certificate Revocation Lists in Trusted CA Objects

You can upload CRLs to a trusted CA object. If you reference that trusted CA object in an SSL policy, you can control encrypted traffic based on whether the CA that issued the session encryption certificate subsequently revoked the certificate. You can upload files encoded in one of the following supported formats:

- Distinguished Encoding Rules (DER)
- Privacy-enhanced Electronic Mail (PEM)

After you add the CRL, you can view the list of revoked certificates. If you want to modify a CRL you have uploaded to an object, you must delete the object and recreate it.

You can upload only files that contain a proper CRL. There is no limit to the number of CRLs you can add to a trusted CA object. However, you must save the object each time you upload a CRL, before adding another CRL.

Note Adding a CRL to an object has no effect when the object is used in your ISE integration configuration.

Adding a Certificate Revocation List to a Trusted CA Object

You can use these objects with any device type except NGIPSv.

In a multidomain deployment, the system displays objects created in the current domain, which you can edit. It also displays objects created in ancestor domains, which in most cases you cannot edit. To view and edit objects in a descendant domain, switch to that domain.

Note Adding a CRL to an object has no effect when the object is used in your ISE integration configuration.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Trusted CAs**.

Step 3 Click **Edit** ( ) next to a trusted CA object.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Click **Add CRL** to upload a DER or PEM-encoded CRL file.

Step 5 Click **OK**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

External Certificate Objects

Each external certificate object you configure represents a server public key certificate that does not belong to your organization. The object consists of the object name and certificate. You can use external certificate objects and groups in SSL rules to control traffic encrypted with the server certificate. For example, you can upload a self-signed server certificate that you trust, but cannot verify with a trusted CA certificate.

You can configure an external certificate object by uploading an X.509 v3 server certificate. You can upload a file in one of the following supported formats:

- Distinguished Encoding Rules (DER)
- Privacy-enhanced Electronic Mail (PEM)

You can upload only files that contains proper server certificate information; the system validates the file before saving the object. If the certificate is encoded in the PEM format, you can also copy and paste the information.

Adding External Certificate Objects

You can use these objects with any device type except NGIPSv.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **External Certs**.

Step 3 Click **Add External Cert**.

Step 4 Enter a **Name**.

In a multidomain deployment, object names must be unique within the domain hierarchy. The system may identify a conflict with the name of an object you cannot view in your current domain.

Step 5 Above the **Certificate Data** field, click **Browse** to upload a DER or PEM-encoded X.509 v3 server certificate file.

Step 6 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Internal Certificate Objects

Each internal certificate object you configure represents a server public key certificate belonging to your organization. The object consists of the object name, public key certificate, and paired private key. You can use internal certificate objects and groups in:

- your SSL rules to decrypt traffic incoming to one of your organization's servers using the known private key.
- your ISE connection. Select an internal certificate object for the **MC Server Certificate** field.
- your captive portal configuration to authenticate the identity of your captive portal device when connecting to users' web browsers. Select an internal certificate object for the **Server Certificate** field.

You can configure an internal certificate object by uploading an X.509 v3 RSA-based or elliptic curve-based server certificate and paired private key. You can upload a file in one of the following supported formats:

- Distinguished Encoding Rules (DER)
- Privacy-enhanced Electronic Mail (PEM)

If the file is password-protected, you must supply the decryption password. If the certificate and key are encoded in the PEM format, you can also copy and paste the information.

You can upload only files that contain proper certificate or key information, and that are paired with each other. The system validates the pair before saving the object.

After you create the internal certificate object, you can modify the name, but cannot modify other object properties.

You cannot delete an internal certificate object that is in use. Additionally, after you edit an internal certificate object that is in use, the associated access control policy goes out-of-date. You must re-deploy the access control policy for your changes to take effect.

Adding Internal Certificate Objects

You can use these objects with any device type except NGIPSv.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Internal Certs**.

Step 3 Click **Add Internal Cert**.

Step 4 Enter a **Name**.

In a multidomain deployment, object names must be unique within the domain hierarchy. The system may identify a conflict with the name of an object you cannot view in your current domain.

Step 5 Above the **Certificate Data** field, click **Browse** to upload a DER or PEM-encoded X.509 v3 server certificate file.

Step 6 Above the **Key** field, or click **Browse** to upload a DER or PEM-encoded paired private key file.

Step 7 If the uploaded private key file is password-protected, check the **Encrypted, and the password is:** check box, and enter the password.

Step 8 Click **Save**.

Certificate Enrollment Objects

Trustpoints let you manage and track CAs and certificates. A trustpoint is a representation of a CA or identity pair. A trustpoint includes the identity of the CA, CA-specific configuration parameters, and an association with one, enrolled identity certificate.

A certificate enrollment object contains the Certification Authority (CA) server information and enrollment parameters that are required for creating Certificate Signing Requests (CSRs) and obtaining Identity Certificates from the specified CA. These activities occur in your Private Key Infrastructure (PKI).

The certificate enrollment object may also includes certificate revocation information. For more information on PKI, digital certificates, and certificate enrollment see PKI Infrastructure and Digital Certificates , on page 845.

How to Use Certificate Enrollment Objects

Certificate Enrollment Objects are used to enroll your managed devices into your PKI infrastructure, and create trustpoints (CA objects) on devices that support VPN connections by doing the following:

1. Define parameters for CA authentication and enrollment in a Certificate Enrollment Object. Specify shared parameters and use the override facility to specify unique object setting for different devices.

2. Associate and install this object on each managed device that requires the identity certificate. On the device, it becomes a *trustpoint*.

   When a certificate enrollment object is associated with and then installed on a device, the process of certificate enrollment starts immediately. The process is automatic for self-signed and SCEP enrollment types, meaning it does not require any additional administrator action. Manual certificate enrollment and importing a PKCS12 file requires extra administrator action.

3. Specify the created trustpoint in your VPN configuration.

Managing Certificate Enrollment Objects

To manage certificate enrollment objects, go to **Objects > Object Management**, then from the navigation pane choose **PKI > Cert Enrollment**. The following information is shown:

- Existing certificate enrollment objects are listed in the **Name** column.

  Use the search field (the magnifying glass) to filter the list.

- The enrollment type of each object is shown in the **Type** column. The following enrollment methods can be used:
  - **Self Signed**—The managed device generates its own self signed root certificate.
  - **SCEP**—(Default) Simple Certificate Enrollment Protocol is used by the device to obtain an identity certificate from the CA.
  - **Manual**—The process of enrolling is carried out manually by the administrator.
  - **PKCS12 File**—Import a PKCS12 file on a Firepower Threat Defense managed device that supports VPN connectivity. A PKCS#12, or PFX, file holds the server certificate, any intermediate certificates, and the private key in one encrypted file.

- The **Override** column indicates whether the object allows overrides (a green check mark) or not (a red X). If a number is displayed, it is the number of overrides in place.

  Use the Override option to customize the object settings for each device that is part of the VPN configuration. Overriding makes each device's trustpoint details unique. Typically the Common Name or Subject is overridden for each device in the VPN configuration.

  See Object Overrides, on page 399 for details and procedures on overriding objects of any type.

- **Edit** a previously created certificate enrollment object by clicking on the edit icon (a pencil). Editing can only be done if the enrollment object is not associated with any managed devices. Refer to the adding instructions for editing a certificate enrollment object.

- **Delete** a previously created certificate enrollment object by clicking on the delete icon (a trash can). You cannot delete a certificate enrollment object if it is associated with any managed device.

Press **(+) Add Cert Enrollment** to open the **Add Cert Enrollment** dialog and configure a Certificate Enrollment Object, see Adding Certificate Enrollment Objects, on page 458. Then install the certificate on each managed, headend device.

Related Topics

Installing a Certificate Using Self-Signed Enrollment , on page 483
Installing a Certificate Using SCEP Enrollment, on page 484
Installing a Certificate Using Manual Enrollment, on page 485

Installing a Certificate by Importing a PKCS12 File, on page 486

Adding Certificate Enrollment Objects

You can use these objects with Firepower Threat Defense devices. You must have Admin or Network Admin privileges to do this task.

Procedure

Step 1 Open the **Add Cert Enrollment** dialog:

- Directly from Object Management: In the **Objects > Object Management** screen, choose **PKI > Cert Enrollment** from the navigation pane, and press **Add Cert Enrollment**.
- While configuring a managed device: In the **Devices > Certificates** screen, choose **Add > Add New Certificate** and click **(+)** for the **Certificate Enrollment** field.

Step 2 Enter the **Name**, and optionally, a **Description** of this enrollment object.

When enrollment is complete, this name is the name of the trustpoint on the managed devices with which it is associated.

Step 3 Open the **CA Information** tab and choose the **Enrollment Type**.

- **Self-Signed Certificate**—The managed device, acting as a CA, generates its own self-signed root certificate. No other information is needed in this pane.

  Note When enrolling a self-signed certificate you must specify the Common Name (CN) in the certificate parameters.

- **SCEP**—(Default) Simple Certificate Enrollment Protocol. Specify the SCEP information. See Certificate Enrollment Object SCEP Options, on page 459.
- **Manual**—Paste an obtained CA certificate in the **CA Certificate** box. You can also obtain a CA certificate by copying it from another device.
- **PKCS12 File**—Import a PKCS12 file on a Firepower Threat Defense managed device that supports VPN connectivity. A PKCS#12, or PFX, file holds a server certificate, intermediate certificates, and a private key in one encrypted file.

Step 4 (Optional) Open the **Certificate Parameters** tab and specify the certificate contents. See Certificate Enrollment Object Certificate Parameters, on page 460.

This information is placed in the certificate and is readable by any party who receives the certificate from the router.

Step 5 (Optional) Open the **Key** tab and specify the Key information. See Certificate Enrollment Object Key Options, on page 460.

Step 6 (Optional) Click the **Revocation** tab, and specify the revocation options: See Certificate Enrollment Object Revocation Options, on page 461.

Step 7 **Allow Overrides** of this object if desired. See Object Overrides, on page 399 for a full description of object overrides.

What to do next

Associate and install the enrollment object on a device to create a trustpoint on that device.

Related Topics

Installing a Certificate Using Self-Signed Enrollment , on page 483
Installing a Certificate Using SCEP Enrollment, on page 484
Installing a Certificate Using Manual Enrollment, on page 485
Installing a Certificate by Importing a PKCS12 File, on page 486

Certificate Enrollment Object SCEP Options

Firepower Management Center Navigation Path

Objects > Object Management, then from the navigation pane choose **PKI > PKI Enrollment**. Press **(+) Add PKI Enrollment** to open the **Add PKI Enrollment** dialog, and select the **CA Information** tab.

Fields

Enrollment Type—set to **SCEP**.

Enrollment URL—The URL of the CA server to which devices should attempt to enroll.

Use an HTTP URL in the form of **`http://CA_name:port,`** where CA_name is the host DNS name or IP address of the CA server. The port number is mandatory.

If the CA cgi-bin script location at the CA is not the default (/cgi-bin/pkiclient.exe), you must also include the nonstandard script location in the URL, in the form of http://CA_name:port/script_location, where script_location is the full path to the CA scripts.

Challenge Password / Confirm Password—The password used by the CA server to validate the identity of the device. You can obtain the password by contacting the CA server directly or by entering the following address in a web browser: **`http://URLHostName/certsrv/mscep/mscep.dll`**. The password is good for 60 minutes from the time you obtain it from the CA server. Therefore, it is important that you deploy the password as soon as possible after you create it.

Retry Period—The interval between certificate request attempts, in minutes. Value can be 1 to 60 minutes. The default is 1 minute.

Retry Count—The number of retries that should be made if no certificate is issued upon the first request. Value can be 1 to 100. The default is 10.

CA Certificate Source—Specify how the CA certificate will be obtained.

- **Retrieve Using SCEP** (Default, and only supported option)—Retrieve the certificate from the CA server using the Simple Certificate Enrollment Process (SCEP). Using SCEP requires a connection between your device and the CA server. Ensure there is a route from your device to the CA server before beginning the enrollment process.

Fingerprint—When retrieving the CA certificate using SCEP, you may enter the fingerprint for the CA server. Using the fingerprint to verify the authenticity of the CA server's certificate helps prevent an unauthorized party from substituting a fake certificate in place of the real one. Enter the **Fingerprint** for the CA server in hexadecimal format. If the value you enter does not match the fingerprint on the certificate, the certificate is rejected. Obtain the CA's fingerprint by contacting the server directly, or by entering the following address in a web browser: **`http://<URLHostName>/certsrv/mscep/mscep.dll`**.

Certificate Enrollment Object Certificate Parameters

Specify additional information in certificate requests sent to the CA server. This information is placed in the certificate and can be viewed by any party who receives the certificate from the router.

Firepower Management Center Navigation Path

Objects > Object Management, then from the navigation pane choose **PKI > PKI Enrollment**. Press **(+) Add PKI Enrollment** to open the **Add PKI Enrollment** dialog, and select the **Certificate Parameters** tab.

Fields

Enter all information using the standard LDAP X.500 format.

- **Include FQDN**—Whether to include the device's fully qualified domain name (FQDN) in the certificate request. Choices are:
  - **Use Device Hostname as FQDN**
  - **Don't use FQDN in certificate**
  - **Custom FQDN**—Select this and then specify it in the **Custom FQDN** field that displays.
- **Include Device's IP Address**—The interface whose IP address is included in the certificate request.
- **Common Name (CN)**—The X.500 common name to include in the certificate.



Note When enrolling a self-signed certificate you must specify the Common Name (CN) in the certificate parameters.

- **Organization Unit (OU)**—The name of the organization unit (for example, a department name) to include in the certificate.
- **Organization (O)**—The organization or company name to include in the certificate.
- **Locality (L)**—The locality to include in the certificate.
- **State (ST)**—The state or province to include in the certificate.
- **County Code (C)**—The country to include in the certificate. These codes conform to ISO 3166 country abbreviations, for example "US" for the United States of America.
- **Email (E)**—The email address to include in the certificate.
- **Include Device's Serial Number**—Whether to include the serial number of the device in the certificate. The CA uses the serial number to either authenticate certificates or to later associate a certificate with a particular device. If you are in doubt, include the serial number, as it is useful for debugging purposes.

Certificate Enrollment Object Key Options

Firepower Management Center Navigation Path

Objects > Object Management, then from the navigation pane choose **PKI > Cert Enrollment**. Press **(+) Add Cert Enrollment** to open the **Add Cert Enrollment** dialog, and select the **Key** tab.

Fields

- **Key Type**—RSA, ECDSA.
- **Key Name**—If the key pair you want to associate with the certificate already exists, this field specifies the name of that key pair. If the key pair does not exist, this field specifies the name to assign to the key pair that will be generated during enrollment. If you do not specify a name, the fully qualified domain name (FQDN) key pair is used instead.
- **Key Size**—If the key pair does not exist, defines the desired key size (modulus), in bits. The recommended size is 2048 bits. The larger the modulus size, the more secure the key. However, keys with larger modulus sizes take longer to generate (a minute or more when larger than 512 bits) and longer to process when exchanged.

Certificate Enrollment Object Revocation Options

Specify whether to check the revocation status of a certificate by choosing and configuring the method. Revocation checking is off by default, neither method (CRL or OCSP) is checked.

Firepower Management Center Navigation Path

Objects > Object Management, then from the navigation pane choose **PKI > PKI Enrollment**. Press **(+) Add PKI Enrollment** to open the **Add PKI Enrollment** dialog, and select the **Revocation** tab.

Fields

- **Enable Certificate Revocation Lists**—Check to enable CRL checking.
  - **Use CRL distribution point from the certificate**—Check to obtain the revocation lists ditribution URL from the certificate.
  - **Use static URL configured**—Check this to add a static, pre-defined distribution URL for revocation lists. Then add the URLs.

    CRL Server URLs—The URL of the LDAP server from which the CRL can be downloaded. This URL must start with **`ldap://`**, and include a port number in the URL.
- **Enable Online Certificate Status Protocol (OCSP)**—Check to enable OCSP checking.

  OCSP Server URL—The URL of the OCSP server checking for revocation if you require OCSP checks. This URL must start with **`http://`**.
- **Consider the certificate valid if revocation information can not be reached**—Checked by default. Uncheck if you do not want to allow this.

SLA Monitor Objects

Each Internet Protocol Service Level Agreement (SLA) monitor defines a connectivity policy to a monitored address and tracks the availability of a route to the address. The route is periodically checked for availability by sending ICMP echo requests and waiting for the response. If the requests time out, the route is removed from the routing table and replaced with a backup route. SLA monitoring jobs start immediately after deployment and continue to run unless you remove the SLA monitor from the device configuration (that is, they do not age out). The Internet Protocol Service Level Agreement (SLA) Monitor Object is used in the

Route Tracking field of an IPv4 Static Route Policy. IPv6 routes do not have the option to use SLA monitor via route tracking.

You can use these objects with Firepower Threat Defense devices.

Procedure

Step 1 Select **Objects** > **Object Management** and choose **SLA Monitor** from the table of contents.

Step 2 Click **Add SLA Monitor**.

Step 3 Enter a name for the object in the **Name** field.

Step 4 (Optional) Enter a description for the object in the **Description** field.

Step 5 Enter the frequency of ICMP echo request transmissions, in seconds, in the **Frequency** field. Valid values range from 1 to 604800 seconds (7 days). The default is 60 seconds.

Note The frequency cannot be less than the timeout value; you must convert frequency to milliseconds to compare the values.

Step 6 Enter the ID number of the SLA operation in the **SLA Monitor ID** field. Values range from 1 to 2147483647. You can create a maximum of 2000 SLA operations on a device. Each ID number must be unique to the policy and the device configuration.

Step 7 Enter the amount of time that must pass after an ICMP echo request before a rising threshold is declared, in milliseconds, in the **Threshold** field. Valid values range from 0 to 2147483647 milliseconds. The default is 5000 milliseconds. The threshold value is used only to indicate events that exceed the defined value. You can use these events to evaluate the proper timeout value. It is not a direct indicator of the reachability of the monitored address.

Note The threshold value should not exceed the timeout value.

Step 8 Enter the amount of time that the SLA operation waits for a response to the ICMP echo requests, in milliseconds, in the **Timeout** field. Values range from 0 to 604800000 milliseconds (7 days). The default is 5000 milliseconds. If a response is not received from the monitored address within the amount of time defined in this field, the static route is removed from the routing table and replaced by the backup route.

Note The timeout value cannot exceed the frequency value (adjust the frequency value to milliseconds to compare the numbers).

Step 9 Enter the size of the ICMP request packet payload, in bytes, in the **Data Size** field. Values range from 0 to 16384 bytes. The default is 28 bytes, which creates a total ICMP packet of 64 bytes. Do not set this value higher than the maximum allowed by the protocol or the Path Maximum Transmission Unit (PMTU). For purposes of reachability, you might need to increase the default data size to detect PMTU changes between the source and the target. A low PMTU can affect session performance and, if detected, might indicate that the secondary path should be used.

Step 10 Enter a value for type of service (ToS) defined in the IP header of the ICMP request packet in the **ToS** field. Values range from 0 to 255. The default is 0. This field contains information such as delay, precedence, reliability, and so on. It can be used by other devices on the network for policy routing and features such as committed access rate.

Step 11 Enter the number of packets that are sent in the **Number of Packets** field. Values range from 1 to 100. The default is 1 packet.

Note Increase the default number of packets if you are concerned that packet loss might falsely cause the Firepower Threat Defense device to believe that the monitored address cannot be reached.

Step 12 Enter the IP address that is being monitored for availability by the SLA operation, in the **Monitored Address** field.

Step 13 The **Available Zones** list displays both zones and interface groups. In the **Zones/Interfaces** list, add the zones or interface groups that contain the interfaces through which the device communicates with the management station. To specify a single interface, you need to create a zone or the interface groups for the interface; see Creating Security Zone and Interface Group Objects, on page 411. The host will be configured on a device only if the device includes the selected interfaces or zones.

Step 14 Click **Save**.

Prefix Lists

You can create prefix list objects for IPv4 and IPv6 to use when you are configuring route maps, policy maps, OSPF Filtering, or BGP Neighbor Filtering.

Configure IPv6 Prefix List

Use the Configure IPv6 Prefix list page to create, copy and edit prefix list objects. You can create prefix list objects to use when you are configuring route maps, policy maps, OSPF Filtering, or BGP Neighbor Filtering.

You can use this object with Firepower Threat Defense devices.

Procedure

Step 1 Select **Objects** > **Object Management** and choose **Prefix Lists** > **IPv6 Prefix List** from the table of contents.

Step 2 Click **Add Prefix List**.

Step 3 Enter a name for the prefix list object in the **Name** field on the **New Prefix List Object** window.

Step 4 Click **Add** on the**New Prefix List Object** window.

Step 5 Select the appropriate action, Allow or Block from the **Action** drop-down list, to indicate the redistribution access.

Step 6 Enter a unique number that indicates the position a new prefix list entry will have in the list of prefix list entries already configured for this object, in the **Sequence No.** field. If left blank, the sequence number will default to five more than the largest sequence number currently in use.

Step 7 Specify the IPv6 address in the IP address/mask length format in the **IP address** field. The mask length must be a valid value between 1-128.

Step 8 Enter the minimum prefix length in the **Minimum Prefix Length** field. The value must be greater than the mask length and less than or equal to the Maximum Prefix Length, if specified.

Step 9 Enter the maximum prefix length in the **Maximum Prefix Length** field. The value must be greater than or equal to the Minimum Prefix Length, if present, or greater than the mask length if the Minimum Prefix Length is not specified.

Step 10 Click **Add**.

Step 11 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 401.

Step 12 Click **Save**.

Configure IPv4 Prefix List

Use the Configure IPv4 Prefix list page to create, copy and edit prefix list objects. You can create prefix list objects to use when you are configuring route maps, policy maps, OSPF Filtering, or BGP Neighbor Filtering.

You can use this object with Firepower Threat Defense devices.

Procedure

Step 1 Select **Objects** > **Object Management** and choose **Prefix Lists** > **IPv4 Prefix List** from the table of contents.

Step 2 Click **Add Prefix List**.

Step 3 Enter a name for the prefix list object in the **Name** field on the **New Prefix List Object** window.

Step 4 Click **Add**.

Step 5 Select the appropriate action, Allow or Block from the **Action** drop-down list, to indicate the redistribution access.

Step 6 Enter a unique number that indicates the position a new prefix list entry will have in the list of prefix list entries already configured for this object, in the **Sequence No.** field. If left blank, the sequence number will default to five more than the largest sequence number currently in use.

Step 7 Specify the IPv4 address in the IP address/mask length format in the **IP address** field. The mask length must be a valid value between 1- 32.

Step 8 Enter the minimum prefix length in the **Minimum Prefix Length** field. The value must be greater than the mask length and less than or equal to the Maximum Prefix Length, if specified.

Step 9 Enter the maximum prefix length in the **Maximum Prefix Length** field. The value must be greater than or equal to the Minimum Prefix Length, if present, or greater than the mask length if the Minimum Prefix Length is not specified.

Step 10 Click **Add**.

Step 11 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 401.

Step 12 Click **Save**.

Route Maps

Route maps are used when redistributing routes into any routing process. They are also used when generating a default route into a routing process. A route map defines which of the routes from the specified routing protocol are allowed to be redistributed into the target routing process. Configure a route map, to create a new route map entry for a Route Map object or to edit an existing one.

You can use this object with Firepower Threat Defense devices.

Before you begin

A Route Map may use one or mores of these objects; it is not mandatory to add all these objects. Create and use any of these objects as required, to configure your route map.

- Add ACLs.
- Add Prefix Lists.
- Add AS Path.
- Add Community Lists.
- Add Policy Lists.

Procedure

Step 1 Select **Objects** > **Object Management** and choose **Route Map** from the table of contents.

Step 2 Click **Add Route Map**.

Step 3 Click **Add** on the**New Route Map Object** window.

Step 4 In the **Sequence No.** field, enter a number, between 0 and 65535, that indicates the position a new route map entry will have in the list of route maps entries already configured for this route map object.

Note We recommend that you number clauses in intervals of at least 10 to reserve numbering space in case you need to insert clauses in the future.

Step 5 Select the appropriate action, Allow or Block from the **Redistribution** drop-down list, to indicate the redistribution access.

Step 6 Click the **Match Clauses** tab to match (routes/traffic) based on the following criteria, which you select in the table of contents:

- **Security Zones** — Match traffic based on the (ingress/egress) interfaces. You can select zones and add them, or type in interface names and add them.
- **IPv4** — Match IPv4 (routes/traffic) based on the following criteria; select the tab to define the criteria.
  a. Click the **Address** tab to match routes based on the route address. For IPv4 addresses, choose whether to use an Access list or Prefix list for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching.
  b. Click the **Next Hop** tab to match routes based on the next hop address of a route. For IPv4 addresses, choose whether to use an access list or Prefix list for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching.
  c. Click the **Route Source** tab to match routes based on the advertising source address of the route. For IPv4 addresses, choose whether to use an access list or Prefix list for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching.
- **IPv6** — Match IPv6 (routes/traffic) based on the route address, next-hop address or advertising source address of route.
- **BGP** — Match BGP (routes/traffic) based on the following criteria; select the tab to define the criteria.

a. Click the **AS Path** tab to enable matching the BGP autonomous system path access list with the specified path access list. If you specify more than one path access list, then the route can match either path access list.

b. Click the **Community List** tab to enable matching the BGP community with the specified community. If you specify more than one community, then the route can match either community. Any route that does not match at least one Match community will not be advertised for outbound route maps.

c. Click the **Policy List** tab to configure a route map to evaluate and process a BGP policy. When multiple policy lists perform matching within a route map entry, all policy lists match on the incoming attribute only.

- **Others** — Match routes or traffic based on the following criteria.

a. Enter the metric values to use for matching in the **Metric Route Value** field, to enable matching the metric of a route. You can enter multiple values separated by commas. This setting allows you to match any routes that have a specified metric. The metric values can range from 0 to 4294967295.

b. Enter the tag values to use for matching in the **Tag Values** field. You can enter multiple values separated by commas. This setting allows you to match any routes that have a specified security group tag. The tag values can range from 0 to 4294967295.

c. Check the appropriate **Route Type** option to enable matching of the route type. Valid route types are External1, External2, Internal, Local, NSSA-External1, and NSSA-External2. You can choose more than one route type from the list.

Step 7 Click the **Set Clauses** tab to set routes/traffic based on the following criteria, which you select in the table of contents:

- **Metric Values** — Set either Bandwidth, all of the values or none of the values.

a. Enter a metric value or bandwidth in Kbits per second in the **Bandwidth** field. Valid values are an integer value in the range from 0 to 4294967295.

b. Select to specify the type of metric for the destination routing protocol, from the **Metric Type** drop-down list. Valid values are : internal, type-1, or type-2.

- **BGP Clauses** — Set BGP routes based on the following criteria; select the tab to define the criteria.

a. Click the **AS Path** tab to modify an autonomous system path for BGP routes.

1. Enter an AS path number in the **Prepend AS Path** field to prepend an arbitrary autonomous system path string to BGP routes. Usually the local AS number is prepended multiple times, increasing the autonomous system path length. If you specify more than one AS path number then the route can prepend either AS number.

2. Enter an AS path number in the **Prepend Last AS to AS Path** field to prepend the AS path with the last AS number. Enter a value for the AS number from 1 to 10.

3. Check the **Convert route tag into AS path** check box to convert the tag of a route into an autonomous system path.

b. Click the **Community List** tab to set the community attributes.

1. Click the **None** radio button, to remove the community attribute from the prefixes that pass the route map.

2. Click the **Specific Community** radio button, to enter a community number, if applicable. Valid values are from 1 to 4294967295.
3. Check the **Add to existing communities** check box, to add the community to the already existing communities.
4. Select the **Internet**, **No-Advertise**, or **No-Export** check-boxes to use one of the well-known communities.

c. Click the **Others** tab to set additional attributes.

1. Check the **Set Automatic Tag** check-box to automatically compute the tag value.
2. Enter a preference value for the autonomous system path in the **Set Local Preference** field. Enter a value between 0 and 4294967295.
3. Enter a BGP weight for the routing table in the **Set Weight** field. Enter a value between 0 and 65535.
4. Select to specify the BGP origin code. Valid values are **Local IGP** Local IGP and **Incomplete**.
5. In the IPv4 Settings section, specify a next hop IPv4 address of the next hop to which packets are output. It need not be an adjacent router. If you specify more than one IPv4 address then the packets can output at either IP address.

   Select to specify an IPv4 prefix list in the **Prefix List** drop-down list.
6. In the IPv6 Settings section, specify a next hop IPv6 address of the next hop to which packets are output. It need not be an adjacent router. If you specify more than one IPv6 address then the packets can output at either IP address.

   Select to specify an IPv6 prefix in the **Prefix List** drop-down list.

Step 8 Click **Add**.

Step 9 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 401.

Step 10 Click **Save**.

Access List

An access list object, also known as an access control list (ACL), selects the traffic to which a service will apply. You use these objects when configuring particular features, such as route maps, for Firepower Threat Defense devices. Traffic identified as allowed by the ACL is provided the service, whereas “blocked” traffic is excluded from the service. Excluding traffic from a service does not necessarily mean that it is dropped altogether.

You can configure the following types of ACL:

- Extended—Identifies traffic based on source and destination address and ports. Supports IPv4 and IPv6 addresses, which you can mix in a given rule.
- Standard—Identifies traffic based on destination address only. Supports IPv4 only.

An ACL is composed of one or more access control entry (ACE), or rule. The order of ACEs is important. When the ACL is evaluated to determine if a packet matches an “allowed” ACE, the packet is tested against each ACE in the order in which the entries are listed. After a match is found, no more ACEs are checked. For example, if you want to “allow” 10.100.10.1, but “block” the rest of 10.100.10.0/24, the allow entry must come before the block entry. In general, place more specific rules at the top of an ACL.

Packets that do not match an “allow” entry are considered to be blocked.

The following topics explain how to configure ACL objects.

Configure Extended ACL Objects

Use extended ACL objects when you want to match traffic based on source and destination addresses, protocol and port, or if the traffic is IPv6.

Procedure

Step 1 Select **Objects** > **Object Management** and choose **Access List** > **Extended** from the table of contents.

Step 2 Do one of the following:

- Click **Add Extended Access List** to create a new object.
- Click **Edit** ( ) to edit an existing object.

Step 3 In the Extended ACL Object dialog box, enter a name for the object (no spaces allowed), and configure the access control entries:

a) Do one of the following:

- Click **Add** to create a new entry.
- Click **Edit** ( ) to edit an existing entry.

b) Select the **Action**, whether to Allow (match) or Block (not match) the traffic criteria.

Note The **Logging**, **Log Level**, and **Log Interval** options are used for access rules only (ACLs attached to interfaces or applied globally). Because ACL objects are not used for access rules, leave these values at their defaults.

c) Configure the source and destination addresses on the **Network** tab using any of the following techniques:

- Select the desired network objects or groups from the Available list and click **Add to Source** or **Add to Destination**. You can create new objects by clicking the + button above the list. You can mix IPv4 and IPv6 addresses.
- Type an address in the edit box below the source or destination list and click **Add**. You can specify a single host address (such as 10.100.10.5 or 2001:DB8::0DB8:800:200C:417A), or a subnet (in 10.100.10.0/24 or 10.100.10.0 255.255.255.0 format, or for IPv6, 2001:DB8:0:CD30::/60).

d) Click the **Port** tab and configure the service using any of the following techniques.

- Select the desired port objects from the Available list and click **Add to Source** or **Add to Destination**. You can create new objects by clicking the + button above the list. The object can specify TCP/UDP ports, ICMP/ICMPv6 message types, or other protocols (including "any"). However, the source port, which you typically would leave empty, accepts TCP/UDP only. You cannot select port groups.
- Type or select a port or protocol in the edit box below the source or destination list and click **Add**.

Note To get an entry that applies to all IP traffic, select a destination port object that specifies "all" protocols.

e) Click **Add** to add the entry to the object.
f) If necessary, click and drag the entry to move it up or down in the rule order to the desired location.

Repeat the process to create or edit additional entries in the object.

Step 4 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 401.

Step 5 Click **Save**.

Configure Standard ACL Objects

Use standard ACL objects when you want to match traffic based on destination IPv4 address only. Otherwise, use extended ACLs.

Procedure

Step 1 Select **Objects** > **Object Management** and choose **Access List** > **Standard** from the table of contents.

Step 2 Do one of the following:

- Click **Add Standard Access List** to create a new object.
- Click **Edit** ( ) to edit an existing object.

Step 3 In the Standard ACL Object dialog box, enter a name for the object (no spaces allowed), and configure the access control entries:

a) Do one of the following:

- Click **Add** to create a new entry.
- Click **Edit** ( ) to edit an existing entry.

b) For each access control entry, configure the following properties:

- **Action**—Whether to Allow (match) or Block (not match) the traffic criteria.
- **Network**—Add the IPv4 network objects or groups that identify the destination of the traffic.

c) Click **Add** to add the entry to the object.

d) If necessary, click and drag the entry to move it up or down in the rule order to the desired location.

Repeat the process to create or edit additional entries in the object.

Step 4 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 401.

Step 5 Click **Save**.

AS Path Objects

An AS Path is a mandatory attribute to set up BGP. It is a sequence of AS numbers through which a network can be accessed. An AS-PATH is a sequence of intermediate AS numbers between source and destination routers that form a directed route for packets to travel. Neighboring autonomous systems (ASes) use BGP to exchange and update messages about how to reach different AS prefixes. After each router makes a new local decision on the best route to a destination, it will send that route, or path information, along with the accompanying distance metrics and path attributes, to each of its peers. As this information travels through the network, each router along the path prepends its unique AS number to a list of ASes in the BGP message. This list is the route's AS-PATH. An AS-PATH along with an AS prefix, provides a specific handle for a one-way transit route through the network. Use the Configure AS Path page to create, copy and edit autonomous system (AS) path policy objects. You can create AS path objects to use when you are configuring route maps, policy maps, or BGP Neighbor Filtering. An AS path filter allows you to filter the routing update message by using regular expressions.

You can use this object with Firepower Threat Defense devices.

Procedure

Step 1 Select **Objects** > **Object Management** and choose **AS Path** from the table of contents.

Step 2 Click **Add AS Path**.

Step 3 Enter a name for the AS Path object in the **Name** field. Valid values are between 1 and 500.

Step 4 Click **Add** on the **New AS Path Object** window.

a) Select the Allow or Block options from the **Action** drop-down list to indicate redistribution access.

b) Specify the regular expression that defines the AS path filter in the **Regular Expression** field.

c) Click **Add**.

Step 5 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 401.

Step 6 Click **Save**.

Community Lists

A Community is an optional transitive BGP attribute. A community is a group of destinations that share some common attribute. It is used for route tagging. The BGP community attribute is a numerical value that can be assigned to a specific prefix and advertised to other neighbors. Communities can be used to mark a set of

prefixes that share a common attribute. Upstream providers can use these markers to apply a common routing policy such as filtering or assigning a specific local preference or modifying other attributes. Use the Configure Community Lists page to create, copy and edit community list policy objects. You can create community list objects to use when you are configuring route maps or policy maps. You can use community lists to create groups of communities to use in a match clause of a route map. The community list is an ordered list of matching statements. Destinations are matched against the rules until a match is found.

You can use this object with Firepower Threat Defense devices.

Procedure

Step 1 Select **Objects** > **Object Management** and choose **Community List** from the table of contents.

Step 2 Click **Add Community List**.

Step 3 In the **Name** field, specify a name for the community list object.

Step 4 Click **Add** on the **New Community List Object** window.

Step 5 Select the **Standard** radio button to indicate the community rule type.

Standard community lists are used to specify well-known communities and community numbers.

Note You cannot have entries using Standard and entries using Expanded community rule types in the same Community List object.

a) Select the Allow or Block options from the **Action** drop-down list to indicate redistribution access.
b) In the **Communities** field, specify a community number. Valid values can be from 1 to 4294967295 or from 0:1 to 65534:65535.
c) Select the appropriate **Route Type**.

- **Internet** — Select to specify the Internet well-known community. Routes with this community are advertised to all peers (internal and external).
- **No Advertise** — Select to specify the no-advertise well-known community. Routes with this community are not advertised to any peer (internal or external).
- **No Export** — Select to specify the no-export well-known community. Routes with this community are advertised to only peers in the same autonomous system or to only other sub-autonomous systems within a confederation. These routes are not advertised to external peers.

Step 6 Select the **Expanded** radio button to indicate the community rule type.

Expanded community lists are used to filter communities using a regular expression. Regular expressions are used to specify patterns to match COMMUNITIES attributes.

a) Select the Allow or Block options from the **Action** drop-down list to indicate redistribution access.
b) Specify the regular expression in the **Expressions** field.

Step 7 Click **Add**.

Step 8 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 401.

Step 9 Click **Save**.

Policy Lists

Use the Configure Policy List page to create, copy, and edit policy list policy objects. You can create policy list objects to use when you are configuring route maps. When a policy list is referenced within a route map, all of the match statements within the policy list are evaluated and processed. Two or more policy lists can be configured with a route map. A policy list can also coexist with any other preexisting match and set statements that are configured within the same route map but outside of the policy list. When multiple policy lists perform matching within a route map entry, all policy lists match on the incoming attribute only.

You can use this object with Firepower Threat Defense devices.

Procedure

Step 1 Select **Objects** > **Object Management** and choose **Policy List** from the table of contents.

Step 2 Click **Add Policy List**.

Step 3 Enter a name for the policy list object in the **Name** field. Object names are not case-sensitive.

Step 4 Select whether to allow or block access for matching conditions from the **Action** drop-down list.

Step 5 Click the **Interface** tab to distribute routes that have their next hop out of one of the interfaces specified.

In the **Zones/Interfaces** list, add the zones that contain the interfaces through which the device communicates with the management station. For interfaces not in a zone, you can type the interface name into the field below the **Selected Zone/Interface** list and click **Add**. The host will be configured on a device only if the device includes the selected interfaces or zones.

Step 6 Click the **Address** tab to redistribute any routes that have a destination address that is permitted by a standard access list or prefix list.

Choose whether to use an **Access List** or **Prefix List** for matching and then enter or select the Standard Access List Objects or Prefix list objects you want to use for matching.

Step 7 Click the **Next Hop** tab to redistribute any routes that have a next hop router address passed by one of the access lists or prefix lists specified.

Choose whether to use an **Access List** or **Prefix List** for matching and then enter or select the Standard Access List Objects or Prefix list objects you want to use for matching.

Step 8 Click the **Route Source** tab to redistribute routes that have been advertised by routers and access servers at the address specified by the access lists or prefix list.

Choose whether to use an **Access List** or **Prefix List** for matching and then enter or select the Standard Access List Objects or Prefix list objects you want to use for matching.

Step 9 Click the **AS Path** tab to match a BGP autonomous system path. If you specify more than one AS path, then the route can match either AS path.

Step 10 Click the **Community Rule** tab to enable matching the BGP community with the specified community. If you specify more than one community, then the route can match either community. To enable matching the BGP community exactly with the specified community, check the **Match the specified community exactly** check box.

Step 11 Click the **Metric & tag** tab to match the metric and security group tag of a route.

a) Enter the metric values to use for matching in the **Metric** field. You can enter multiple values separated by commas. This setting allows you to match any routes that have a specified metric. The metric values can range from 0 to 4294967295.
b) Enter the tag values to use for matching in the **Tag** field. You can enter multiple values separated by commas. This setting allows you to match any routes that have a specified security group tag. The tag values can range from 0 to 4294967295.

Step 12 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 401.

Step 13 Click **Save**.

VPN Objects

You can use the following VPN objects on Firepower Threat Defense devices. To use these objects, you must have Admin privileges, and your Smart License account must satisfy export controls. You can configure these objects in leaf domains only.

Firepower Threat Defense IKE Policies

Internet Key Exchange (IKE) is a key management protocol that is used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and automatically establish IPsec security associations (SAs). The IKE negotiation comprises two phases. Phase 1 negotiates a security association between two IKE peers, which enables the peers to communicate securely in Phase 2. During Phase 2 negotiation, IKE establishes SAs for other applications, such as IPsec. Both phases use proposals when they negotiate a connection. An IKE proposal is a set of algorithms that two peers use to secure the negotiation between them. IKE negotiation begins by each peer agreeing on a common (shared) IKE policy. This policy states which security parameters are used to protect subsequent IKE negotiations.

For IKEv1, IKE proposals contain a single set of algorithms and a modulus group. You can create multiple, prioritized policies to ensure that at least one policy matches a remote peer's policy. Unlike IKEv1, in an IKEv2 proposal, you can select multiple algorithms and modulus groups in one policy. Since peers choose during the Phase 1 negotiation, this makes it possible to create a single IKE proposal, but consider multiple, different proposals to give higher priority to your most desired options. For IKEv2, the policy object does not specify authentication, other policies must define the authentication requirements.

An IKE policy is required when you configure a site-to-site IPsec VPN. For more information, see Firepower Threat Defense VPN, on page 837.

Configure IKEv1 Policy Objects

Use the IKEv1 Policy page to create, delete, or edit an IKEv1 policy object. These policy objects contain the parameters required for IKEv1 policies.

Procedure

Step 1 Choose **Objects** > **Object Management** and then **VPN** > **IKEv1 Policy** from the table of contents.

Previously configured policies are listed including system defined defaults. Depending on your level of access, you may **Edit** (), View **View** (), or **Delete** () a proposal.

Step 2 (Optional) Choose **Add** ()**Add IKEv1 Policy** to create a new policy object.

Step 3 Enter a **Name** for this policy. A maximum of 128 characters is allowed.

Step 4 (Optional) Enter a **Description** for this proposal. A maximum of 1,024 characters is allowed.

Step 5 Enter the **Priority** value of the IKE policy.

The priority value determines the order of the IKE policy compared by the two negotiating peers when attempting to find a common security association (SA). If the remote IPsec peer does not support the parameters selected in your first priority policy, it tries to use the parameters defined in the next lowest priority. Valid values range from 1 to 65,535. The lower the number, the higher the priority. If you leave this field blank, Management Center assigns the lowest unassigned value starting with 1, then 5, then continuing in increments of 5.

Step 6 Choose the **Encryption** method.

When deciding which encryption and Hash Algorithms to use for the IKEv1 policy, your choice is limited to algorithms supported by the peer devices. For an extranet device in the VPN topology, you must choose the algorithm that matches both peers. For IKEv1, select one of the options. For a full explanation of the options, see Deciding Which Encryption Algorithm to Use, on page 843.

Step 7 Choose the **Hash** Algorithm that creates a Message Digest, which is used to ensure message integrity.

When deciding which encryption and Hash Algorithms to use for the IKEv1 proposal, your choice is limited to algorithms supported by the managed devices. For an extranet device in the VPN topology, you must choose the algorithm that matches both peers. For a full explanation of the options, see Deciding Which Hash Algorithms to Use, on page 843.

Step 8 Set the **DH Group**.

The Diffie-Hellman group to use for encryption. A larger modulus provides higher security but requires more processing time. The two peers must have a matching modulus group. Select the group that you want to allow in the VPN.For a full explanation of the options, see Deciding Which Diffie-Hellman Modulus Group to Use, on page 844.

Step 9 Set the **Lifetime**of the security association (SA), in seconds. You can specify a value from 120 to 2,147,483,647 seconds. The default is 86400.

When the lifetime is exceeded, the SA expires and must be renegotiated between the two peers. Generally, the shorter the lifetime (up to a point), the more secure your IKE negotiations. However, with longer lifetimes, future IPsec security associations can be set up more quickly than with shorter lifetimes.

Step 10 Set the **Authentication Method** to use between the two peers.

- **Preshared Key**—Preshared keys allow for a secret key to be shared between two peers and to be used by IKE during the authentication phase. If one of the participating peers is not configured with the same preshared key, the IKE SA cannot be established.
- **Certificate**—When you use Certificates as the authentication method for VPN connections, peers obtain digital certificates from a CA server in your PKI infrastructure, and trade them to authenticate each other.

Note In a VPN topology that supports IKEv1, the **Authentication Method** specified in the chosen IKEv1 Policy object becomes the default in the IKEv1 **Authentication Type** setting. These values must match, otherwise, your configuration will error.

Step 11 Click **Save**
The new IKEv1 policy is added to the list.

Configure IKEv2 Policy Objects

Use the IKEv2 policy dialog box to create, delete, and edit an IKEv2 policy object. These policy objects contain the parameters required for IKEv2 policies.

Procedure

Step 1 Choose **Objects** > **Object Management** and then **VPN** > **IKEv2 Policy** from the table of contents.

Previously configured policies are listed including system defined defaults. Depending on your level of access, you may **Edit** ( ), **View** ( ), or **Delete** ( ) a policy.

Step 2 Choose **Add** ( )**Add IKEv2 Policy** to create a new policy.

Step 3 Enter a **Name** for this policy.

The name of the policy object. A maximum of 128 characters is allowed.

Step 4 Enter a **Description** for this policy.

A description of the policy object. A maximum of 1024 characters is allowed.

Step 5 Enter the **Priority**.

The priority value of the IKE proposal. The priority value determines the order of the IKE proposals compared by the two negotiating peers when attempting to find a common security association (SA). If the remote IPsec peer does not support the parameters selected in your first priority policy, it tries to use the parameters defined in the next lowest priority policy. Valid values range from 1 to 65535. The lower the number, the higher the priority. If you leave this field blank, Management Center assigns the lowest unassigned value starting with 1, then 5, then continuing in increments of 5.

Step 6 Set the **Lifetime**of the security association (SA), in seconds. You can specify a value from 120 to 2,147,483,647 seconds. The default is 86400.

When the lifetime is exceeded, the SA expires and must be renegotiated between the two peers. Generally, the shorter the lifetime (up to a point), the more secure your IKE negotiations. However, with longer lifetimes, future IPsec security associations can be set up more quickly than with shorter lifetimes.

Step 7 Choose the **Integrity Algorithms** portion of the Hash Algorithm used in the IKE policy. The Hash Algorithm creates a Message Digest, which is used to ensure message integrity.

When deciding which encryption and Hash Algorithms to use for the IKEv2 proposal, your choice is limited to algorithms supported by the managed devices. For an extranet device in the VPN topology, you must choose the algorithm that matches both peers. Select all the algorithms that you want to allow in the VPN.For a full explanation of the options, see Deciding Which Hash Algorithms to Use, on page 843.

Step 8 Choose the **Encryption Algorithm** used to establish the Phase 1 SA for protecting Phase 2 negotiations.

When deciding which encryption and Hash Algorithms to use for the IKEv2 proposal, your choice is limited to algorithms supported by the managed devices. For an extranet device in the VPN topology, you must choose the algorithm that matches both peers. Select all the algorithms that you want to allow in the VPN. For a full explanation of the options, see Deciding Which Encryption Algorithm to Use, on page 843.

Step 9 Choose the **PRF Algorithm**.

The pseudorandom function (PRF) portion of the Hash Algorithm used in the IKE policy. In IKEv1, the Integrity and PRF algorithms are not separated, but in IKEv2, you can specify different algorithms for these elements. Select all of the algorithms that you want to allow in the VPN. For a full explanation of the options, see Deciding Which Hash Algorithms to Use, on page 843.

Step 10 Select and **Add** a **DH Group**.

The Diffie-Hellman group used for encryption. A larger modulus provides higher security but requires more processing time. The two peers must have a matching modulus group. Select the groups that you want to allow in the VPN. For a full explanation of the options, see Deciding Which Diffie-Hellman Modulus Group to Use, on page 844.

Step 11 Click **Save**
If a valid combination of choices has been selected the new IKEv2 policy is added to the list. If not, errors are displayed and you must make changes accordingly to successfully save this policy.

Firepower Threat Defense IPsec Proposals

IPsec Proposals (or Transform Sets) are used when configuring VPN topologies. During the IPsec security association negotiation with ISAKMP, the peers agree to use a particular proposal to protect a particular data flow. The proposal must be the same for both peers.

There are separate IPsec proposal objects based on the IKE version, IKEv1, or IKEv2:

- When you create an IKEv1 IPsec Proposal (Transform Set) object, you select the mode in which IPsec operates, and define the required encryption and authentication types. You can select single options for the algorithms. If you want to support multiple combinations in a VPN, create multiple IKEv1 IPsec Proposal objects.
- When you create an IKEv2 IPsec Proposal object, you can select all of the encryption and Hash Algorithms allowed in a VPN. During IKEv2 negotiations, the peers select the most appropriate options that each support.

The Encapsulating Security Protocol (ESP) is used for both IKEv1 and IKEv2 IPsec Proposals. It provides authentication, encryption, and antireplay services. ESP is IP protocol type 50.

Note We recommend using both encryption and authentication on IPsec tunnels.

Configure IKEv1 IPsec Proposal Objects

Procedure

Step 1 Choose **Objects** > **Object Management** and then **VPN** > **IPsec IKev1 Proposal** from the table of contents.

Previously configured Proposals are listed including system defined defaults. Depending on your level of access, you may **Edit** ( ), View **View** ( ), or **Delete** ( ) a Proposal.

Step 2 Choose **Add** ( )**Add IPsec IKEv1 Proposal** to create a new Proposal.

Step 3 Enter a **Name** for this Proposal

The name of the policy object. A maximum of 128 characters is allowed.

Step 4 Enter a **Description** for this Proposal.

A description of the policy object. A maximum of 1024 characters is allowed.

Step 5 Choose the **ESP Encryption** method. The Encapsulating Security Protocol (ESP) encryption algorithm for this Proposal.

For IKEv1, select one of the options. When deciding which encryption and Hash Algorithms to use for the IPsec proposal, your choice is limited to algorithms supported by the devices in the VPN. For a full explanation of the options, see Deciding Which Encryption Algorithm to Use, on page 843.

Step 6 Select an option for **ESP Hash**.

For a full explanation of the options, see Deciding Which Hash Algorithms to Use, on page 843.

Step 7 Click **Save**
The new Proposal is added to the list.

Configure IKEv2 IPsec Proposal Objects

Procedure

Step 1 Choose **Objects** > **Object Management** and then **VPN** > **IKEv2 IPsec Proposal** from the table of contents.

Previously configured Proposals are listed including system defined defaults. Depending on your level of access, you may Edit **Edit** ( ), View **View** ( ), or Delete **Delete** ( ) a Proposal.

Step 2 Choose **Add** ( )**Add IKEv2 IPsec Proposal** to create a new Proposal.

Step 3 Enter a **Name** for this Proposal

The name of the policy object. A maximum of 128 characters is allowed.

Step 4 Enter a **Description** for this Proposal.

A description of the policy object. A maximum of 1024 characters is allowed.

Step 5 Choose the **ESP Hash** method, the hash or integrity algorithm to use in the Proposal for authentication.

For IKEv2, select all the options you want to support for **ESP Hash**. For a full explanation of the options, see Deciding Which Hash Algorithms to Use, on page 843.

Step 6 Choose the **ESP Encryption** method. The Encapsulating Security Protocol (ESP) encryption algorithm for this Proposal.

For IKEv2, click Select to open a dialog box where you can select all of the options you want to support. When deciding which encryption and Hash Algorithms to use for the IPsec proposal, your choice is limited to algorithms supported by the devices in the VPN. For a full explanation of the options, see Deciding Which Encryption Algorithm to Use, on page 843.

Step 7 Click **Save**
The new Proposal is added to the list.

Address Pools

You can configure IP address pools for both IPv4 and IPv6 that can be used for the Diagnostic interface with clustering.

You can use this object with Firepower Threat Defense devices.

Procedure

Step 1 Select **Objects** > **Object Management** > **Address Pools** > **IPv4 Pools**.

Step 2 Click **Add IPv4 Pools**, and configure the following fields:

- **Name**—Enter the name of the address pool. It can be up to 64 characters
- **Description**—Add an optional description for this pool.
- **IP Address**—Enter a range of addresses available in the pool. Use dotted decimal notation and a dash between the beginning and the end address, for example: 10.10.147.100-10.10.147.177.
- **Mask**—Identifies the subnet on which this IP address pool resides.
- **Allow Overrides**—Check this check box to enable object overrides. Click the expand arrow to show the **Overrides** table. You can add a new override by clicking **Add**. See Object Overrides, on page 399 for more information.

Step 3 Click **Save**.

Step 4 Click **Add IPv6 Pools**, and configure the following fields:

- **Name**—Enter the name of the address pool. It can be up to 64 characters
- **Description**—Add an optional description for this pool.
- **IPv6 Address**—Enter the first IP address available in the configured pool and the prefix length in bits. For example: 2001:DB8::1/64.

- **Number of Addresses**—Identifies the number of IPv6 addresses, starting at the Starting IP Address, that are in the pool.
- **Allow Overrides**—Check this check box to enable overrides. Click the expand arrow to show the **Overrides** table. You can add a new override by clicking **Add**. See Object Overrides, on page 399 for more information.

Step 5 Click **Save**.

FlexConfig Objects

Use FlexConfig policy objects in FlexConfig policies to provide customized configuration of features on Firepower Threat Defense devices that you cannot otherwise configure using Firepower Management Center. For more information on FlexConfig policies, see FlexConfig Policy Overview, on page 575.

You can configure the following types of objects for FlexConfig.

Text Objects

Text objects define free-form text strings that you use as variables in a FlexConfig object. These objects can have single values or be a list of multiple values.

There are several predefined text objects that are used in the predefined FlexConfig objects. If you use the associated FlexConfig object, you simply need to edit the contents of the text object to customize how the FlexConfig object configures a given device. When editing a predefined object, it is in general a better option to create device overrides for each device you are configuring, rather than directly change the default values of these objects. This helps avoid unintended consequences if another user wants to use the same FlexConfig object for a different set of devices.

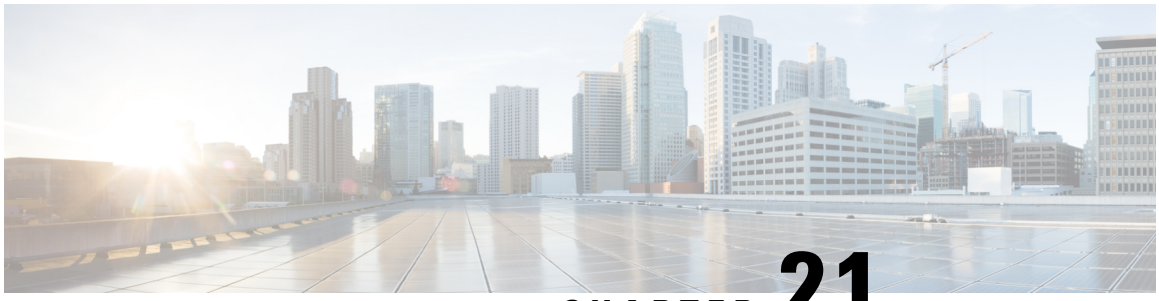
For information on configuring text objects, see Configure FlexConfig Text Objects, on page 597.

FlexConfig Objects

FlexConfig Objects include device configuration commands, variables, and scripting language instructions. During configuration deployment, these instructions are processed to create a sequence of configuration commands with customized parameters to configure specific features on the target devices.

These instructions are either configured before (prepended) the system configures features defined in regular Firepower Management Center policies and settings, or after (appended). Any FlexConfig that depends on Firepower Management Center-configured objects (for example, a network object) must be appended to the configuration deployment, or the needed objects would not be configured before the FlexConfig needed to refer to the objects.

For more information on configuring FlexConfig objects, see Configure FlexConfig Objects, on page 594.



CHAPTER 21

Firepower Threat Defense Certificate-Based Authentication

- Requirements and Prerequisites for FTD Certificate-Based Authentication, on page 481
- Firepower Threat Defense VPN Certificate Guidelines and Limitations, on page 482
- Managing Firepower Threat Defense VPN Certificates, on page 482
- Installing a Certificate Using Self-Signed Enrollment , on page 483
- Installing a Certificate Using SCEP Enrollment, on page 484
- Installing a Certificate Using Manual Enrollment, on page 485
- Installing a Certificate by Importing a PKCS12 File, on page 486
- Troubleshooting Firepower Threat Defense VPN Certificates, on page 487

Requirements and Prerequisites for FTD Certificate-Based Authentication

Model Support

FTD

Supported Domains

Any

User Roles

Admin

Network Admin

Firepower Threat Defense VPN Certificate Guidelines and Limitations

- When a certificate enrollment object is associated with and then installed on a device, the process of certificate enrollment starts immediately. The process is automatic for self-signed and SCEP enrollment types, meaning it does not require any additional administrator action. Manual certificate enrollment and importing a PKCS12 file requires extra administrator action.
- When the certificate enrollment is complete, a trustpoint exists on the device with the same name as the certificate enrollment object. Use this trustpoint in the configuration of your VPN Authentication Method.
- Firepower Threat Defense currently supports RSA keys only, not ECDSA, even though the choice appears in the user interface.
- Since Firepower Threat Defense VPNs are not supported in a clustered environment, PKI is also not supported in a clustered environment.
- Firepower Threat Defense devices support certificate enrollment using Microsoft Certificate Authority(CA) Service, and CA Services provided on Cisco Adaptive Security Appliances(ASA) and Cisco IOS Router.
- Firepower Threat Defense devices cannot be configured as a certificate authority (CA).

Guidlelines for Certificate Management Across Domains and Devices

- The Firepower Threat Defense devices support, and have been verified, certificate enrollment using, Microsoft CA Service, and CA Services provided on Cisco Adaptive Security Appliances and Cisco IOS Routes.
- Certificate enrollment can be done in a child or parent domain.
- When enrollment is done from a parent domain, the certificate enrollment object also needs to be in the same domain. If the trustpoint on a device is overridden in the child domain, the overridden value will be deployed on the device.
- When the certificate enrollment is done on a device in a leaf domain, the enrollment will not be visible to the parent domain another child domain.
- When a leaf domain is deleted, certificate enrollments on the contained devices have to be removed.
- Once a device has certificates enrolled in one domain, it will not be allowed to be enrolled in any other domain. However, the certificates can be viewed in the the other domain.
- When a device is moved from one domain to another domain, or from no domain into a domain, certificate enrollments on that device have to be removed and reconfigured in the new domain. You will receive an alert to delete the enrollments on these devices.

Managing Firepower Threat Defense VPN Certificates

See PKI Infrastructure and Digital Certificates , on page 845 for an introduction to Digital Certificates.

See Certificate Enrollment Objects, on page 456 for a description of the objects used to enroll and obtain certificates on managed devices.

Procedure

Step 1 Select **Devices** > **Certificates**.

You can see the following columns for each device listed on this screen:

- **Name**—Lists the devices that already have trustpoints associated with them. Expand the device to see the list of associated trustpoints.
- **Enrollment Type**—Displays the type of enrollment used for a trustpoint.
- The additional columns provide the status of the **CA Certificate** and **Identity Certificate**. In each column, the certificate contents, when `Available`, can be viewed by clicking the magnifying glass.

  The values of these columns depend on the enrollment type and change during the process of enrollment. The CA Certificate can be `Available`, `Not Available`, and `Not Applicable`. The Identity Certificate status can be `Available`, `Pending`, and `Available and Pending` during a refresh.
- Refresh (circling arrows) a certificate on a managed device. Refreshing a certificate would synchronize the Firepower Threat Defense device certificate status to the Firepower Management Center.
- Delete (trash can) a configured certificate.

Step 2 Choose **(+) Add** > **Add New Certificate** to associate and install an enrollment object on a device. Continue based on the type of enrollment.

Note When a certificate enrollment object is associated with and then installed on a device, the process of certificate enrollment starts immediately. The process is automatic for self-signed and SCEP enrollment types, meaning it does not require any additional administrator action. Manual certificate enrollment and importing a PKCS12 file requires extra administrator action.

Related Topics

Installing a Certificate Using Self-Signed Enrollment , on page 483
Installing a Certificate Using SCEP Enrollment, on page 484
Installing a Certificate Using Manual Enrollment, on page 485
Installing a Certificate by Importing a PKCS12 File, on page 486

Installing a Certificate Using Self-Signed Enrollment

Procedure

Step 1 On the **Devices** > **Certificates** screen, choose **Add** > **Add New Certificate** to open the **Add New Certificate** dialog.

Step 2 Choose a device from the **Device** drop down list.

Step 3 Associate a certificate enrollment object with this device in one of the following ways:

- Choose a Certificate Enrollment Object of the appropriate type from the drop-down list.

- Click **(+)**, to add a new Certificate Enrollment Object, see Adding Certificate Enrollment Objects, on page 458.

Step 4 Press **Install**, to start the Self Signed, automatic, enrollment process.

For self signed enrollment type trustpoints, the **CA Certificate** status will always be `NotApplicable` since the managed device is acting as its own CA and does not need a CA certificate to generate its own Identity Certificate.

The **Identity Certificate** will go from InProgress to Available as the device creates its own self signed identity certificate.

Step 5 Click the magnifying glass to view the self-signed Identity Certificate created for this device.

What to do next

When the certificate enrollment is complete, a trustpoint exists on the device with the same name as the certificate enrollment object. Use this trustpoint in the configuration of your VPN Authentication Method.

Installing a Certificate Using SCEP Enrollment

Before you begin

Note Using SCEP enrollment establishes a direct connection between the managed device and the CA server. So be sure your device is connected to the CA server before beginning the enrollment process.

Procedure

Step 1 Press **Install**, to start the automatic enrollment process.

Step 2 On the **Devices > Certificates** screen, choose **Add > Add New Certificate** to open the **Add New Certificate** dialog.

Step 3 Choose a device from the **Device** drop down list.

Step 4 Associate a certificate enrollment object with this device in one of the following ways:

- Choose a Certificate Enrollment Object of the appropriate type from the drop-down list.
- Click **(+)**, to add a new Certificate Enrollment Object, see Adding Certificate Enrollment Objects, on page 458.

Step 5 Press **Install**, to start the automatic enrollment process.

For SCEP enrollment type trustpoints, the **CA Certificate** status will transition from InProgress to Available as the CA Certificate is obtained from the CA server and installed on the device.

The **Identity Certificate** will go from `InProgress` to `Available` as the device obtains its identity certificate using SCEP from the specified CA.

Note The SCEP certificate enrollment could possibly fail because of some error messages. For example:

```
Error:
crypto ca authenticate scep1 nointeractive:[error]:ERROR:receiving Certificate
Authority certificate: status = FAIL, cert length = 0
Possible
```

Some recommendations to overcome this situaton:

- Make sure that the connectivity to the SCEP Server from the FTD - Route is added to the SCEP Server.
- If the SCEP Server is referred with hostname/FQDN, configure DNS Server using FlexConfig object.
- Ensure that the SCEP Server and the FTD device are in time-sync by configuring the same NTP Server.

Step 6 Click the magnifying glass to view the Identity Certificate created and installed on this device.

What to do next

When the certificate enrollment is complete, a trustpoint exists on the device with the same name as the certificate enrollment object. Use this trustpoint in the configuration of your VPN Authentication Method.

Installing a Certificate Using Manual Enrollment

Procedure

Step 1 Press **Install** to start the enrollment process.

Step 2 On the **Devices > Certificates** screen, choose **Add > Add New Certificate** to open the **Add New Certificate** dialog.

Step 3 Choose a device from the **Device** drop down list.

Step 4 Associate a certificate enrollment object with this device in one of the following ways:

- Choose a Certificate Enrollment Object of the appropriate type from the drop-down list.
- Click **(+)**, to add a new Certificate Enrollment Object, see Adding Certificate Enrollment Objects, on page 458.

Step 5 Press **Import**, to initiate the manual enrollment process.

The **CA Certificate** status will go from `InProgress` to `Available` as the Firepower Management Center installs the CA certificate (provided in the enrollment object) on the managed device, authenticates the CA Server, and creates a trustpoint on the managed device.

The **Identity Certificate** status could throw a warning message that CSR generation and identity certificate import is pending.

Step 6 Execute the appropriate activity with your PKI CA Server to obtain an identity certificate.

a) Click **Identity Certificate** warning to view and copy the CSR.

b) Execute the appropriate activity with your PKI CA Server to obtain an identity certificate using this CSR.

This activity is completely independent of the Firepower Management Center or the managed device. When complete, you will have an Identity Certificate for the managed device. You can copy it or place it in a file.

c) To finish the manual process, install the obtained identity certificate onto the managed device.

Return to the Firepower Management Center dialog to paste the Identity Certificate into its field. Or, select **Browse** to choose the identity certificate file.

Step 7 Select **Import** to import the Identity Certificate.

The Identity Certificate status will be `Available` when the import complete.

Step 8 Click the magnifying glass to view the **Identity Certificate** for this device.

What to do next

When the certificate enrollment is complete, a trustpoint exists on the device with the same name as the certificate enrollment object. Use this trustpoint in the configuration of your VPN Authentication Method.

Installing a Certificate by Importing a PKCS12 File

Procedure

Step 1 Go to **Devices** > **Certificates** screen, then click **+ Add** > **Import PKCS12 File** to open the **Import PKCS12 File** dialog.

Step 2 Choose a pre-configured managed device from the **Device** drop down list.

Step 3 Specify a **Certificate Enrollment** type of **PKCS12**.

Step 4 Select **Browse** to find and choose your PKCS#12 Certificate file.

Step 5 Enter the **Passphrase** for decryption.

Step 6 Press **Add**.

For file import, the CA Certificate and Identity Certificate status will go from `In Progress` to `Available` as it installs the PKCS12 file on the device.

Step 7 Once `Available`, click the magnifying glass to view the Identity Certificate for this device.

What to do next

The certificate (trustpoint) on the managed device is named the same as the PKCS#12 file. Use this certificate in your VPN authentication configuration.

Troubleshooting Firepower Threat Defense VPN Certificates

See Firepower Threat Defense VPN Certificate Guidelines and Limitations, on page 482 to determine if variations in your certificate enrollment environment may be causing a problem. Then consider the following:

- Ensure there is a route to the CA Server from the device.

  If the CA Server's host name is given in the Enrollment Object, use Flex Config to configure DNS appropriately to reach the server. Alternatively, use the IP Address of the CA Server.

- If you are using a Microsoft 2012 CA Server, the default IPsec Template is not accepted by the managed device and must be changed.

  To configure a working template, follow these steps as you use MS CA documentation as a reference.

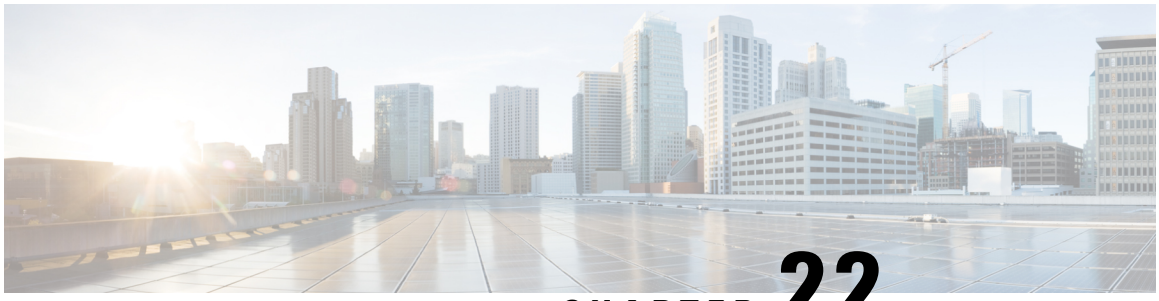
  1. Duplicate the IPsec (Offline Request) template.
  2. In **Extensions > Application policies**, select *IP security end system*, instead of the *IP security IKE intermediate*.
  3. Set the permissions and the template name.
  4. Add the new template and change the registry settings to reflect the new template name.



PART V

Classic Device Configuration Basics

- Classic Device Management Basics, on page 491
- IPS Device Deployments and Configuration, on page 503



CHAPTER 22

Classic Device Management Basics

The following topics describe how to manage Classic devices (7000 and 8000 Series/ASA with FirePOWER Services/NGIPSv) in the Firepower System:

- Requirements and Prerequisites for Classic Device Management, on page 491
- Remote Management Configuration (Classic Devices), on page 491
- Interface Configuration Settings, on page 494

Requirements and Prerequisites for Classic Device Management

Model Support

Classic models as indicated in the procedures.

Supported Domains

Leaf unless indicated otherwise.

User Roles

- Admin
- Network Admin

Remote Management Configuration (Classic Devices)

All Devices Except 7000 and 8000 Series

For information on configuring remote management for devices that use Classic licenses, see the quick start guide for your device.

7000 and 8000 Series Devices

Configure remote management of a 7000 or 8000 Series device using its local web interface, before you register the device to the FMC.

Before you can manage a Firepower device, you must set up a two-way, SSL-encrypted communication channel between the device and the Firepower Management Center. The appliances use the channel to share configuration and event information. High availability peers also use the channel, which is by default on port 8305/tcp.

To enable communications between two appliances, you must provide a way for the appliances to recognize each other, as follows:

- The hostname or IP address of the appliance with which you are trying to establish communication.

  In NAT environments, even if the other appliance does not have a routable address, you must provide a hostname or an IP address either when you are configuring remote management, or when you are adding the managed appliance.

- A self-generated alphanumeric registration key up to 37 characters in length that identifies the connection.
- An optional unique alphanumeric NAT ID that can help establish communications in a NAT environment.

  The NAT ID *must* be unique among all NAT IDs used to register managed appliances.

Configuring Remote Management on a Managed Device

This procedure applies to 7000 & 8000 Series devices.

Procedure

Step 1 On the web interface for the device you want to manage, choose **System** > **Integration** > **Remote Management**.

Step 2 Click **Remote Management**, if it is not already displaying.

Step 3 Click **Add Manager**.

Step 4 In the **Management Host** field, enter one of the following for the Firepower Management Center that you want to use to manage this appliance:

- The IP address
- The fully qualified domain name or the name that resolves through the local DNS to a valid IP address (that is, the host name)

Caution Use a host name rather than an IP address if your network uses DHCP to assign IP addresses.

In a NAT environment, you do not need to specify an IP address or host name here if you plan to specify it when you add the managed appliance. In this case, the Firepower System uses the NAT ID you will provide later to identify the remote manager on the managed appliance's web interface.

Step 5 In the **Registration Key** field, enter the registration key that you want to use to set up communications between appliances.

Step 6 For NAT environments, in the **Unique NAT ID** field, enter a **unique** alphanumeric NAT ID that you want to use to set up communications between appliances.

Step 7 Click **Save**.

What to do next

- Wait until the appliances confirm that they can communicate with each other and the Pending Registration status appears.
- Add this device to the Firepower Management Center; see Add a Device to the FMC, on page 229.

Editing Remote Management on a Managed Device

This procedure applies to 7000 & 8000 Series devices.

When editing a remote manager, note that:

- The **Host** field specifies the fully qualified domain name or the name that resolves through the local DNS to a valid IP address (that is, the host name).
- The **Name** field specifies the display name of the managing appliance, which is used only within the context of the Firepower System. Entering a different display name does not change the host name for the managing device.

Procedure

Step 1 On the web interface for the device, choose **System** > **Integration**.

Step 2 Click **Remote Management**, if it is not already displaying.

Step 3 You can:

- Disable remote management — Click the slider next to the manager to enable or disable it. Disabling management blocks the connection between the Firepower Management Center and the device, but does **not** delete the device from the Firepower Management Center. If you no longer want to manage a device, see Delete a Device from the FMC, on page 231.
- Edit manager information — Click **Edit** ( ) next to the manager you want to modify, modify the **Name** and **Host** fields, and click **Save**.

Changing the Management Port

Appliances communicate using a two-way, SSL-encrypted communication channel, which by default is on port 8305.

Although Cisco *strongly* recommends that you keep the default setting, you can choose a different port if the management port conflicts with other communications on your network. Usually, changes to the management port are made during installation.

Caution If you change the management port, you must change it for all appliances in your deployment that need to communicate with each other.

You must perform this task in the global domain.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **Management Interfaces**.

Step 3 In the **Shared Settings** section, enter the port number that you want to use in the **Remote Management Port** field.

Step 4 Click **Save**.

What to do next

Repeat this procedure for every appliance in your deployment that must communicate with this appliance.

Interface Configuration Settings

The Interfaces page of the appliance editor displays detailed interface configuration information. The page is composed of the physical hardware view and the interfaces table view, which allow you to drill down to configuration details. You can add and edit interfaces from this page.

The Physical Hardware View

The top of the Interfaces page provides a graphical representation of the physical hardware view of a 7000 or 8000 Series device.

Use the physical hardware view to:

- view a network module's type, part number, and serial number
- select an interface in the interfaces table view
- open an interface editor
- view the name of the interface, the type of interface, whether the interface has link, the interface's speed setting, and whether the interface is currently in bypass mode
- view the details about an error or warning

The Interfaces Page

The interfaces page lists all the available interfaces you have on a device. The table includes an expandable navigation tree you can use to view all configured interfaces. You can click the arrow icon next to an interface to collapse or expand the interface to hide or view its subcomponents. The interfaces table view also provides summarized information about each interface.

Field	Description
Name	Each interface type is represented by a unique icon that indicates its type and link state (if applicable). You can hover your pointer over the name or the icon to view a tooltip with additional information. The interface icons are described in Interface Icons, on page 496. The icons use a badging convention to indicate the current link state of the interface, which may be one of three states: • **Error** • **Fault** • **Not available** Logical interfaces have the same link state as their parent physical interface. ASA FirePOWER modules do not display link state. Note that disabled interfaces are represented by semi-transparent icons. Interface names, which appear to the right of the icons, are auto-generated with the exception of hybrid and ASA FirePOWER interfaces, which are user-defined. Note that for ASA FirePOWER interfaces, the system displays only interfaces that are enabled, named, and have link. Physical interfaces display the name of the physical interface. Logical interfaces display the name of the physical interface and the assigned VLAN tag. ASA FirePOWER interfaces display the name of the security context and the name of the interface if there are multiple security contexts. If there is only one security context, the system displays only the name of the interface.
Security Zone	The security zone where the interface is assigned. To add or edit a security zone, click **Edit** ().
Used by	The inline set, virtual switch, or virtual router where the interface is assigned.
MAC Address	The MAC address displayed for the interface when it is enabled for switched and routed features. For NGIPSv devices, the MAC address is displayed so that you can match the network adapters configured on your device to the interfaces that appear on the Interfaces page.
IP Addresses (7000/8000 series only)	IP addresses assigned to the interface. Hover your pointer over an IP address to view whether it is active. Inactive IP addresses are also grayed out.

Interface Icons

Table 60: Interface Icon Types and Descriptions

Icon	Interface Type	Description	See
Physical	Physical	Unconfigured physical interface.	Configuring Physical Switched Interfaces, on page 1107 or Configuring Physical Routed Interfaces, on page 1117
Passive	Passive	Sensing interface configured to analyze traffic in a passive deployment.	Configuring Passive Interfaces, on page 505
Inline	Inline	Sensing interface configured to handle traffic in an inline deployment.	Configuring Inline Interfaces, on page 508
Switched	Switched	Interface configured to switch traffic in a Layer 2 deployment.	Switched Interface Configuration, on page 1105
Routed	Routed	Interface configured to route traffic in a Layer 3 deployment.	Routed Interfaces, on page 1116
Aggregate	Aggregate	Multiple physical interfaces configured as a single logical link.	About Aggregate Interfaces, on page 1147
Aggregate Switched	Aggregate Switched	Multiple physical interfaces configured as a single logical link in a Layer 2 deployment.	Adding Aggregate Switched Interfaces, on page 1153
Aggregate Routed	Aggregate Routed	Multiple physical interfaces configured as a single logical link in a Layer 3 deployment.	Adding Aggregate Routed Interfaces, on page 1155
Hybrid	Hybrid	Logical interface configured to bridge traffic between a virtual router and a virtual switch.	Logical Hybrid Interfaces, on page 1161
HA	HA	Interface configured on each member of a pair of devices to act as a redundant communications channel between the devices. Also called a high availability link interface.	Configuring HA Link Interfaces, on page 498
ASA FirePOWER	ASA FirePOWER	Interface configured on an ASA device with the ASA FirePOWER module installed.	Managing Cisco ASA FirePOWER Interfaces, on page 500

Using the Physical Hardware View

This task applies to 7000 & 8000 Series devices.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click edit **Edit** ( ) next to the device you want to manage.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Use the graphical interface to:

- Choose — If you want to choose an interface, click interface. The system highlights the related entry in the interface table.
- Edit — If you want to open an interface editor, double-click interface.
- View error or warning information — If you want to view the details about an error or warning, hover your cursor over the affected port on the network module.
- View interface information — If you want to view the name of the interface, the type of interface, whether the interface has link, the interface's speed setting, and whether the interface is currently in bypass mode, hover your cursor over the interface.
- View network module information — If you want to view a network module's type, part number, and serial number, hover your cursor over the dark circle in the lower left corner of the network module.

Configuring Sensing Interfaces

You can configure the sensing interfaces of a managed device, according to your Firepower deployment, from the Interfaces page of the appliance editor. Note that you can only configure a total of 1024 interfaces on a managed device.

Note The Firepower Management Center does not display ASA interfaces when the ASA FirePOWER is deployed in SPAN port mode.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to configure an interface, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Edit** ( ) next to the interface you want to configure.

Step 4 Use the interface editor to configure the sensing interface:

- HA Link — If you want an interface configured on each member of a high-availability pair of 7000/8000 series devices to act as a redundant communications channel between the devices; also called a high availability link interface, click **HA Link** and proceed as described in Configuring HA Link Interfaces, on page 498.

- Inline — If you want an interface configured to handle traffic in an inline deployment, click **Inline** and proceed as described in Configuring Inline Interfaces, on page 508.
- Passive — If you want an interface configured to analyze traffic in a passive deployment, click **Passive** and proceed as described in Configuring Passive Interfaces, on page 505.
- Routed — If you want an interface configured to route traffic in a 7000/8000 series Layer 3 deployment, click **Routed** and proceed as described in Routed Interfaces, on page 1116.
- Switched — If you want an interface configured to switch traffic in a 7000/8000 series Layer 2 deployment, click **Switched** and proceed as described in Switched Interface Configuration, on page 1105.

Step 5 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configuring HA Link Interfaces

After you establish a 7000 or 8000 Series device high-availability pair, you should configure a physical interface as a high availability (HA) link interface. This link acts as a redundant communications channel for sharing health information between the paired devices. When you configure an HA link interface on one device, you automatically configure an interface on the second device. You must configure both HA links on the same broadcast domain.

Dynamic NAT relies on dynamically allocating IP addresses and ports to map to other IP addresses and ports. Without an HA link, these mappings are lost in a failover, causing all translated connections to fail as they are routed through the now-active device in the high-availability pair.

Similarly, 7000 or 8000 Series devices with high-availability state sharing, dynamic NAT, or VPN require an HA link interface.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the peer where you want to configure the HA link interface, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Next to the interface you want to configure as a HA link interface, click **Edit** ( ).

Step 4 Click **HA Link**.

Step 5 Check the **Enabled** check box.

If you clear the check box, the system administratively takes down the interface, disabling it.

Step 6 From the **Mode** drop-down list, choose an option to designate the link mode, or choose **Autonegotiation** to specify that the interface is configured to autonegotiate speed and duplex settings.

Step 7 From the **MDI/MDIX** drop-down list, choose an option to designate whether the interface is configured for MDI (medium dependent interface), MDIX (medium dependent interface crossover), or Auto-MDIX.

Normally, MDI/MDIX is set to **Auto-MDIX**, which automatically handles switching between MDI and MDIX to attain link.

Step 8 Enter a maximum transmission unit (MTU) in the **MTU** field.

The range of MTU values can vary depending on the model of the managed device and the interface type. See MTU Ranges for 7000 and 8000 Series Devices and NGIPSv, on page 500 for more information.

Caution Changing the highest MTU value among all non-management interfaces on the device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all non-management interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. See Snort® Restart Traffic Behavior, on page 350 for more information.

Step 9 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Snort® Restart Scenarios, on page 348
MTU Ranges for 7000 and 8000 Series Devices and NGIPSv, on page 500

Disabling Interfaces

You can disable an interface by setting the interface type to **None**. Disabled interfaces appear grayed out in the interface list.

This procedure applies to NGIPSv and 7000 & 8000 Series devices.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to disable the interface, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Next to the interface you want to disable, click **Edit** ( ).

Step 4 Click **None**.

Step 5 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Managing Cisco ASA FirePOWER Interfaces

When editing an ASA FirePOWER interface, you can configure only the interface's security zone from the Firepower Management Center.

You fully configure ASA FirePOWER interfaces using the ASA-specific software and CLI. If you edit an ASA FirePOWER and switch from multiple context mode to single context mode (or visa versa), the ASA FirePOWER renames all of its interfaces. You must reconfigure all Firepower System security zones, correlation rules, and related configurations to use the updated ASA FirePOWER interface names. For more information about ASA FirePOWER interface configuration, see the ASA documentation.

Note You cannot change the type of ASA FirePOWER interface, nor can you disable the interface from the Firepower Management Center.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to edit the interface, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Interfaces** if it is not already displaying.

Step 4 Next to the interface you want to edit, click **Edit** ( ).

Step 5 Choose an existing security zone from the **Security Zone** drop-down list, or choose **New** to add a new security zone.

Step 6 Click **Save** to configure the security zone.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

MTU Ranges for 7000 and 8000 Series Devices and NGIPSv

Changing the highest MTU value among all non-management interfaces on the device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all non-management interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. See Snort® Restart Traffic Behavior, on page 350 for more information.

Note The system trims 18 bytes from the configured MTU value. Do not set the IPv4 MTU lower than 594 or the IPv6 MTU lower than 1298.

Platform	MTU Range
7000 & 8000 Series	576-9234 (management interface) 576-10172 (inline sets, passive interface) 576-9922 (all others)
NGIPSv	576-9018 (all interfaces, inline sets)

Related Topics

About the MTU, on page 642

Synchronizing Security Zone Object Revisions

When you update a security zone object, the system saves a new revision of the object. As a result, if you have managed devices in the same security zone that have different revisions of the security zone object configured in the interfaces, you may log what appear to be duplicate connections.

If you notice duplicate connection reporting, you can update all managed devices to use the same revision of the object.

This procedure applies to NGIPSv and 7000 & 8000 Series devices.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to update the security zone selection, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

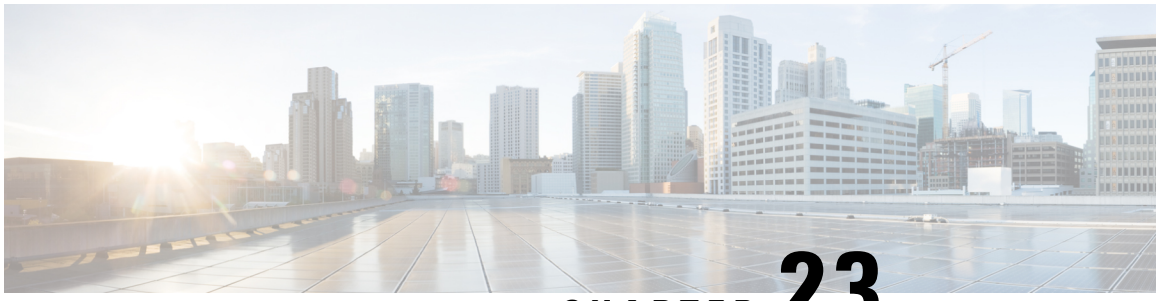
Step 3 For each interface logging duplicate connection events, change the **Security Zone** to another zone, click **Save**, then change it back to the desired zone, and click **Save** again.

Step 4 Repeat steps 2 through 3 for each device logging duplicate events. You must edit all devices before you continue.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Caution Do not deploy configuration changes to any device until you edit the zone setting for interfaces on *all* devices you want to sync. You must deploy to all managed devices at the same time.



CHAPTER 23

IPS Device Deployments and Configuration

The following topics describe how to configure your device in an IPS deployment:

- Introduction to IPS Device Deployment and Configuration, on page 503
- License Requirements for IPS Device Deployment, on page 503
- Requirements and Prerequisites for IPS Device Deployment, on page 503
- Passive IPS Deployments, on page 504
- Inline IPS Deployments, on page 506

Introduction to IPS Device Deployment and Configuration

You can configure your device in either a passive or inline IPS deployment. In a passive deployment, you deploy the system out of band from the flow of network traffic. In an inline deployment, you configure the system transparently on a network segment by binding two ports together.

License Requirements for IPS Device Deployment

FTD License

Threat

Classic License

Protection

Requirements and Prerequisites for IPS Device Deployment

Model Support

Any.

Supported Domains

Leaf.

User Roles

- Admin
- Network Admin

Passive IPS Deployments

In a passive IPS deployment, the Firepower System monitors traffic flowing across a network using a switch SPAN (or mirror) port. The SPAN port allows for traffic to be copied from other ports on the switch. This provides the system visibility within the network without being in the flow of network traffic. When configured in a passive deployment, the system cannot take certain actions such as blocking or shaping traffic. Passive interfaces receive all traffic unconditionally, and no traffic received on these interfaces is retransmitted. Passive interfaces support both local SPAN and remote SPAN (RSPAN) traffic.

Note Outbound traffic includes flow control packets. Because of this, passive interfaces on your appliances may show outbound traffic and, depending on your configuration, generate events; this is expected behavior.

Passive Interfaces on the Firepower System

You can configure one or more physical ports on a managed device as passive interfaces.

When you enable a passive interface to monitor traffic, you designate mode and MDI/MDIX settings, which are available only for copper interfaces. Interfaces on 8000 Series appliances do not support half-duplex options.

When you disable a passive interface, users can no longer access it for security purposes.

The range of MTU values can vary depending on the model of the managed device and the interface type.

Caution Changing the highest MTU value among all non-management interfaces on the device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all non-management interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. See Snort® Restart Traffic Behavior, on page 350 for more information.

Related Topics

MTU Ranges for 7000 and 8000 Series Devices and NGIPSv, on page 500
Snort® Restart Scenarios, on page 348

Configuring Passive Interfaces

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **Edit** ( ) next to the device where you want to configure the passive interface.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Edit** ( ) next to the interface you want to configure as a passive interface.

Step 4 Click **Passive**.

Step 5 If you want to associate the passive interface with a security zone, do one of the following:

- Choose an existing security zone from the **Security Zone** drop-down list.
- Choose **New** to add a new security zone; see Creating Security Zone and Interface Group Objects, on page 411.

Step 6 Check the **Enabled** check box.

If you clear the check box, the interface becomes disabled so that users cannot access it for security purposes.

Step 7 7000 & 8000 Series only: From the **Mode** drop-down list, designate the link mode, or choose **Autonegotiation** to specify that the interface is configured to automatically negotiate speed and duplex settings.

Mode settings are available only for copper interfaces.

Interfaces on 8000 Series appliances do not support half-duplex options.

Step 8 7000 & 8000 Series only: From the **MDI/MDIX** drop-down list, designate whether the interface is configured for MDI (medium dependent interface), MDIX (medium dependent interface crossover), or Auto-MDIX.

MDI/MDIX settings are available only for copper interfaces.

By default, MDI/MDIX is set to **Auto-MDIX**, which automatically handles switching between MDI and MDIX to attain link.

Step 9 Enter a maximum transmission unit (MTU) in the **MTU** field.

The range of MTU values can vary depending on the model of the managed device and the interface type.

Caution Changing the highest MTU value among all non-management interfaces on the device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all non-management interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. See Snort® Restart Traffic Behavior, on page 350 for more information.

Step 10 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Inline IPS Deployments

In an inline IPS deployment, you configure the Firepower System transparently on a network segment by binding two ports together. This allows the system to be installed in any network environment without the configuration of adjacent network devices. Inline interfaces receive all traffic unconditionally, but all traffic received on these interfaces is retransmitted out of an inline set unless explicitly dropped.

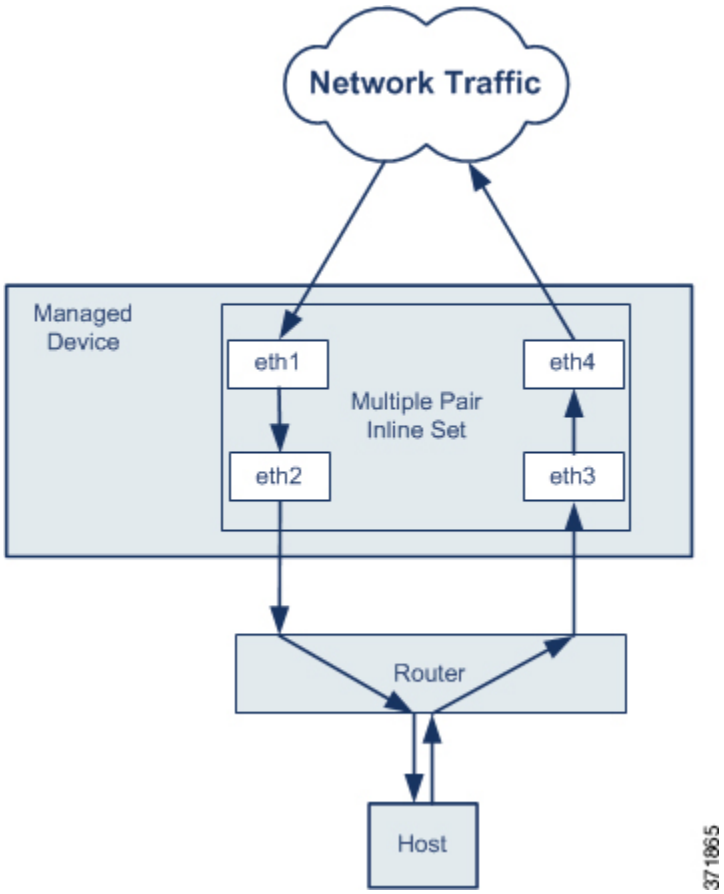


Note For the system to affect traffic, you must deploy relevant configurations to managed devices using routed, switched, or transparent interfaces, or inline interface pairs.

You can configure the interfaces on your managed device to route traffic between a host on your network and external hosts through different inline interface pairs, depending on whether the device traffic is inbound or outbound. This is an *asynchronous routing* configuration. If you deploy asynchronous routing but you include only one interface pair in an inline set, the device might not correctly analyze your network traffic because it might see only half of the traffic.

Adding multiple inline interface pairs to the same inline interface set allows the system to identify the inbound and outbound traffic as part of the same traffic flow. For passive interfaces only, you can also achieve this by including the interface pairs in the same security zone.

When the system generates a connection event from traffic passing through an asynchronous routing configuration, the event may identify an ingress and egress interface from the same inline interface pair. The configuration in the following diagram, for example, would generate a connection event identifying **eth3** as the ingress interface and **eth2** as the egress interface. This is expected behavior in this configuration.



Note If you assign multiple interface pairs to a single inline interface set but you experience issues with duplicate traffic, reconfigure to help the system uniquely identify packets. For example, you could reassign your interface pairs to separate inline sets or modify your security zones.

For devices with inline sets, a software bridge is automatically set up to transport packets after the device restarts. If the device is restarting, there is no software bridge running anywhere. If you enable bypass mode on the inline set, it goes into hardware bypass while the device is restarting. In that case, you may lose a few seconds of packets as the system goes down and comes back up, due to renegotiation of link with the device. However, the system will pass traffic while Snort is restarting.

Related Topics

MTU Ranges for 7000 and 8000 Series Devices and NGIPSv, on page 500
Snort® Restart Scenarios, on page 348

Inline Interfaces on the Firepower System

You can configure one or more physical ports on a managed device as inline interfaces. You must assign a pair of inline interfaces to an inline set before they can handle traffic in an inline deployment.

Note:

- The system warns you if you set the interfaces in an inline pair to different speeds or if the interfaces negotiate to different speeds.
- If you configure an interface as an inline interface, the adjacent port on its NetMod automatically becomes an inline interface as well to complete the pair.
- To configure inline interfaces on an NGIPSv device, you must create the inline pair using adjacent interfaces.

Configuring Inline Interfaces

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **Edit** ( ) next to the device where you want to configure the interface.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Edit** ( ) next to the interface you want to configure.

Step 4 Click **Inline**.

Step 5 If you want to associate the inline interface with a security zone, do one of the following:

- Choose an existing security zone from the **Security Zone** drop-down list.
- Choose **New** to add a new security zone; see Creating Security Zone and Interface Group Objects, on page 411.

Step 6 Choose an existing inline set from the **Inline Set** drop-down list, or choose **New** to add a new inline set.

Note If you add a new inline set, you must configure it after you set up the inline interface; see Adding Inline Sets, on page 511.

Step 7 Check the **Enabled** check box.

If you clear the check box, the interface becomes disabled so that users cannot access it for security purposes.

Step 8 7000 & 8000 Series only: From the **Mode** drop-down list, designate the link mode, or choose **Autonegotiation** to specify that the interface is configured to automatically negotiate speed and duplex settings.

Mode settings are available only for copper interfaces.

Interfaces on 8000 Series appliances do not support half-duplex options.

Step 9 7000 & 8000 Series only: From the **MDI/MDIX** drop-down list, designate whether the interface is configured for MDI (medium dependent interface), MDIX (medium dependent interface crossover), or Auto-MDIX.

MDI/MDIX settings are available only for copper interfaces.

By default, MDI/MDIX is set to **Auto-MDIX**, which automatically handles switching between MDI and MDIX to attain link.

Step 10 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Inline Sets

Before you can use inline interfaces in an inline deployment, you must configure inline sets and assign inline interface pairs to them. An inline set is a grouping of one or more inline interface pairs on a device; an inline interface pair can belong to only one inline set at a time.

The **Inline Sets** tab of the Device Management page displays a list of all inline sets you have configured on a device.

You can add inline sets from the **Inline Sets** tab of the Device Management page or you can add inline sets as you configure inline interfaces.

You can assign **only** inline interface pairs to an inline set. If you want to create an inline set before you configure the inline interfaces on your managed devices, you can create an empty inline set and add interfaces to it later. You can use alphanumeric characters and spaces when you type a name for an inline set.

Note Create inline sets before you add security zones for the interfaces in the inline set; otherwise security zones are removed and you must add them again.

Name

The name of the inline set.

Interfaces

A list of all inline interface pairs assigned to the inline set. A pair is not available when you disable either interface in the pair from the Interfaces tab.

MTU

The maximum transmission unit for the inline set. The range of MTU values can vary depending on the model of the managed device and the interface type.

Caution Changing the highest MTU value among all non-management interfaces on the device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all non-management interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. See Snort® Restart Traffic Behavior, on page 350 for more information.

Failsafe

Behavior of the interface on a 7000 or 8000 Series or NGIPSv device when the Snort process is busy or down.

- Enabled—New and existing flows pass without inspection when the Snort process is busy or down.
- Disabled—New and existing flows drop when the Snort process is busy and pass without inspection when the Snort process is down.

The Snort process can be busy when traffic buffers are full, indicating that there is more traffic than the managed device can handle, or because of other software issues.

The Snort process goes down when you deploy a configuration that requires it to restart. See Configurations that Restart the Snort Process When Deployed or Activated, on page 352 for more information.

Note When traffic passes without inspection, features that rely on the Snort process do not function. These include application control and deep inspection. The system performs only basic access control using simple, easily determined transport and network layer characteristics.

Bypass Mode

Firepower 7000 or 8000 Series only: The configured bypass mode of the inline set. This setting determines how the relays in the inline interfaces respond when an interface fails. The bypass mode allows traffic to continue to pass through the interfaces. The non-bypass mode blocks traffic.

Caution In bypass mode, you may lose a few packets when you reboot the appliance. You cannot configure bypass mode for inline sets on 7000 or 8000 Series devices in a high-availability pair, for non-bypass NetMods on 8000 Series devices, or for SFP modules on Firepower 7115 or 7125 devices.

Related Topics

MTU Ranges for 7000 and 8000 Series Devices and NGIPSv, on page 500
Snort® Restart Scenarios, on page 348

Viewing Inline Sets

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **Edit** ( ) next to the device where you want to view the inline sets.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Inline Sets**.

Adding Inline Sets

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **Edit** ( ) next to the device where you want to add the inline set.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Inline Sets**.

Step 4 Click **Add Inline Set**.

Step 5 Enter a **Name**.

Step 6 Next to **Interfaces**, choose one or more inline interface pairs, then click **Add Selected**. To add all interface pairs to the inline set, click **Add All**.

Tip To remove inline interfaces from the inline set, choose one or more inline interface pairs and click **Remove Selected**. To remove all interface pairs from the inline set, click **Remove All**. Disabling either interface in a pair from **Interfaces** also removes the pair.

Step 7 Enter a maximum transmission unit (MTU) in the **MTU** field.

The range of MTU values can vary depending on the model of the managed device and the interface type.

Caution Changing the highest MTU value among all non-management interfaces on the device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all non-management interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. See Snort® Restart Traffic Behavior, on page 350 for more information.

Step 8 If you want to specify that traffic is allowed to bypass detection and continue through the device when the Snort process is busy or down, choose **Failsafe**. See Inline Sets, on page 509 for more information.

Enabling **Failsafe** on a device with inline sets greatly decreases the risk of dropped packets if the internal traffic buffers are full, but your device may still drop packets in certain conditions. In the worst case, the device may experience a temporary network outage.

Step 9 (7000/8000 series only) Specify the bypass mode.

- Click **Bypass** to allow traffic to continue to pass through the interfaces.
- Click **Non-Bypass** to block traffic.

Note You cannot configure bypass mode for inline sets on 7000 or 8000 Series devices in high-availability pairs, inline sets on an NGIPSv device, for non-bypass NetMods on 8000 Series devices, or for SFP modules on Firepower 7115 or 7125 devices.

Step 10 Optionally, configure advanced settings; see Advanced Inline Set Options, on page 512.

Step 11 Click **OK**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

MTU Ranges for 7000 and 8000 Series Devices and NGIPSv, on page 500
Snort® Restart Scenarios, on page 348

Advanced Inline Set Options

There are a number of advanced options you may consider as you configure inline sets.

Tap Mode

Tap mode is available on 7000 and 8000 Series devices when you create an inline or inline with fail-open interface set.

With tap mode, the device is deployed inline, but instead of the packet flow passing through the device, a copy of each packet is sent to the device and the network traffic flow is undisturbed. Because you are working with copies of packets rather than the packets themselves, rules that you set to drop and rules that use the replace keyword do not affect the packet stream. However, rules of these types do generate intrusion events when they are triggered, and the table view of intrusion events indicates that the triggering packets would have dropped in an inline deployment.

There are benefits to using tap mode with devices that are deployed inline. For example, you can set up the cabling between the device and the network as if the device were inline and analyze the kinds of intrusion events the device generates. Based on the results, you can modify your intrusion policy and add the drop rules that best protect your network without impacting its efficiency. When you are ready to deploy the device inline, you can disable tap mode and begin dropping suspicious traffic without having to reconfigure the cabling between the device and the network.

Note that you cannot enable this option and strict TCP enforcement on the same inline set.

Propagate Link State

Note Link state propagation is not supported on virtual devices. Only 7000 and 8000 Series devices support link state propagation.

Link state propagation is a feature for inline sets configured in bypass mode and non-bypass mode so both pairs of an inline set track state. Link state propagation is available for both copper and fiber configurable bypass interfaces.

Link state propagation automatically brings down the second interface in the inline interface pair when one of the interfaces in an inline set goes down. When the downed interface comes back up, the second interface automatically comes back up, also. In other words, if the link state of one interface changes, the appliance senses the change and updates the link state of the other interface to match it. Note that appliances require up to 4 seconds to propagate link state changes.

Link state propagation is especially useful in resilient network environments where routers are configured to reroute traffic automatically around network devices that are in a failure state.

You cannot disable link state propagation for inline sets configured on 7000 and 8000 Series devices in high-availability pairs.

Transparent Inline Mode

Transparent Inline Mode option allows the device to act as a “bump in the wire” and means that the device forwards all the network traffic it sees, regardless of its source and destination. You cannot disable this option on 7000 and 8000 Series devices.

Strict TCP Enforcement

Note Strict TCP enforcement is not supported on virtual devices. Only 7000 and 8000 Series devices support this option. In addition, you cannot enable this option and tap mode on the same inline set.

To maximize TCP security, you can enable strict enforcement, which blocks connections where the three-way handshake was not completed. Strict enforcement also blocks:

- non-SYN TCP packets for connections where the three-way handshake was not completed
- non-SYN/RST packets from the initiator on a TCP connection before the responder sends the SYN-ACK
- non-SYN-ACK/RST packets from the responder on a TCP connection after the SYN but before the session is established
- SYN packets on an established TCP connection from either the initiator or the responder

Configuring Advanced Inline Set Options

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **Edit** ( ) next to the device where you want to edit the inline set.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Inline Sets**.

Step 4 Click **Edit** ( ) next to the inline set you want to edit.

Step 5 Click **Advanced**.

Step 6 Configure options as described in Advanced Inline Set Options, on page 512.

Note Link state propagation and strict TCP enforcement are not supported on virtual devices.

Step 7 Click **OK**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Deleting Inline Sets

When you delete an inline set, any inline interfaces assigned to the set become available for inclusion in another set. The interfaces are not deleted.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to delete the inline set, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Inline Sets**.

Step 4 Next to the inline set you want to delete, click **Delete** ( ).

Step 5 When prompted, confirm that you want to delete the inline set.

What to do next

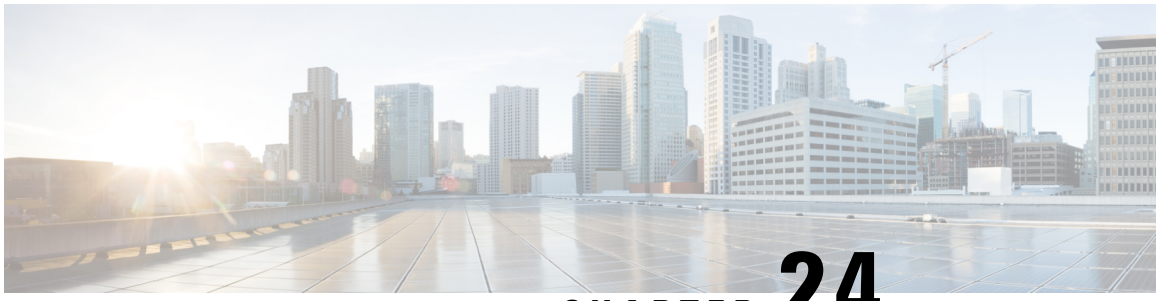
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.



PART VI

Classic Device High Availability and Scalability

- 7000 and 8000 Series Device High Availability, on page 517
- 8000 Series Device Stacking, on page 533



CHAPTER 24

7000 and 8000 Series Device High Availability

The following topics describe how to configure high availability for Firepower 7000 Series and 8000 Series devices in the Firepower System:

- About 7000 and 8000 Series Device High Availability, on page 517
- Establishing Firepower 7000/8000 Series High Availability, on page 521
- Editing Device High Availability, on page 522
- Configuring Individual Devices in a High-Availability Pair, on page 523
- Configuring Individual Device Stacks in a High-Availability Pair, on page 523
- Configuring Interfaces on a Device in a High-Availability Pair, on page 524
- Switching the Active Peer in a Device High-Availability Pair, on page 525
- Placing a High-Availability Peer into Maintenance Mode, on page 525
- Replacing a Device in a Stack in a High-Availability Pair, on page 526
- Device High Availability State Sharing, on page 526
- Device High Availability State Sharing Statistics for Troubleshooting, on page 529
- Separating Device High-Availability Pairs, on page 532

About 7000 and 8000 Series Device High Availability

With 7000 and 8000 Series device high availability, you can establish redundancy of networking functionality and configuration data between two peer devices or two peer device stacks.

You achieve configuration redundancy by configuring two peer devices or two peer device stacks into a high-availability pair to act as a single logical system for policy deploys, system updates, and registration. The system automatically synchronizes other configuration data.

Note Static routes, non-SFRP IP addresses, and routing priorities are not synchronized between the peer devices or peer device stacks. Each peer device or peer device stack maintains its own routing intelligence.

Related Topics

SFRP

Advanced Virtual Switch Settings, on page 1111

Device High Availability Requirements

Before you can configure a 7000 and 8000 Series device high-availability pair, the following must be true:

- You can only pair single devices with single devices or device stacks with device stacks.
- Both devices or device stacks must have normal health status, be running the same software, and have the same licenses. See Using the Health Monitor, on page 291 for more information. In particular, the devices cannot have hardware failures that would cause them to enter maintenance mode and trigger a failover.

Note After you pair the devices, you cannot change the license options for individual paired devices, but you can change the license for the entire high-availability pair.

- Interfaces must be configured on each device or each primary device in a stack.
- Both devices or the primary members of the device stacks must be the same model and have identical copper or fiber interfaces.
- Device stacks must have identical hardware configurations, except for an installed malware storage pack. For example, you can pair a Firepower 8290 with another 8290. None, one, or all devices in either stack might have a malware storage pack.

Caution Do not attempt to install a hard drive that was not supplied by Cisco in your device. Installing an unsupported hard drive may damage the device. Malware storage pack kits are available for purchase **only** from Cisco, and are for use **only** with 8000 Series devices. Contact Support if you require assistance with the malware storage pack. See the *Firepower System Malware Storage Pack Guide* for more information.

- If the devices are targeted by NAT policies, both peers must have the same NAT policy.
- In a multidomain deployment, you can only establish 7000 or 8000 Series device high-availability or device stacks within a leaf domain.

Note After failover and recovery, SFRP preempts to the primary node.

Device High Availability Failover and Maintenance Mode

With a 7000 and 8000 Series device high availability, the system fails over either manually or automatically. You manually trigger failover by placing one of the paired devices or stacks in maintenance mode.

Automatic failover occurs after the health of the active device or stack becomes compromised, during a system update, or after a user with Administrator privileges shuts down the device. Automatic failover also occurs after an active device or device stack experiences NMSB failure, NFE failure, hardware failure, firmware

failure, critical process failure, a disk full condition, or link failure between two stacked devices. If the health of the standby device or stack becomes similarly compromised, the system does not fail over and enters a degraded state. The system also does not fail over when one of the devices or device stacks is in maintenance mode. Note that disconnecting the stacking cable from an active stack sends that stack into maintenance mode. Shutting down the secondary device in an active stack also sends that stack into maintenance mode.



Note If the active member of the high-availability pair goes into maintenance mode and the active role fails over to the other pair member, when the original active pair member is restored to normal operation it does not automatically reclaim the active role.

Configuration Deployment and Upgrade Behavior for High-Availability Pairs

This topic describes upgrade and deployment behavior for 7000 and 8000 Series devices (and stacks) in high availability pairs.

Behavior During Deploy

You deploy configuration changes to the members of a high availability pair at the same time. Deploy either succeeds or fails for both peers. The Firepower Management Center deploys to the active device; if that succeeds then changes are deployed to the standby.



Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 and Configurations that Restart the Snort Process When Deployed or Activated, on page 352.

Behavior During Upgrade

You should not experience interruptions in traffic flow or inspection while upgrading devices (or device stacks) in high availability pairs. To ensure continuity of operations, they upgrade one at a time. Devices operate in maintenance mode while they upgrade.

Which peer upgrades first depends on your deployment:

- Routed or switched—Standby upgrades first. The devices switch roles, then the new standby upgrades. When the upgrade completes, the devices' roles remain switched. If you want to preserve the active/standby roles, manually switch the roles before you upgrade. That way, the upgrade process switches them back.
- Access control only— Active upgrades first. When the upgrade completes, the active and standby maintain their old roles.

Deployment Types and Device High Availability

You determine how to configure 7000 or 8000 Series device high availability depending on your Firepower System deployment: passive, inline, routed, or switched. You can also deploy your system in multiple roles at once. Of the four deployment types, only passive deployments require that you configure devices or stacks

using high availability to provide redundancy. You can establish network redundancy for the other deployment types with or without device high availability. For a brief overview on high availability in each deployment type, see the sections below.



Note You can achieve Layer 3 redundancy without using device high availability by using the Cisco Redundancy Protocol (SFRP). SFRP allows devices to act as redundant gateways for specified IP addresses. With network redundancy, you configure two devices or stacks to provide identical network connections, ensuring connectivity for other hosts on the network.

Passive Deployment Redundancy

Passive interfaces are generally connected to tap ports on central switches, which allows them to analyze all of the traffic flowing across the switch. If multiple devices are connected to the same tap feed, the system generates events from each of the devices. When configured in a high-availability pair, devices act as either active or standby, which allows the system to analyze traffic even in the event of a system failure while also preventing duplicate events.

Inline Deployment Redundancy

Because an inline set has no control over the routing of the packets being passed through it, it must always be active in a deployment. Therefore, redundancy relies on external systems to route traffic correctly. You can configure redundant inline sets with or without 7000 or 8000 Series device high availability.

To deploy redundant inline sets, you configure the network topology so that it allows traffic to pass through only one of the inline sets while preventing circular routing. If one of the inline sets fails, the surrounding network infrastructure detects the loss of connectivity to the gateway address and adjusts the routes to send traffic through the redundant set.

Routed Deployment Redundancy

Hosts in an IP network must use a well-known gateway address to send traffic to different networks. Establishing redundancy in a routed deployment requires that routed interfaces share the gateway addresses so that only one interface handles traffic for that address at any given time. To accomplish this, you must maintain an equal number of IP addresses on a virtual router. One interface advertises the address. If that interface goes down, the standby interface begins advertising the address.

In devices that are not members of a high-availability pair, you use SFRP to establish redundancy by configuring gateway IP addresses shared between multiple routed interfaces. You can configure SFRP with or without 7000 or 8000 Series device high availability. You can also establish redundancy using dynamic routing such as OSPF or RIP.

Switched Deployment Redundancy

You establish redundancy in a switched deployment using the Spanning Tree Protocol (STP), one of the advanced virtual switch settings. STP is a protocol that manages the topology of bridged networks. It is specifically designed to allow redundant links to provide automatic standby for switched interfaces without configuring standby links. Devices in a switched deployment rely on STP to manage traffic between redundant interfaces. Two devices connected to the same broadcast network receive traffic based on the topology calculated by STP.

Note Cisco strongly recommends that you enable STP when configuring a virtual switch that you plan to deploy in a 7000 or 8000 Series device high-availability pair.

7000/8000 Series High Availability Configuration

When establishing 7000 or 8000 Series device high availability, you designate one of the devices or stacks as active and the other as standby. The system applies a merged configuration to the paired devices. If there is a conflict, the system applies the configuration from the device or stack you designated as active.

After you pair the devices, you cannot change the license options for individual paired devices, but you can change the license for the entire high availability pair. If there are interface attributes that need to be set on switched interfaces or routed interfaces, the system establishes the high availability pair, but sets it to a pending status. After you configure the necessary attributes, the system completes the high availability pair and sets it to a normal status.

After you establish a high availability pair, the system treats the peer devices or stacks as a single device on the Device Management page. Device high availability pairs display the High Availability icon ( ) in the appliance list. Any configuration changes you make are synchronized between the paired devices. The Device Management page displays which device or stack in the high availability pair is active, which changes after manual or automatic failover.

Removing registration of a device high availability pair from a Firepower Management Center removes registration from both devices or stacks. You remove a device high availability pair from the FMC as you would an individual managed device.

You can then register the high availability pair on another FMC. To register single devices from a high availability pair, you add remote management to the active device in the pair and then add that device to the FMC, which adds the whole pair. To register stacked devices in a high availability pair, you add remote management to the primary device of the either stack and then add that device to the FMC, which adds the whole pair.

After you establish a device high availability pair, you should configure a high-availability link interface.

Note If you plan to set up dynamic NAT, HA state sharing, or VPN using the devices in the high availability pair, you must configure a high-availability link interface. For more information, see Configuring HA Link Interfaces, on page 498.

Establishing Firepower 7000/8000 Series High Availability

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Any	Admin/Network Admin

Note This procedure describes establishing a 7000 & 8000 Series device high-availability pair. For information on establishing Firepower Threat Defense high availability, see Add a Firepower Threat Defense High Availability Pair, on page 699.

When establishing a 7000 & 8000 Series device high-availability pair, you designate one of the devices or stacks as active and the other as standby. The system applies a merged configuration to the paired devices. If there is a conflict, the system applies the configuration from the device or stack you designated as active.

In a multidomain deployment, devices in a high-availability pair must belong to the same domain.

Before you begin

Confirm that all requirements are met; see Device High Availability Requirements, on page 518.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 From the **Add** drop-down menu, choose **Add High Availability**.

Step 3 Enter a **Name**.

Step 4 Under **Device Type**, choose **Firepower**.

Step 5 Assign roles for the devices or stacks:

a) Choose the **Active Peer** device or stack for the high-availability pair.

b) Choose the **Standby Peer** device or stack for the high-availability pair.

Step 6 Click **Add**. The process takes a few minutes as the system synchronizes data.

What to do next

Create an HA Link interface on each of the devices in the high-availability pair if you plan to set up HA state sharing, dynamic NAT, or VPN with the devices. For more information on HA link interfaces, see Configuring HA Link Interfaces, on page 498.

Editing Device High Availability

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

After you establish a 7000 or 8000 Series device high-availability pair, most changes you make to the device configuration also change the configuration of the whole high-availability pair.

You can view the status of the high-availability pair by hovering your pointer over the status icon in the General section. You can also view which device or stack is the active peer and standby peer in the pair.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device high availability pair where you want to edit the configuration, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Use the sections on the High Availability page to make changes to the high-availability pair configuration as you would a single device configuration.

Configuring Individual Devices in a High-Availability Pair

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

After you establish a 7000 or 8000 Series device high-availability pair, you can still configure some attributes for each device within the pair. You can make changes to a paired device just as you would to a single device.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device high-availability pair where you want to edit the configuration, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Devices** tab.

Step 4 From the **Selected Device** drop-down list, choose the device you want to modify.

Step 5 Use the sections on the Devices page to make changes to the individual paired device as you would a single device.

Configuring Individual Device Stacks in a High-Availability Pair

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	Firepower 8140, Firepower 8200 family, Firepower 8300 family	Leaf only	Admin/Network Admin

After you configure stacked 8000 Series devices into a high-availability pair, the system limits the stack attributes that you can edit. You can edit the name of a stack in a paired stack. In addition, you can edit the

network configuration of the stack, as described in Configuring Interfaces on a Device in a High-Availability Pair, on page 524.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device high-availability pair where you want to edit the configuration, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Stacks** tab.

Step 4 From the **Selected Device** drop-down list, choose the stack you want to modify.

Step 5 Next to the **General** section, click the edit icon ( ).

Step 6 Enter a **Name**.

Step 7 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configuring Interfaces on a Device in a High-Availability Pair

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

You can configure interfaces on individual devices in a 7000 or 8000 Series device high-availability pair. However, you must also configure an equivalent interface on the peer device in the pair. For paired stacks, you configure identical interfaces on the primary devices of the stacks. When you configure virtual routers, you select the stack where you want to configure the routers.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device high-availability pair where you want to configure interfaces, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Interfaces** tab.

Step 4 From the **Selected Device** drop-down list, choose the device you want to modify.

Step 5 Configure interfaces as you would on an individual device.

Related Topics

Virtual Router Configuration, on page 1123

Switching the Active Peer in a Device High-Availability Pair

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Any	Admin/Network Admin

After you establish a 7000 or 8000 Series device high-availability pair, you can manually switch the active and standby peer devices or stacks.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device high-availability pair where you want to change the active peer, click the Switch Active Peer icon ( ).

Step 3 You can:

- Click **Yes** to immediately make the standby peer the active peer in the high-availability pair.
- Click **No** to cancel and return to the Device Management page.

Placing a High-Availability Peer into Maintenance Mode

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Any	Admin/Network Admin

After you establish a 7000 or 8000 Series device high-availability pair, you can manually trigger failover by placing one of the peers into maintenance mode to perform maintenance on the devices. In maintenance mode, the system administratively takes down all interfaces except for the management interface. After maintenance is completed, you can re-enable the peer to resume normal operation.

Note You should not place both peers in a high-availability pair into maintenance mode at the same time. Doing so will prevent that pair from inspecting traffic.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the peer you want to place in maintenance mode, click the toggle maintenance mode icon ().

Step 3 Click **Yes** to confirm maintenance mode.

What to do next

- When maintenance is complete, click the toggle maintenance mode icon () again to bring the peer out of maintenance mode.

Replacing a Device in a Stack in a High-Availability Pair

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	Firepower 8140, 8200 family, 8300 family	Any	Admin/Network Admin

After you place a stack that is a member of a high-availability pair into maintenance mode, you can replace a secondary device in the stack for another device. You can only select devices that are not currently stacked or paired. The new device must follow the same guidelines for establishing a device stack.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the stack member you want to place into maintenance mode, click the toggle maintenance mode icon ().

Step 3 Click **Yes** to confirm maintenance mode.

Step 4 Click the replace device icon ().

Step 5 Choose the **Replacement Device** from the drop-down list.

Step 6 Click **Replace** to replace the device.

Step 7 Click the toggle maintenance mode icon () again to bring the stack immediately out of maintenance mode.

Note You do not need to re-deploy the device configuration.

Device High Availability State Sharing

Device high availability state sharing allows devices or stacks in high-availability pairs to synchronize as much state as necessary, so that if either device or stack fails, the other peer can take over with no interruption to traffic flow. Without state sharing, the following features may not fail over properly:

- Strict TCP enforcement

- Unidirectional access control rules
- Blocking persistence

Note, however, that enabling state sharing slows system performance.

You must configure and enable HA link interfaces on both devices or the primary stacked devices in the high-availability pair before you can configure high availability state sharing. Firepower 82xx Family and 83xx Family devices require a 10G HA link, while other model devices require a 1G HA link.

You must disable state sharing before you can modify the HA link interfaces.



Note If paired devices fail over, the system terminates all existing SSL-encrypted sessions on the active device. Even if you establish high availability state sharing, these sessions must be renegotiated on the standby device. If the server establishing the SSL session supports session reuse and the standby device does not have the SSL session ID, it cannot renegotiate the session.

Strict TCP Enforcement

When you enable strict TCP enforcement for a domain, the system drops any packets that are out of order on TCP sessions. For example, the system drops non-SYN packets received on an unestablished connection. With state sharing, devices in the high-availability pair allow TCP sessions to continue after failover without having to reestablish the connection, even if strict TCP enforcement is enabled. You can enable strict TCP enforcement on inline sets, virtual routers, and virtual switches.

Unidirectional Access Control Rules

If you have configured unidirectional access control rules, network traffic may match a different access control rule than intended when the system reevaluates a connection midstream after failover. For example, consider if you have a policy containing the following two access control rules:

```
Rule 1: Allow from 192.168.1.0/24 to 192.168.2.0/24
Rule 2: Block all
```

Without state sharing, if an allowed connection from 192.168.1.1 to 192.168.2.1 is still active following a failover and the next packet is seen as a response packet, the system denies the connection. With state sharing, a midstream pickup would match the existing connection and continue to be allowed.

Blocking Persistence

While many connections are blocked on the first packet based on access control rules or other factors, there are cases where the system allows some number of packets through before determining that the connection should be blocked. With state sharing, the system immediately blocks the connection on the peer device or stack as well.

When establishing state sharing for a high-availability pair, you can configure the following options:

Enabled

Click the check box to enable state sharing. Clear the check box to disable state sharing.

Minimum Flow Lifetime

Specify the minimum time (in milliseconds) for a session before the system sends any synchronization messages for it. You can use any integer from 0 to 65535. The system does not synchronize any sessions that have not met the minimum flow lifetime, and the system synchronizes only when a packet is received for the connection.

Minimum Sync. Interval

Specify the minimum time (in milliseconds) between update messages for a session. You can use any integer from 0 to 65535. The minimum synchronization interval prevents synchronization messages for a given connection from being sent more frequently than the configured value after the connection reaches the minimum lifetime.

Maximum HTTP URL Length

Specify the maximum characters for the URL the system synchronizes between the paired devices. You may use any integer from 0 to 225.

Related Topics

Configuring HA Link Interfaces, on page 498

Establishing Device High-Availability State Sharing

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

Device high-availability state sharing allows 7000 or 8000 Series devices or stacks in high-availability pairs to synchronize as much state as necessary, so that if either device or stack fails, the other peer can take over with no interruption to traffic flow.

Caution Modifying a high-availability state sharing option on a 7000 or 8000 Series device restarts the Snort process on the primary and secondary devices, temporarily interrupting traffic inspection on both devices. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Procedure

Step 1 Configure HA link interfaces for each device in the device high-availability pair; see Configuring HA Link Interfaces, on page 498.

Step 2 Choose **Devices** > **Device Management**.

Step 3 Next to the device high-availability pair you want to edit, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 4 In the **State Sharing** section, click the edit icon ( ).

Step 5 Decrease the state sharing values to improve paired peer readiness, or increase the values to allow better performance.

We recommend you use the default values, unless your deployment presents a good reason to change them.

Step 6 Click **OK**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Configuring HA Link Interfaces, on page 498
Snort® Restart Scenarios, on page 348

Device High Availability State Sharing Statistics for Troubleshooting

The sections below describe the statistics you can view for each device and how you can use them to troubleshoot your state sharing configuration for 7000 and 8000 Series device high-availability pairs.

Messages Received (Unicast)

Messages received are the number of high availability synchronization messages received from the paired peer.

The value should be close to the number of messages sent by the peer. During active use, the values may not match, but should be close. If traffic stops, the values should become stable and the messages received will match the messages sent.

For troubleshooting, you should view both the messages received and the messages sent, compare the rate of increase, and make sure the values are close. The sent value on each peer should be incrementing at approximately the same rate as the received value on the opposite peer.

Contact Support if the received messages stop incrementing or increment slower than the messages sent by the peer.

Packets Received

The system batches multiple messages into single packets in order to decrease overhead. The Packets Received counter displays the total number of these data packets, as well as other control packets that have been received by a device.

The value should be close to the number of packets sent by the peer device. During active use, the values may not match, but should be close. Because the number of messages received should be close and incrementing at the same rate as the number of messages sent by the peer, the number of packets received should have the same behavior.

For troubleshooting, you should view both the packets received and the messages sent, compare the rate of increase, and make sure the values are increasing at the same rate. If the sent value on the paired peer is incrementing, the received value on the device should also increase at the same rate.

Contact Support if the received packets stop incrementing or increment slower than the messages sent by the peer.

Total Bytes Received

Total bytes received are the number of bytes that make up the packets received by the peer.

The value should be close to the number of bytes sent by the other peer. During active use, the values may not match, but should be close.

For troubleshooting, you should view both the total bytes received and the messages sent, compare the rate of increase, and make sure the values are increasing at the same rate. If the sent value on the paired peer is incrementing, the received value on the device should also increase at the same rate.

Contact Support if the received bytes stop incrementing or increment slower than the messages sent by the peer.

Protocol Bytes Received

Protocol bytes received are the number of bytes of protocol overhead received, which includes everything but the payload of session state synchronization messages.

The value should be close to the number of bytes sent by the peer. During active use, the values may not match, but should be close.

For troubleshooting, you should view the total bytes received to discover how much actual state data is being shared in comparison to protocol data. If the protocol data is a large percentage of the data being sent, you can adjust the minimum sync interval.

Contact Support if the protocol bytes received increment at a similar rate to the total bytes received. Protocol bytes received should be minimal in relation to the total bytes received.

Messages Sent

Messages sent are the number of high availability synchronization messages sent to the paired peer.

This data is useful in comparison to the number of messages received. During active use, the values may not match, but should be close.

For troubleshooting, you should view both the messages received and the messages sent, compare the rate of increase, and make sure the values are close.

Contact Support if the messages sent increment at a similar rate to the total bytes received.

Bytes Sent

Bytes sent are the total number of bytes sent that make up the high availability synchronization messages sent to the peer.

This data are useful in comparison to the number of messages received. During active use, the values may not match, but should be close. The number of bytes received on the peer should be close to, but not more than this value.

Contact Support if the total bytes received is not incrementing at about the same rate as the bytes sent.

Tx Errors

Tx errors are the number of memory allocation failures the system encounters when trying to allocate space for messages to be sent to the paired peer.

This value should be zero at all times on both peers. Contact Support if this number is not zero or if the number steadily increases, which indicates the system has encountered an error where it cannot allocate memory.

Tx Overruns

Tx overruns are the number of times the system attempts and fails to place a message into the transit queue.

This value should be zero at all times on both peers. When the value is not zero or is steadily increasing, it indicates that the system is sharing too much data across the HA link that cannot be sent quickly enough.

You should increase the HA link MTU if it was previously set below the default value (9918 or 9922). You can change the minimum flow lifetime and minimum synchronization interval settings to reduce the amount of data shared across the HA link to prevent the number from incrementing.

Contact Support if this value persists or continues to increase.

Recent Logs

The system log displays the most recent high availability synchronization messages. The log should not display any ERROR or WARN messages. It should remain comparable between the peers, such as the same number of sockets being connected.

However, the data displayed may be opposite in some instances, for example, one peer reports that it received a connection from the other peer and references different IP addresses. The log provides a comprehensive view of the high availability state sharing connection, and any errors within the connection.

Contact Support if the log displays an ERROR or WARN message, or any message that does not appear to be purely informational.

Viewing Device High Availability State Sharing Statistics

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

After you establish state sharing, you can view the following information about the configuration in the **State Sharing** section of the High Availability page:

- The HA link interface that is being used and its current link state
- Detailed synchronization statistics for troubleshooting issues

The state sharing statistics are primarily counters for different aspects of the high availability synchronization traffic sent and received, along with some other error counters. In addition, you can view the latest system logs for each device in the high-availability pair.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device high-availability pair you want to edit, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 In the **State Sharing** section, click the view statistics icon ( ).

Step 4 Choose a **Device** to view if your high-availability pair is composed of device stacks.

Step 5 You can:

- Click **Refresh** to update the statistics.
- Click **View** to view the latest data log for each device in the high-availability pair.

Separating Device High-Availability Pairs

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Any	Admin/Network Admin

When you separate, or "break," a 7000 or 8000 Series device high-availability pair:

- The active peer (device or stack) retains full deployment functionality
- The standby peer (device or stack) loses its interface configurations and fails over to the active peer, unless you choose to leave the interface configurations active, in which case the standby peer resumes normal operation.
- The standby peer always loses the configuration of passive interfaces.
- Any peer in maintenance mode resumes normal operation.

Procedure

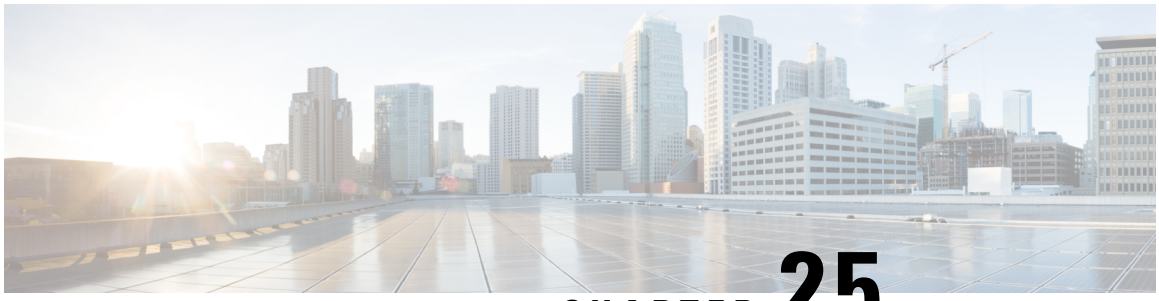
Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the high-availability pair you want to break, click the Break HA icon ( ).

Step 3 Optionally, check the check box to remove the interface configurations on the standby peer.

This step administratively takes down all interfaces except for the management interface.

Step 4 Click **Yes**.



CHAPTER 25

8000 Series Device Stacking

The following topics describe how to work with Firepower 8000 Series device stacks in the Firepower System:

- About Device Stacks, on page 533
- Device Stack Configuration, on page 535
- Establishing Device Stacks, on page 536
- Editing Device Stacks, on page 537
- Replacing a Device in a Stack, on page 537
- Replacing a Device in a Stack in a High-Availability Pair, on page 538
- Configuring Individual Devices in a Stack, on page 539
- Configuring Interfaces on a Stacked Device, on page 539
- Separating Stacked Devices, on page 540
- Replacing a Device in a Stack, on page 541

About Device Stacks

You can increase the amount of traffic inspected on a network segment by using devices in a stacked configuration. For each stacked configuration, all devices in the stack must have the same hardware. However, none, some, or all devices might have an installed malware storage pack. The devices must also be from the same device family based on the following stacked configurations:

The stacked configuration is supported for Firepower 8140, Firepower 8200 family, Firepower 8300 family devices.

For the 81xx Family:

- two Firepower 8140s

For the 82xx Family:

- up to four Firepower 8250s
- a Firepower 8260 (a primary device and a secondary device)
- a Firepower 8270 (a primary device with 40G capacity and two secondary devices)
- a Firepower 8290 (a primary device with 40G capacity and three secondary devices)

For the 83xx Family:

- up to four Firepower 8350s
- up to four AMP8350s
- a Firepower 8360 (a primary device with 40G capacity and a secondary device)
- an AMP8360 (a primary device with 40G capacity and a secondary device)
- a Firepower 8370 (a primary device with 40G capacity and two secondary devices)
- an AMP8370 (a primary device with 40G capacity and two secondary devices)
- a Firepower 8390 (a primary device with 40G capacity and three secondary devices)
- an AMP8390 (a primary device with 40G capacity and three secondary devices)

For more information about stacked configurations, see the Cisco Firepower 8000 Series Getting Started Guide. For more information about the malware storage pack, see the *Firepower System Malware Storage Pack Guide*. *Firepower System Malware Storage Pack Guide*.



Caution Do not attempt to install a hard drive that was not supplied by Cisco in your device. Installing an unsupported hard drive may damage the device. Malware storage pack kits are available for purchase **only** from Cisco, and are for use **only** with 8000 Series devices. Contact Support if you require assistance with the malware storage pack. See the *Firepower System Malware Storage Pack Guide* for more information.

When you establish a stacked configuration, you combine the resources of each stacked device into a single, shared configuration.

You designate one device as the *primary* device, where you configure the interfaces for the entire stack. You designate the other devices as *secondary*. Secondary devices must not be currently sensing any traffic and must not have link on any interface.

Connect the primary device to the network segment you want to analyze in the same way you would configure a single device. Connect the secondary devices to the primary device using the stacked device cabling instructions found in the Cisco Firepower 8000 Series Getting Started Guide.

All devices in the stacked configuration must have the same hardware, run the same software version, and have the same licenses. If the devices are targeted by NAT policies, both the primary and secondary device must have the same NAT policy. You must deploy updates to the entire stack from the Firepower Management Center. If an update fails on one or more devices in the stack, the stack enters a mixed-version state. You cannot deploy policies to or update a stack in a mixed-version state. To correct this state, you can break the stack or remove individual devices with different versions, update the individual devices, then reestablish the stacked configuration. After you stack the devices, you can change the licenses only for the entire stack at once.

After you establish the stacked configuration, the devices act like a single, shared configuration. If the primary device fails, no traffic is passed to the secondary devices. Health alerts are generated indicating that the stacking heartbeat has failed on the secondary devices.

If the secondary device in a stack fails, inline sets with configurable bypass enabled go into bypass mode on the primary device. For all other configurations, the system continues to load balance traffic to the failed secondary device. In either case, a health alert is generated to indicate loss of link.

You can use a device stack as you would a single device in your deployment, with a few exceptions. If you have 7000 or 8000 Series devices in a high-availability pair, you cannot stack a device high-availability pair or a device in a high-availability pair. You also cannot configure NAT on a device stack.

Note If you use eStreamer to stream event data from stacked devices to an external client application, collect the data from each device and ensure that you configure each device identically. The eStreamer settings are not automatically synchronized between stacked devices.

In a multidomain deployment, you can only stack devices that belong to the same domain.

Related Topics

About Health Monitoring, on page 275

Device Stack Configuration

You can increase the amount of traffic inspected on a network segment by stacking two Firepower 8140 devices, up to four Firepower 8250s, a Firepower 8260, a Firepower 8270, a Firepower 8290, up to four Firepower 8350s, a Firepower 8360, a Firepower 8370, or a Firepower 8390 and using their combined resources in a single, shared, configuration. If you have 7000 or 8000 Series devices in a high-availability pair, you cannot stack a device high-availability pair or a device in a high-availability pair. However, you can configure two device stacks into a high-availability pair.

After you establish a device stack, the system treats the devices as a single device on the Device Management page. Device stacks display the stack icon () in the appliance list.

Removing registration of a device stack from a Firepower Management Center also removes registration from both devices. You delete stacked devices from the Firepower Management Center as you would a single managed device; you can then register the stack on another Firepower Management Center. You only need to register one of the stacked devices on the new Firepower Management Center for the entire stack to appear.

After you establish the device stack, you cannot change which devices are primary or secondary unless you break and reestablish the stack. However, you can:

- add secondary devices to an existing stack of two or three Firepower 8250s, a Firepower 8260, or a Firepower 8270 up to the limit of four Firepower 8250s in a stack
- add secondary devices to an existing stack of two or three Firepower 8350s, a Firepower 8360, or a Firepower 8370 up to the limit of four Firepower 8350s in a stack

For additional devices, the primary device in the stack must have the necessary stacking NetMods for additional cabled devices. For example, if you have a Firepower 8260 where the primary only has a single stacking NetMod, you cannot add another secondary device to this stack. You add secondary devices to an existing stack in the same manner that you initially establish a stacked device configuration.

Establishing Device Stacks

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Any	Firepower 8140, 8200 family, 8300 family	Any	Admin/Network Admin

All devices in a stack must be of the same hardware model (for example, a Firepower 8140 with another 8140). You can stack a total of four devices (one primary device and up to three secondary devices) in the 8200 family and in the 8300 family.

In a multidomain deployment, all devices in the stack must belong to the same domain.

Before you begin

- Decide which unit will be the primary device.
- Confirm that the units are cabled properly before designating the primary/secondary relationship. For information about cabling, see the Cisco Firepower 8000 Series Getting Started Guide.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 From the **Add** drop-down menu, choose **Add Stack**.

Step 3 From the **Primary** drop-down list, choose the device that you cabled for primary operation.

Note If you choose a device that is not cabled as the primary device, you cannot perform the next series of steps.

Step 4 Enter a **Name**.

Step 5 Click **Add** to choose the devices you want to include in the stack.

Step 6 From the **Slot on Primary Device** drop-down list, choose the stacking network module that connects the primary device to the secondary device.

Step 7 From the **Secondary Device** drop-down list, choose the device you cabled for secondary operation.

Step 8 From the **Slot on Secondary Device** drop-down list, choose the stacking network module that connects the secondary device to the primary device.

Step 9 Click **Add**.

Step 10 Repeat steps 5 through 9 if you are adding secondary devices to an existing stack of Firepower 8250s, a Firepower 8260, a Firepower 8270, an existing stack of Firepower 8350s, a Firepower 8360, or a Firepower 8370.

Step 11 Click **Stack** to establish the device stack or to add secondary devices. Note that this process takes a few minutes as the process synchronizes system data.

Related Topics

About 7000 and 8000 Series Device High Availability, on page 517

Delete a Device from the FMC, on page 231
Add a Device to the FMC, on page 229

Editing Device Stacks

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Any	Firepower 8140, Firepower 8200 family, Firepower 8300 family	Leaf only	Admin/Network Admin

After you establish a device stack, most changes you make to the device configuration also change the configuration of the entire stack. On the Stack page of the appliance editor, you can make changes to the stack configuration as on the Device page of a single device.

You can change the display name of the stack, enable and disable licenses, view system and health policies, and configure advanced settings.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the stacked device where you want to edit the configuration, click the edit icon ().

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Use the sections on the Stack page to make changes to the stacked configuration as you would a single device configuration.

Replacing a Device in a Stack

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Any	FirePOWER 8140, 8200 family, 8300 family	Any	Admin/Network Admin

If the Firepower Management Center cannot communciate with the device, you must connect to the device and use CLI commands to separate the stack and unregister the device. For more information, see **stacking disable** and **delete** CLI commands in the relevant chapter: Classic Device CLI Configuration Commands, on page 2464.

To replace a device within a stack:

Procedure

Step 1 Select the stack with the device to replace and break that stack. For more information, see Separating Stacked Devices, on page 540.

Step 2 Unregister the device from the Firepower Management Center. For more information, see Delete a Device from the FMC, on page 231.

Step 3 Register the replacement device to the Firepower Management Center. For more information, see Add a Device to the FMC, on page 229.

Step 4 Create a device stack that includes the replacement deivce. for more information, see Establishing Device Stacks, on page 536.

Replacing a Device in a Stack in a High-Availability Pair

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	Firepower 8140, 8200 family, 8300 family	Any	Admin/Network Admin

After you place a stack that is a member of a high-availability pair into maintenance mode, you can replace a secondary device in the stack for another device. You can only select devices that are not currently stacked or paired. The new device must follow the same guidelines for establishing a device stack.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the stack member you want to place into maintenance mode, click the toggle maintenance mode icon ( ).

Step 3 Click **Yes** to confirm maintenance mode.

Step 4 Click the replace device icon ( ).

Step 5 Choose the **Replacement Device** from the drop-down list.

Step 6 Click **Replace** to replace the device.

Step 7 Click the toggle maintenance mode icon ( ) again to bring the stack immediately out of maintenance mode.

Note You do not need to re-deploy the device configuration.

Configuring Individual Devices in a Stack

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Any	Firepower 8140, Firepower 8200 family, Firepower 8300 family	Leaf only	Admin/Network Admin

After you establish a device stack, you can still configure some attributes for an individual device within the stack. You can make changes to a device configured in a stack as you would for a single device. You can change the display name of a device, view system settings, shut down or restart a device, view health information, and edit device management settings.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the stacked device where you want to edit the configuration, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Device** tab.

Step 4 From the **Selected Device** drop-down list, choose the device you want to modify.

Step 5 Use the sections on the Devices page to make changes to the individual stacked device as you would a single device.

Configuring Interfaces on a Stacked Device

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Any	Firepower 8140, Firepower 8200 family, Firepower 8300 family	Leaf only	Admin/Network Admin

With the exception of the management interface, you configure stacked device interfaces on the Interfaces page of the primary device in the stack. You can choose any device in the stack to configure the management interface.

The Interfaces page of a Firepower stacked device includes the hardware and interfaces views that you find on an individual device.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the primary stacked device, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Interfaces** tab.

Step 4 From the **Selected Device** drop-down list, choose the device you want to modify.

Step 5 Configure interfaces as you would on an individual device; see Configuring Sensing Interfaces, on page 497.

Related Topics

Management Interfaces, on page 881

Separating Stacked Devices

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Any	FirePOWER 8140, 8200 family, 8300 family	Any	Admin/Network Admin

If you no longer need to use a stacked configuration for your devices, you can break the stack and separate the devices.

Note If a stacked device fails, or if communication fails between member devices of a stack, you cannot separate the stacked devices using the Firepower Management Center web interface. In this case, use the auxiliary CLI command `configure stacking disable` to remove the stack configuration from each device individually.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device stack you want to break, click the break stack icon ( ).

Tip To remove a secondary device from a stack of three or more Firepower 8250 devices without breaking the stack, click the remove from stack icon ( ). Removing the secondary device causes a brief disruption of traffic inspection, traffic flow, or link state as the system reconfigures the stack for operation without the extra device.

Step 3 Click **Yes** to separate the device stack.

Replacing a Device in a Stack

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Any	FirePOWER 8140, 8200 family, 8300 family	Any	Admin/Network Admin

If the Firepower Management Center cannot communciate with the device, you must connect to the device and use CLI commands to separate the stack and unregister the device. For more information, see **stacking disable** and **delete** CLI commands in the relevant chapter: Classic Device CLI Configuration Commands, on page 2464.

To replace a device within a stack:

Procedure

Step 1 Select the stack with the device to replace and break that stack. For more information, see Separating Stacked Devices, on page 540.

Step 2 Unregister the device from the Firepower Management Center. For more information, see Delete a Device from the FMC, on page 231.

Step 3 Register the replacement device to the Firepower Management Center. For more information, see Add a Device to the FMC, on page 229.

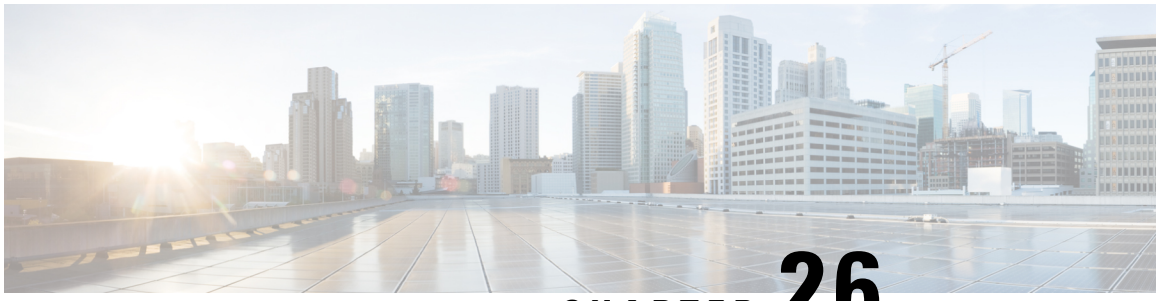
Step 4 Create a device stack that includes the replacement deivce. for more information, see Establishing Device Stacks, on page 536.



PART VII

Firepower Threat Defense Getting Started

- Transparent or Routed Firewall Mode for Firepower Threat Defense, on page 545
- Logical Devices for the Firepower Threat Defense on the Firepower 4100/9300, on page 557
- FlexConfig Policies for FTD, on page 575



CHAPTER 26

Transparent or Routed Firewall Mode for Firepower Threat Defense

This chapter describes how to set the firewall mode to routed or transparent, as well as how the firewall works in each firewall mode.

Note The firewall mode only affects regular firewall interfaces, and not IPS-only interfaces such as inline sets or passive interfaces. IPS-only interfaces can be used in both firewall modes. See Inline Sets and Passive Interfaces for Firepower Threat Defense, on page 653 for more information about IPS-only interfaces. Inline sets might be familiar to you as "transparent inline sets," but the inline interface type is unrelated to the transparent firewall mode described in this chapter or the firewall-type interfaces.

- About the Firewall Mode, on page 545
- Default Settings, on page 553
- Guidelines for Firewall Mode, on page 553
- Set the Firewall Mode, on page 554

About the Firewall Mode

The FTD supports two firewall modes for regular firewall interfaces: Routed Firewall mode and Transparent Firewall mode.

About Routed Firewall Mode

In routed mode, the FTD device is considered to be a router hop in the network. Each interface that you want to route between is on a different subnet.

With Integrated Routing and Bridging, you can use a "bridge group" where you group together multiple interfaces on a network, and the FTD device uses bridging techniques to pass traffic between the interfaces. Each bridge group includes a Bridge Virtual Interface (BVI) to which you assign an IP address on the network. The FTD device routes between BVIs and regular routed interfaces. If you do not need clustering or EtherChannel member interfaces, you might consider using routed mode instead of transparent mode. In routed mode, you can have one or more isolated bridge groups like in transparent mode, but also have normal routed interfaces as well for a mixed deployment.

About Transparent Firewall Mode

Traditionally, a firewall is a routed hop and acts as a default gateway for hosts that connect to one of its screened subnets. A transparent firewall, on the other hand, is a Layer 2 firewall that acts like a “bump in the wire,” or a “stealth firewall,” and is not seen as a router hop to connected devices. However, like any other firewall, access control between interfaces is controlled, and all of the usual firewall checks are in place.

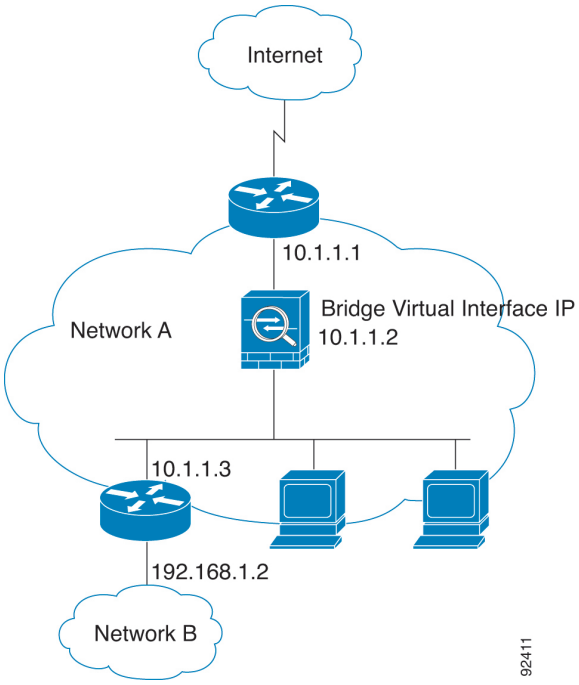
Layer 2 connectivity is achieved by using a "bridge group" where you group together the inside and outside interfaces for a network, and the FTD device uses bridging techniques to pass traffic between the interfaces. Each bridge group includes a Bridge Virtual Interface (BVI) to which you assign an IP address on the network. You can have multiple bridge groups for multiple networks. In transparent mode, these bridge groups cannot communicate with each other.

Using the Transparent Firewall in Your Network

The FTD device connects the same network between its interfaces. Because the firewall is not a routed hop, you can easily introduce a transparent firewall into an existing network.

The following figure shows a typical transparent firewall network where the outside devices are on the same subnet as the inside devices. The inside router and hosts appear to be directly connected to the outside router.

Figure 7: Transparent Firewall Network



Interface

In addition to each Bridge Virtual Interface (BVI) IP address, you can add a separate *slot*/*port* interface that is not part of any bridge group, and that allows only management traffic to the FTD device.

Passing Traffic For Routed-Mode Features

For features that are not directly supported on the transparent firewall, you can allow traffic to pass through so that upstream and downstream routers can support the functionality. For example, by using an access rule, you can allow DHCP traffic (instead of the unsupported DHCP relay feature) or multicast traffic such as that created by IP/TV. You can also establish routing protocol adjacencies through a transparent firewall; you can allow OSPF, RIP, EIGRP, or BGP traffic through based on an access rule. Likewise, protocols like HSRP or VRRP can pass through the FTD device.

About Bridge Groups

A bridge group is a group of interfaces that the FTD device bridges instead of routes. Bridge groups are supported in both transparent and routed firewall mode. Like any other firewall interfaces, access control between interfaces is controlled, and all of the usual firewall checks are in place.

Bridge Virtual Interface (BVI)

Each bridge group includes a Bridge Virtual Interface (BVI). The FTD device uses the BVI IP address as the source address for packets originating from the bridge group. The BVI IP address must be on the same subnet as the bridge group member interfaces. The BVI does not support traffic on secondary networks; only traffic on the same network as the BVI IP address is supported.

In transparent mode: Only bridge group member interfaces are named and can be used with interface-based features.

In routed mode: The BVI acts as the gateway between the bridge group and other routed interfaces. To route between bridge groups/routed interfaces, you must name the BVI. For some interface-based features, you can use the BVI itself:

- DHCPv4 server—Only the BVI supports the DHCPv4 server configuration.
- Static routes—You can configure static routes for the BVI; you cannot configure static routes for the member interfaces.
- Syslog server and other traffic sourced from the FTD device—When specifying a syslog server (or SNMP server, or other service where the traffic is sourced from the FTD device), you can specify either the BVI or a member interface.

If you do not name the BVI in routed mode, then the FTD device does not route bridge group traffic. This configuration replicates transparent firewall mode for the bridge group. If you do not need clustering or EtherChannel member interfaces, you might consider using routed mode instead. In routed mode, you can have one or more isolated bridge groups like in transparent mode, but also have normal routed interfaces as well for a mixed deployment.

Bridge Groups in Transparent Firewall Mode

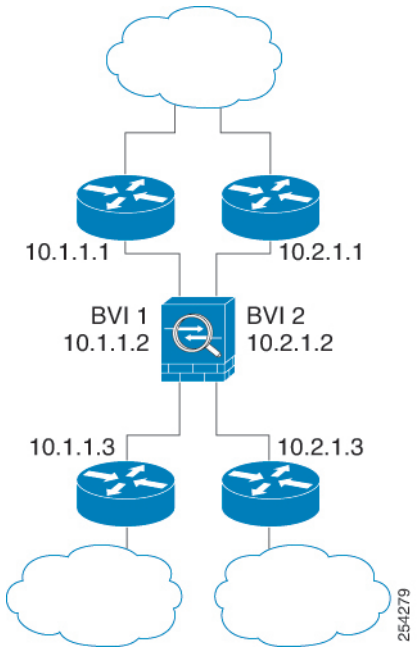
Bridge group traffic is isolated from other bridge groups; traffic is not routed to another bridge group within the FTD device, and traffic must exit the FTD device before it is routed by an external router back to another bridge group in the FTD device. Although the bridging functions are separate for each bridge group, many other functions are shared between all bridge groups. For example, all bridge groups share a syslog server or AAA server configuration.

You can include multiple interfaces per bridge group. See Guidelines for Firewall Mode, on page 553 for the exact number of bridge groups and interfaces supported. If you use more than 2 interfaces per bridge group,

you can control communication between multiple segments on the same network, and not just between inside and outside. For example, if you have three inside segments that you do not want to communicate with each other, you can put each segment on a separate interface, and only allow them to communicate with the outside interface. Or you can customize the access rules between interfaces to allow only as much access as desired.

The following figure shows two networks connected to the FTD device, which has two bridge groups.

Figure 8: Transparent Firewall Network with Two Bridge Groups

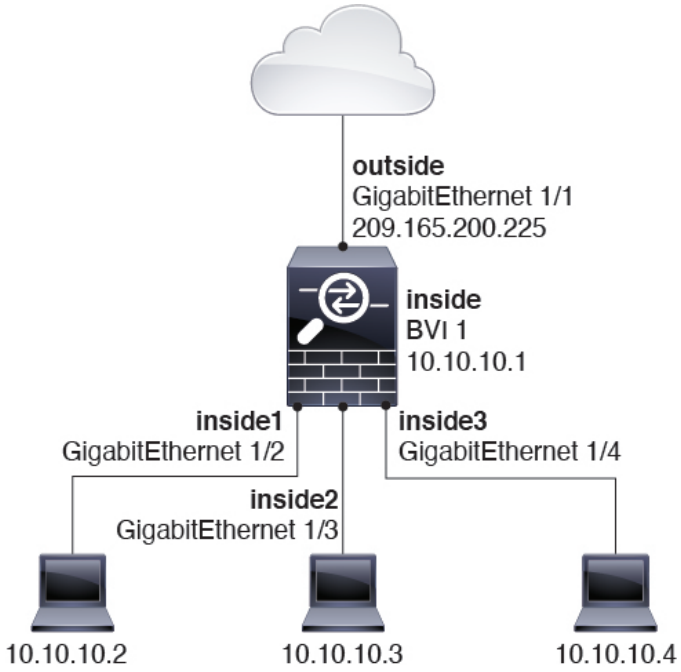


Bridge Groups in Routed Firewall Mode

Bridge group traffic can be routed to other bridge groups or routed interfaces. You can choose to isolate bridge group traffic by not assigning a name to the BVI interface for the bridge group. If you name the BVI, then the BVI participates in routing like any other regular interface.

One use for a bridge group in routed mode is to use extra interfaces on the FTD instead of an external switch. For example, the default configuration for some devices include an outside interface as a regular interface, and then all other interfaces assigned to the inside bridge group. Because the purpose of this bridge group is to replace an external switch, you need to configure an access policy so all bridge group interfaces can freely communicate.

Figure 9: Routed Firewall Network with an Inside Bridge Group and an Outside Routed Interface



Allowing Layer 3 Traffic

- Unicast IPv4 and IPv6 traffic requires an access rule to be allowed through the bridge group.
- ARPs are allowed through the bridge group in both directions without an access rule. ARP traffic can be controlled by ARP inspection.
- IPv6 neighbor discovery and router solicitation packets can be passed using access rules.
- Broadcast and multicast traffic can be passed using access rules.

Allowed MAC Addresses

The following destination MAC addresses are allowed through the bridge group if allowed by your access policy (see Allowing Layer 3 Traffic, on page 549). Any MAC address not on this list is dropped.

- TRUE broadcast destination MAC address equal to FFFF.FFFF.FFFF
- IPv4 multicast MAC addresses from 0100.5E00.0000 to 0100.5EFE.FFFF
- IPv6 multicast MAC addresses from 3333.0000.0000 to 3333.FFFF.FFFF
- BPDU multicast address equal to 0100.0CCC.CCCD

BPDU Handling

To prevent loops using the Spanning Tree Protocol, BPDUs are passed by default.

MAC Address vs. Route Lookups

For traffic within a bridge group, the outgoing interface of a packet is determined by performing a destination MAC address lookup instead of a route lookup.

Route lookups, however, are necessary for the following situations:

- Traffic originating on the FTD device—Add a default/static route on the FTD device for traffic destined for a remote network where a syslog server, for example, is located.
- Voice over IP (VoIP) and TFTP traffic, and the endpoint is at least one hop away—Add a static route on the FTD device for traffic destined for the remote endpoint so that secondary connections are successful. The FTD device creates a temporary "pinhole" in the access control policy to allow the secondary connection; and because the connection might use a different set of IP addresses than the primary connection, the FTD device needs to perform a route lookup to install the pinhole on the correct interface.

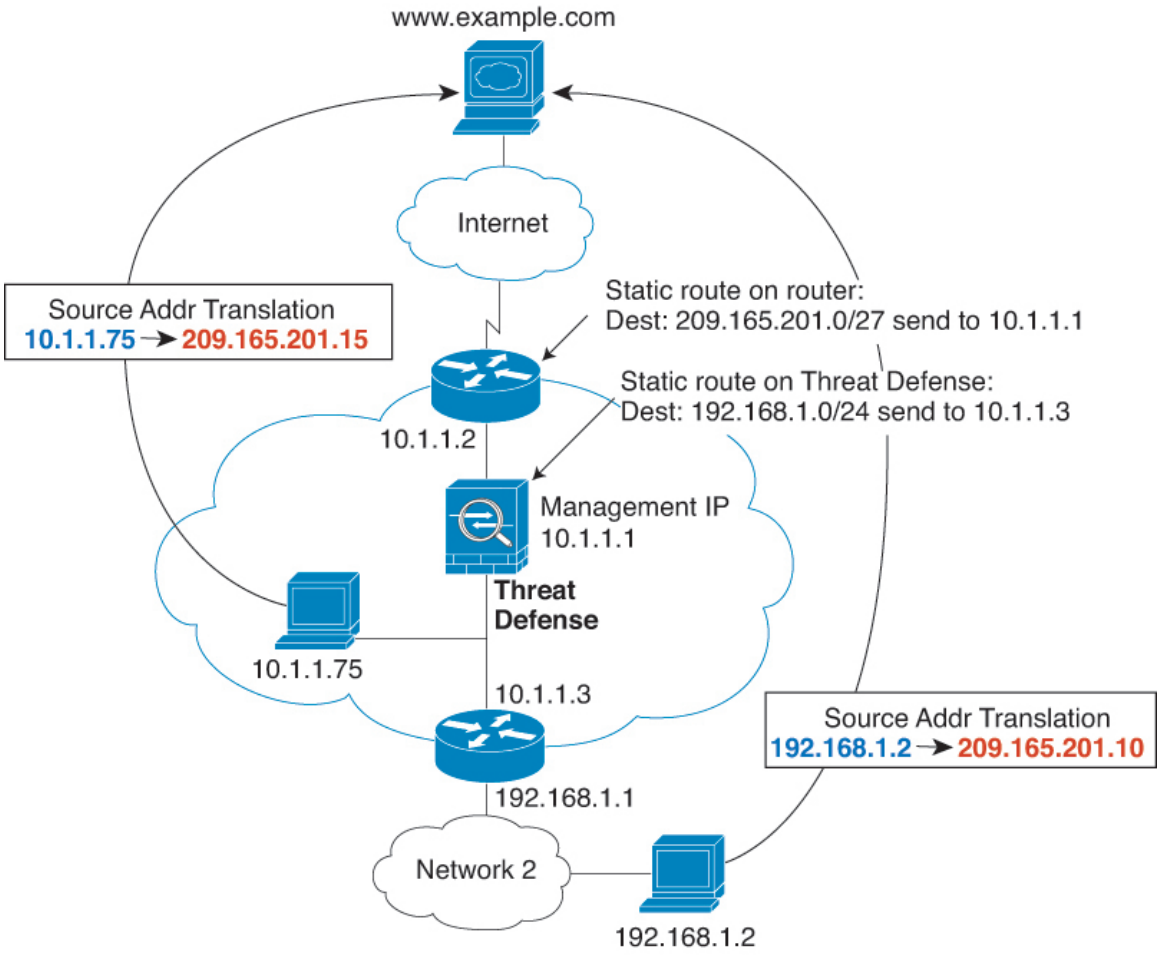
  Affected applications include:

  - H.323
  - RTSP
  - SIP
  - Skinny (SCCP)
  - SQL*Net
  - SunRPC
  - TFTP

- Traffic at least one hop away for which the FTD device performs NAT—Configure a static route on the FTD device for traffic destined for the remote network. You also need a static route on the upstream router for traffic destined for the mapped addresses to be sent to the FTD device.

  This routing requirement is also true for embedded IP addresses for VoIP and DNS with NAT enabled, and the embedded IP addresses are at least one hop away. The FTD device needs to identify the correct egress interface so it can perform the translation.

Figure 10: NAT Example: NAT within a Bridge Group



Unsupported Features for Bridge Groups in Transparent Mode

The following table lists the features are not supported in bridge groups in transparent mode.

Table 61: Unsupported Features in Transparent Mode

Feature	Description
Dynamic DNS	—
DHCP relay	The transparent firewall can act as a DHCPv4 server, but it does not support DHCP relay. DHCP relay is not required because you can allow DHCP traffic to pass through using two access rules: one that allows DCHP requests from the inside interface to the outside, and one that allows the replies from the server in the other direction.
Dynamic routing protocols	You can, however, add static routes for traffic originating on the FTD device for bridge group member interfaces. You can also allow dynamic routing protocols through the FTD device using an access rule.

Feature	Description
Multicast IP routing	You can allow multicast traffic through the FTD device by allowing it in an access rule.
QoS	—
VPN termination for through traffic	The transparent firewall supports site-to-site VPN tunnels for management connections only on bridge group member interfaces. It does not terminate VPN connections for traffic through the FTD device. You can pass VPN traffic through the ASA using an access rule, but it does not terminate non-management connections.

Unsupported Features for Bridge Groups in Routed Mode

The following table lists the features are not supported in bridge groups in routed mode.

Table 62: Unsupported Features in Routed Mode

Feature	Description
EtherChannel member interfaces	Only physical interfaces, redundant interfaces, and subinterfaces are supported as bridge group member interfaces. interfaces are also not supported.
Clustering	Bridge groups are not supported in clustering.
Dynamic DNS	—
DHCP relay	The routed firewall can act as a DHCPv4 server, but it does not support DHCP relay on BVIs or bridge group member interfaces.
Dynamic routing protocols	You can, however, add static routes for BVIs. You can also allow dynamic routing protocols through the FTD device using an access rule. Non-bridge group interfaces support dynamic routing.
Multicast IP routing	You can allow multicast traffic through the FTD device by allowing it in an access rule. Non-bridge group interfaces support multicast routing.
QoS	Non-bridge group interfaces support QoS.
VPN termination for through traffic	You cannot terminate a VPN connection on the BVI. Non-bridge group interfaces support VPN. Bridge group member interfaces support site-to-site VPN tunnels for management connections only. It does not terminate VPN connections for traffic through the FTD device. You can pass VPN traffic through the bridge group using an access rule, but it does not terminate non-management connections.

Default Settings

Bridge Group Defaults

By default, all ARP packets are passed within the bridge group.

Guidelines for Firewall Mode

Bridge Group Guidelines (Transparent and Routed Mode)

- You can create up to 250 bridge groups, with 64 interfaces per bridge group.
- Each directly-connected network must be on the same subnet.
- The FTD device does not support traffic on secondary networks; only traffic on the same network as the BVI IP address is supported.
- An IP address for the BVI is required for each bridge group for to-the-device and from-the-device management traffic, as well as for data traffic to pass through the FTD device. For IPv4 traffic, specify an IPv4 address. For IPv6 traffic, specify an IPv6 address.
- You can only configure IPv6 addresses manually.
- The BVI IP address must be on the same subnet as the connected network. You cannot set the subnet to a host subnet (255.255.255.255).
- Management interfaces are not supported as bridge group members.
- For the FTDv on VMware with bridged ixgbevf interfaces, transparent mode is not supported, and bridge groups are not supported in routed mode.
- In transparent mode, you must use at least 1 bridge group; data interfaces must belong to a bridge group.
- In transparent mode, do not specify the BVI IP address as the default gateway for connected devices; devices need to specify the router on the other side of the FTD as the default gateway.
- In transparent mode, the *default* route, which is required to provide a return path for management traffic, is only applied to management traffic from one bridge group network. This is because the default route specifies an interface in the bridge group as well as the router IP address on the bridge group network, and you can only define one default route. If you have management traffic from more than one bridge group network, you need to specify a regular static route that identifies the network from which you expect management traffic.
- In transparent mode, PPPoE is not supported for the interface.
- Transparent mode is not supported on threat defense virtual instances deployed on Amazon Web Services, Microsoft Azure, Google Cloud Platform, and Oracle Cloud Infrastructure.
- In routed mode, to route between bridge groups and other routed interfaces, you must name the BVI.
- In routed mode, FTD-defined EtherChannel interfaces are not supported as bridge group members. EtherChannels on the Firepower 4100/9300 can be bridge group members.

- Bidirectional Forwarding Detection (BFD) echo packets are not allowed through the FTD when using bridge group members. If there are two neighbors on either side of the FTD running BFD, then the FTD will drop BFD echo packets because they have the same source and destination IP address and appear to be part of a LAND attack.

Set the Firewall Mode

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	N/A	Firepower Threat Defense	Any	Admin Access Admin Network Admin

You can set the firewall mode when you perform the initial system setup at the CLI. We recommend setting the firewall mode during setup because changing the firewall mode erases your configuration to ensure you do not have incompatible settings. If you need to change the firewall mode later, you must do so from the CLI.

Procedure

Step 1 Deregister the Firepower Threat Defense device from the FMC.

You cannot change the mode until you deregister the device.

a) Choose **Devices** > **Device Management**.
b) Select the device from the list of managed devices.
c) Delete the device (click Trash can), confirm, and wait for system to remove the device.

Step 2 Access the Firepower Threat Defense device CLI, preferably from the console port.

If you use SSH to the diagnostic interface, then changing the mode erases your interface configuration and you will be disconnected. You should instead connect to the management interface.

Step 3 Change the firewall mode:

configure firewall [**routed** | **transparent**]

Example:

```
> configure firewall transparent
This will destroy the current interface configurations, are you sure that you want to
proceed? [y/N] y
The firewall mode was changed successfully.
```

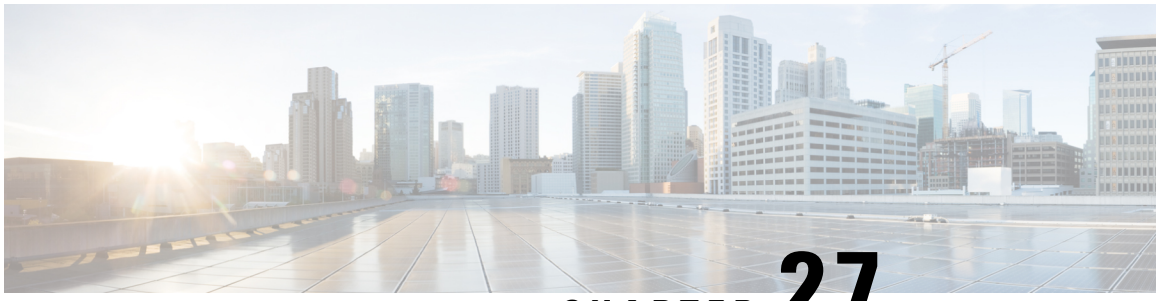
Step 4 Re-register with the FMC:

configure manager add {*hostname* | *ip_address* | **DONTRESOLVE**} *reg_key* [*nat_id*]

where:

- {*hostname* | *ip_address* | **DONTRESOLVE** } specifies either the fully qualified host name or IP address of the FMC. If the FMC is not directly addressable, use **DONTRESOLVE**.

- *reg_key* is the unique alphanumeric registration key required to register a device to the FMC.
- *nat_id* is an optional alphanumeric string used during the registration process between the FMC and the device. It is required if the hostname is set to **DONTRESOLVE**.



CHAPTER 27

Logical Devices for the Firepower Threat Defense on the Firepower 4100/9300

The Firepower 4100/9300 is a flexible security platform on which you can install one or more *logical devices*. Before you can add the Firepower Threat Defense to the FMC, you must configure chassis interfaces, add a logical device, and assign interfaces to the device on the Firepower 4100/9300 chassis using the Firepower Chassis Manager or the FXOS CLI. This chapter describes basic interface configuration and how to add a standalone or High Availability logical device using the Firepower Chassis Manager. To add a clustered logical device, see Clustering for the Firepower Threat Defense, on page 709. To use the FXOS CLI, see the FXOS CLI configuration guide. For more advanced FXOS procedures and troubleshooting, see the FXOS configuration guide.

- About Interfaces, on page 557
- About Logical Devices, on page 559
- Requirements and Prerequisites for Logical Devices, on page 560
- Guidelines and Limitations for Logical Devices, on page 561
- Configure Interfaces, on page 562
- Configure Logical Devices, on page 565
- History for Firepower Threat Defense Logical Devices, on page 572

About Interfaces

The Firepower 4100/9300 chassis supports physical interfaces and EtherChannel (port-channel) interfaces. EtherChannel interfaces can include up to 16 member interfaces of the same type.

Chassis Management Interface

The chassis management interface is used for management of the FXOS Chassis by SSH or Firepower Chassis Manager. This interface appears at the top of the **Interfaces** tab as **MGMT**, and you can only enable or disable this interface on the **Interfaces** tab. This interface is separate from the mgmt-type interface that you assign to the logical devices for application management.

To configure parameters for this interface, you must configure them from the CLI. To view information about this interface in the FXOS CLI, connect to local management and show the management port:

Firepower # **connect local-mgmt**

Firepower(local-mgmt) # **show mgmt-port**

Note that the chassis management interface remains up even if the physical cable or SFP module are unplugged, or if the **mgmt-port shut** command is performed.

Note The chassis management interface does not support jumbo frames.

Interface Types

Physical interfaces and EtherChannel (port-channel) interfaces can be one of the following types:

- Data—Use for regular data. Data interfaces cannot be shared between logical devices, and logical devices cannot communicate over the backplane to other logical devices. For traffic on Data interfaces, all traffic must exit the chassis on one interface and return on another interface to reach another logical device.
- Mgmt—Use to manage application instances. These interfaces can be shared by one or more logical devices to access external hosts; logical devices cannot communicate over this interface with other logical devices that share the interface. You can only assign one management interface per logical device. Depending on your application and manager, you can later enable management from a data interface; but you must assign a Management interface to the logical device even if you don't intend to use it after you enable data management.

Note Mgmt interface change will cause reboot of the logical device, for example one change mgmt from e1/1 to e1/2 will cause the logical device to reboot to apply the new management.

- Eventing—Use as a secondary management interface for Firepower Threat Defense devices. To use this interface, you must configure its IP address and other parameters at the Firepower Threat Defense CLI. For example, you can separate management traffic from events (such as web events). See the management center configuration guide for more information. Eventing interfaces can be shared by one or more logical devices to access external hosts; logical devices cannot communicate over this interface with other logical devices that share the interface.

Note A virtual Ethernet interface is allocated when each application instance is installed. If the application does not use an eventing interface, then the virtual interface will be in an admin down state.

```
Firepower # show interface Vethernet775
Firepower # Vethernet775 is down (Administratively down)
Bound Interface is Ethernet1/10
Port description is server 1/1, VNIC ext-mgmt-nic5
```

- Cluster—Use as the cluster control link for a clustered logical device. By default, the cluster control link is automatically created on Port-channel 48. The Cluster type is only supported on EtherChannel interfaces.

FXOS Interfaces vs. Application Interfaces

The Firepower 4100/9300 manages the basic Ethernet settings of physical interfaces and EtherChannel (port-channel) interfaces. Within the application, you configure higher level settings. For example, you can only create EtherChannels in FXOS; but you can assign an IP address to the EtherChannel within the application.

The following sections describe the interaction between FXOS and the application for interfaces.

VLAN Subinterfaces

For all logical devices, you can create VLAN subinterfaces within the application.

Independent Interface States in the Chassis and in the Application

You can administratively enable and disable interfaces in both the chassis and in the application. For an interface to be operational, the interface must be enabled in both operating systems. Because the interface state is controlled independently, you may have a mismatch between the chassis and application.

Inline Set Link State Propagation for the FTD

An inline set acts like a bump on the wire, and binds two interfaces together to slot into an existing network. This function allows the system to be installed in any network environment without the configuration of adjacent network devices. Inline interfaces receive all traffic unconditionally, but all traffic received on these interfaces is retransmitted out of an inline set unless explicitly dropped.

When you configure an inline set in the Firepower Threat Defense application and enable link state propagation, the Firepower Threat Defense sends inline set membership to the FXOS chassis. Link state propagation means that the chassis automatically brings down the second interface in the inline interface pair when one of the interfaces in an inline set goes down. When the downed interface comes back up, the second interface automatically comes back up, also. In other words, if the link state of one interface changes, the chassis senses the change and updates the link state of the other interface to match it. Note that the chassis requires up to 4 seconds to propagate link state changes. Link state propagation is especially useful in resilient network environments where routers are configured to reroute traffic automatically around network devices that are in a failure state.



Note Do not enable Hardware Bypass and link state propagation for the same inline set.

About Logical Devices

A logical device lets you run one application instance (either ASA or Firepower Threat Defense) and also one optional decorator application (Radware DefensePro) to form a service chain.

When you add a logical device, you also define the application instance type and version, assign interfaces, and configure bootstrap settings that are pushed to the application configuration.

Note For the Firepower 9300, you must install the same application instance type (ASA or Firepower Threat Defense) on all modules in the chassis; different types are not supported at this time. Note that modules can run different versions of an application instance type.

Standalone and Clustered Logical Devices

You can add the following logical device types:

- Standalone—A standalone logical device operates as a standalone unit or as a unit in a High Availability pair.
- Cluster—A clustered logical device lets you group multiple units together, providing all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices. Multiple module devices, like the Firepower 9300, support intra-chassis clustering. For the Firepower 9300, all three modules must participate in the cluster.

Requirements and Prerequisites for Logical Devices

See the following sections for requirements and prerequisites.

Requirements and Prerequisites for Hardware and Software Combinations

The Firepower 4100/9300 supports multiple models, security modules, application types, and high availability and scalability features. See the following requirements for allowed combinations.

Firepower 9300 Requirements

The Firepower 9300 includes 3 security module slots and multiple types of security modules. See the following requirements:

- Security Module Types—All modules in the Firepower 9300 must be the same type.
- Clustering—All security modules in the cluster, whether it is intra-chassis or inter-chassis, must be the same type. You can have different quantities of installed security modules in each chassis, although all modules present in the chassis must belong to the cluster including any empty slots. For example, you can install 2 SM-40s in chassis 1, and 3 SM-40s in chassis 2.
- High Availability—High Availability is only supported between same-type modules on the Firepower 9300.
- ASA and FTD application types—You can only install one application type on the chassis, ASA or FTD.
- ASA or FTD versions—You can run different versions of an application instance type on separate modules. For example, you can install the FTD 6.3 on module 1, FTD 6.4 on module 2, and FTD 6.5 on module 3.

Firepower 4100 Requirements

The Firepower 4100 comes in multiple models. See the following requirements:

- Clustering—All chassis in the cluster must be the same model.
- High Availability—High Availability is only supported between same-type models.
- ASA and FTD application types—The Firepower 4100 can only run a single application type.

Requirements and Prerequisites for High Availability

- The two units in a High Availability Failover configuration must:
  - Be on a separate chassis; intra-chassis High Availability for the Firepower 9300 is not supported.
  - Be the same model.
  - Have the same interfaces assigned to the High Availability logical devices.
  - Have the same number and types of interfaces. All interfaces must be preconfigured in FXOS identically before you enable High Availability.
- For High Availability system requirements, see High Availability System Requirements, on page 683.

Guidelines and Limitations for Logical Devices

See the following sections for guidelines and limitations.

Guidelines and Limitations for Interfaces

Inline Sets for FTD

- Supported for physical interfaces (both regular and breakout ports) and EtherChannels.
- Link state propagation is supported.
- Do not enable Hardware Bypass and link state propagation for the same inline set.

Hardware Bypass

- Supported for the Firepower Threat Defense; you can use them as regular interfaces for the ASA.
- The Firepower Threat Defense only supports Hardware Bypass with inline sets.
- Hardware Bypass-capable interfaces cannot be configured for breakout ports.
- You cannot include Hardware Bypass interfaces in an EtherChannel and use them for Hardware Bypass; you can use them as regular interfaces in an EtherChannel.
- Hardware Bypass is not supported with High Availability.
- Do not enable Hardware Bypass and link state propagation for the same inline set.

Default MAC Addresses

Default MAC address assignments depend on the type of interface.

- Physical interfaces—The physical interface uses the burned-in MAC address.
- EtherChannels—For an EtherChannel, all interfaces that are part of the channel group share the same MAC address. This feature makes the EtherChannel transparent to network applications and users, because they only see the one logical connection; they have no knowledge of the individual links. The port-channel interface uses a unique MAC address from a pool; interface membership does not affect the MAC address.

General Guidelines and Limitations

Firewall Mode

You can set the firewall mode to routed or transparent in the bootstrap configuration for the Firepower Threat Defense.

High Availability

- Configure high availability within the application configuration.
- You can use any data interfaces as the failover and state links.

Configure Interfaces

By default, physical interfaces are disabled. You can enable interfaces, add EtherChannels, and edit interface properties.

Enable or Disable an Interface

You can change the **Admin State** of each interface to be enabled or disabled. By default, physical interfaces are disabled.

Procedure

Step 1 Choose **Interfaces** to open the Interfaces page.

The Interfaces page shows a visual representation of the currently installed interfaces at the top of the page and provides a listing of the installed interfaces in the table below.

Step 2 To enable the interface, click the disabled **Slider disabled** () so that it changes to the enabled **Slider enabled** ().

Click **Yes** to confirm the change. The corresponding interface in the visual representation changes from gray to green.

Step 3 To disable the interface, click the enabled **Slider enabled** () so that it changes to the disabled **Slider disabled** ().

Click **Yes** to confirm the change. The corresponding interface in the visual representation changes from green to gray.

Configure a Physical Interface

You can physically enable and disable interfaces, as well as set the interface speed and duplex. To use an interface, it must be physically enabled in FXOS and logically enabled in the application.

Note For QSFPH40G-CUxM, auto-negotiation is always enabled by default and you cannot disable it.

Before you begin

- Interfaces that are already a member of an EtherChannel cannot be modified individually. Be sure to configure settings before you add it to the EtherChannel.

Procedure

Step 1 Choose **Interfaces** to open the Interfaces page.

The **All Interfaces** page shows a visual representation of the currently installed interfaces at the top of the page and provides a listing of the installed interfaces in the table below.

Step 2 Click **Edit** in the row for the interface you want to edit to open the **Edit Interface** dialog box.

Step 3 To enable the interface, check the **Enable** check box. To disable the interface, uncheck the **Enable** check box.

Step 4 Choose the interface **Type**:

- **Data**
- **Mgmt**
- **Firepower-eventing**—For Firepower Threat Defense only.
- **Cluster**—Do not choose the **Cluster** type; by default, the cluster control link is automatically created on Port-channel 48.

Step 5 (Optional) Choose the speed of the interface from the **Speed** drop-down list.

Step 6 (Optional) If your interface supports **Auto Negotiation**, click the **Yes** or **No** radio button.

Step 7 (Optional) Choose the duplex of the interface from the **Duplex** drop-down list.

Step 8 (Optional) Explicitly configure **Debounce Time (ms)**. Enter a value between 0-15000 milli-seconds.

Step 9 Click **OK**.

Add an EtherChannel (Port Channel)

An EtherChannel (also known as a port channel) can include up to 16 member interfaces of the same media type and capacity, and must be set to the same speed and duplex. The media type can be either RJ-45 or SFP; SFPs of different types (copper and fiber) can be mixed. You cannot mix interface capacities (for example 1GB and 10GB interfaces) by setting the speed to be lower on the larger-capacity interface. The Link Aggregation Control Protocol (LACP) aggregates interfaces by exchanging the Link Aggregation Control Protocol Data Units (LACPDUs) between two network devices.

The Firepower 4100/9300 chassis only supports EtherChannels in Active LACP mode so that each member interface sends and receives LACP updates. An active EtherChannel can establish connectivity with either an active or a passive EtherChannel. You should use the active mode unless you need to minimize the amount of LACP traffic.

LACP coordinates the automatic addition and deletion of links to the EtherChannel without user intervention. It also handles misconfigurations and checks that both ends of member interfaces are connected to the correct channel group.

When the Firepower 4100/9300 chassis creates an EtherChannel, the EtherChannel stays in a **Suspended** state until you assign it to a logical device, even if the physical link is up. The EtherChannel will be brought out of this **Suspended** state in the following situations:

- The EtherChannel is added as a data or management interface for a standalone logical device
- The EtherChannel is added as a management interface or cluster control link for a logical device that is part of a cluster
- The EtherChannel is added as a data interface for a logical device that is part of a cluster and at least one unit has joined the cluster

Note that the EtherChannel does not come up until you assign it to a logical device. If the EtherChannel is removed from the logical device or the logical device is deleted, the EtherChannel will revert to a **Suspended** state.

Procedure

Step 1 Choose **Interfaces** to open the Interfaces page.

The **All Interfaces** page shows a visual representation of the currently installed interfaces at the top of the page and provides a listing of the installed interfaces in the table below.

Step 2 Click **Add Port Channel** above the interfaces table to open the **Add Port Channel** dialog box.

Step 3 Enter an ID for the port channel in the **Port Channel ID** field. Valid values are between 1 and 47.

Port-channel 48 is reserved for the cluster control link when you deploy a clustered logical device. If you do not want to use Port-channel 48 for the cluster control link, you can delete it and configure a Cluster type EtherChannel with a different ID. For intra-chassis clustering, do not assign any interfaces to the Cluster EtherChannel.

Step 4 To enable the port channel, check the **Enable** check box. To disable the port channel, uncheck the **Enable** check box.

Step 5 Choose the interface **Type**:

- **Data**
- **Mgmt**
- **Firepower-eventing**—For Firepower Threat Defense only.
- **Cluster**

Step 6 Set the required **Admin Speed** for the member interfaces from the drop-down list.

If you add a member interface that is not at the specified speed, it will not successfully join the port channel.

Step 7 Set the required **Admin Duplex** for the member interfaces, **Full Duplex** or **Half Duplex**.

If you add a member interface that is configured with the specified duplex, it will not successfully join the port channel.

Step 8 To add an interface to the port channel, select the interface in the **Available Interface** list and click **Add Interface** to move the interface to the Member ID list.

You can add up to 16 member interfaces of the same media type and capacity. The member interfaces must be set to the same speed and duplex, and must match the speed and duplex that you configured for this port channel. The media type can be either RJ-45 or SFP; SFPs of different types (copper and fiber) can be mixed. You cannot mix interface capacities (for example 1GB and 10GB interfaces) by setting the speed to be lower on the larger-capacity interface.

Tip You can add multiple interfaces at one time. To select multiple individual interfaces, click on the desired interfaces while holding down the **Ctrl** key. To select a range of interfaces, select the first interface in the range, and then, while holding down the **Shift** key, click to select the last interface in the range.

Step 9 To remove an interface from the port channel, click the **Delete** button to the right of the interface in the Member ID list.

Step 10 Click **OK**.

Configure Logical Devices

Add a standalone logical device or a High Availability pair on the Firepower 4100/9300 chassis.

For clustering, see Clustering for the Firepower Threat Defense, on page 709.

Add a Standalone FTD

Standalone logical devices work either alone or in a High Availability pair. On the Firepower 9300 with multiple security modules, you can deploy either a cluster or standalone devices. The cluster must use all modules, so you cannot mix and match a 2-module cluster plus a single standalone device, for example.

Before you begin

- Download the application image you want to use for the logical device from Cisco.com, and then upload that image to the Firepower 4100/9300 chassis.

Note For the Firepower 9300, you must install the same application instance type (ASA or FTD) on all modules in the chassis; different types are not supported at this time. Note that modules can run different versions of an application instance type.

- Configure a management interface to use with the logical device. The management interface is required. Note that this management interface is not the same as the chassis management port that is used only for chassis management (and that appears at the top of the **Interfaces** tab as **MGMT**).
- You must also configure at least one Data type interface. Optionally, you can also create a firepower-eventing interface to carry all event traffic (such as web events). See Interface Types, on page 558 for more information.
- Gather the following information:
  - Interface IDs for this device
  - Management interface IP address and network mask
  - Gateway IP address
  - FMC IP address and/or NAT ID of your choosing
  - DNS server IP address
  - Firepower Threat Defense hostname and domain name

Procedure

Step 1 Choose **Logical Devices**.

Step 2 Click **Add Device**, and set the following parameters:

a) Provide a **Device Name**.

This name is used by the chassis supervisor to configure management settings and to assign interfaces; it is not the device name used in the application configuration.

Note You cannot change this name after you add the logical device.

b) For the **Template**, choose **Cisco Firepower Threat Defense**.
c) Choose the **Image Version**.
d) For the **Usage**, click the **Standalone** radio button.
e) Click **OK**.

You see the Provisioning - *device name* window.

Step 3 Expand the **Data Ports** area, and click each interface that you want to assign to the device.

You can only assign data interfaces that you previously enabled on the **Interfaces** page. You will later enable and configure these interfaces in FMC, including setting the IP addresses.

Hardware Bypass-capable ports are shown with the following icon: . For certain interface modules, you can enable the Hardware Bypass feature for Inline Set interfaces only (see the FMC configuration guide). Hardware Bypass ensures that traffic continues to flow between an inline interface pair during a power outage. This feature can be used to maintain network connectivity in the case of software or hardware failures. If you do not assign both interfaces in a Hardware Bypass pair, you see a warning message to make sure your assignment is intentional. You do not need to use the Hardware Bypass feature, so you can assign single interfaces if you prefer.

Step 4 Click the device icon in the center of the screen.

A dialog box appears where you can configure initial bootstrap settings. These settings are meant for initial deployment only, or for disaster recovery. For normal operation, you can later change most values in the application CLI configuration.

Step 5 On the **General Information** page, complete the following:

a) (For the Firepower 9300) Under **Security Module Selection** click the security module that you want to use for this logical device.

b) Choose the **Management Interface**.

This interface is used to manage the logical device. This interface is separate from the chassis management port.

c) Choose the management interface **Address Type**: **IPv4 only**, **IPv6 only**, or **IPv4 and IPv6**.

d) Configure the **Management IP** address.

Set a unique IP address for this interface.

e) Enter a **Network Mask** or **Prefix Length**.

f) Enter a **Network Gateway** address.

Step 6 On the **Settings** tab, complete the following:

a) Enter the **Firepower Management Center IP** of the managing FMC.

b) Enter the **Search Domains** as a comma-separated list.

c) Choose the **Firewall Mode**: **Transparent** or **Routed**.

In routed mode, the Firepower Threat Defense is considered to be a router hop in the network. Each interface that you want to route between is on a different subnet. A transparent firewall, on the other hand, is a Layer 2 firewall that acts like a "bump in the wire," or a "stealth firewall," and is not seen as a router hop to connected devices.

The firewall mode is only set at initial deployment. If you re-apply the bootstrap settings, this setting is not used.

d) Enter the **DNS Servers** as a comma-separated list.

The Firepower Threat Defense uses DNS if you specify a hostname for the FMC, for example.

e) Enter the **Fully Qualified Hostname** for the Firepower Threat Defense.

f) Enter a **Registration Key** to be shared between the FMC and the device during registration.

You can choose any text string for this key between 1 and 37 characters; you will enter the same key on the FMC when you add the Firepower Threat Defense.

g) Enter a **Password** for the Firepower Threat Defense admin user for CLI access.

h) Choose the **Eventing Interface** on which events should be sent. If not specified, the management interface will be used.

This interface must be defined as a Firepower-eventing interface.

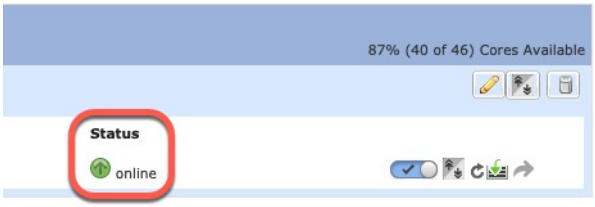
Step 7 On the **Agreement** tab, read and accept the end user license agreement (EULA).

Step 8 Click **OK** to close the configuration dialog box.

Step 9 Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for the status of the new logical device. When the logical device shows its **Status** as **online**, you can start configuring the security policy in the application.





Step 10 See the FMC configuration guide to add the Firepower Threat Defense as a managed device and start configuring your security policy.

Add a High Availability Pair

FTD High Availability (also known as failover) is configured within the application, not in FXOS. However, to prepare your chassis for high availability, see the following steps.

Before you begin

See Requirements and Prerequisites for High Availability, on page 561.

Procedure

Step 1 Allocate the same interfaces to each logical device.

Step 2 Allocate 1 or 2 data interfaces for the failover and state link(s).

These interfaces exchange high availability traffic between the 2 chassis. We recommend that you use a 10 GB data interface for a combined failover and state link. If you have available interfaces, you can use separate failover and state links; the state link requires the most bandwidth. You cannot use the management-type interface for the failover or state link. We recommend that you use a switch between the chassis, with no other device on the same network segment as the failover interfaces.

Step 3 Enable High Availability on the logical devices. See High Availability for FTD, on page 683.

Step 4 If you need to make interface changes after you enable High Availability, perform the changes on the standby unit first, and then perform the changes on the active unit.

Change an Interface on a FTD Logical Device

You can allocate or unallocate an interface, or replace a management interface on the Firepower Threat Defense logical device. You can then sync the interface configuration in the FMCthe .

Adding a new interface, or deleting an unused interface has minimal impact on the Firepower Threat Defense configuration. However, deleting an interface that is used in your security policy will impact the configuration. Interfaces can be referenced directly in many places in the Firepower Threat Defense configuration, including access rules, NAT, SSL, identity rules, VPN, DHCP server, and so on. Policies that refer to security zones are not affected. You can also edit the membership of an allocated EtherChannel without affecting the logical device or requiring a sync on the FMCthe .

Deleting an interface will delete any configuration associated with that interface.

Before you begin

- Configure your interfaces, and add any EtherChannels according to Configure a Physical Interface, on page 563 and Add an EtherChannel (Port Channel), on page 564.
- If you want to add an already-allocated interface to an EtherChannel (for example, all interfaces are allocated by default to a cluster), you need to unallocate the interface from the logical device first, then add the interface to the EtherChannel. For a new EtherChannel, you can then allocate the EtherChannel to the device.
- If you want to replace the management or eventing interface with a management EtherChannel, then you need to create the EtherChannel with at least 1 unallocated data member interface, and then replace the current management interface with the EtherChannel. After the FTD device reboots (management interface changes cause a reboot), and you sync the configuration in the FMCthe , you can add the (now unallocated) management interface to the EtherChannel as well.
- For clustering or High Availability, make sure you add or remove the interface on all units before you sync the configuration in the FMCthe . We recommend that you make the interface changes on the data/standby unit(s) first, and then on the control/active unit. Note that new interfaces are added in an administratively down state, so they do not affect interface monitoring.
- In mult-instance mode, for changing a sub-interface with an another sub-interface with the same vlan tag, you must first remove all the configuration (including nameif config) of the interface and then unalloacte the interface from Firepower Chassis Manager. Once unallocated, add the new interface and then use sync interfaces from the FMC.

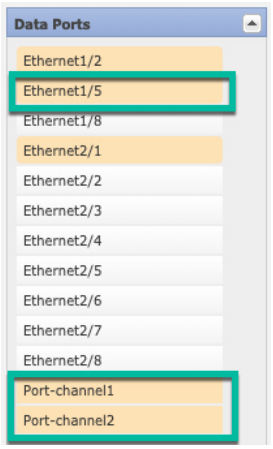
Procedure

Step 1 In the Firepower Chassis Manager, choose **Logical Devices**.

Step 2 Click the **Edit** icon at the top right to edit the logical device.

Step 3 Allocate a new data interface by selecting the interface in the **Data Ports** area.

Do not delete any interfaces yet.



Step 4 Replace the management or eventing interface:

For these types of interfaces, the device reboots after you save your changes.

a) Click the device icon in the center of the page.
b) On the **General** or **Cluster Information** tab, choose the new **Management Interface** from the drop-down list.
c) On the **Settings** tab, choose the new **Eventing Interface** from the drop-down list.
d) Click **OK**.

If you change the IP address of the Management interface, then you must also change the IP address for the device in the FMC: go to **Devices** > **Device Management** > **Device/Cluster**. In the **Management** area, set the IP address to match the bootstrap configuration address.

Step 5 Click **Save**.

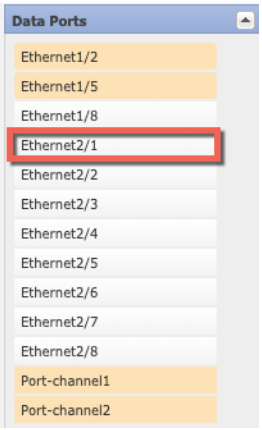
Step 6 Sync the interfaces in the FMC.

a) Log into the FMC.
b) Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.
c) Click the **Sync Device** button on the top left of the **Interfaces** page.
d) After the changes are detected, you will see a red banner on the **Interfaces** page indicating that the interface configuration has changed. Click the **Click to know more** link to view the interface changes.
e) If you plan to delete an interface, manually transfer any interface configuration from the old interface to the new interface.

   Because you have not yet deleted any interfaces, you can refer to the existing configuration.

f) Click **Save**.
g) Select the devices and click **Deploy** to deploy the policy to the assigned devices. The changes are not active until you deploy them.

Step 7 In the Firepower Chassis Manager, unallocate a data interface by de-selecting the interface in the **Data Ports** area.



Step 8 Click **Save**.

Step 9 Sync the interfaces again in the FMCthe .

Connect to the Console of the Application

Use the following procedure to connect to the console of the application.

Procedure

Step 1 Connect to the module CLI.

connect module *slot_number* **console**

To connect to the security engine of a device that does not support multiple security modules, always use **1** as the *slot_number*.

Example:

```
Firepower# connect module 1 console
Telnet escape character is '~'.
Trying 127.5.1.1...
Connected to 127.5.1.1.
Escape character is '~'.

CISCO Serial Over LAN:
Close Network Connection to Exit

Firepower-module1>
```

Step 2 Connect to the application console.

connect ftd

Example:

```
Firepower-module1> connect ftd
Connecting to ftd(ftd-native) console... enter exit to return to bootCLI
[...]
```

>

Step 3 Exit the application console to the FXOS module CLI.

- FTD—Enter **Ctrl-a, d**

Step 4 Return to the supervisor level of the FXOS CLI.

a) Enter **~**

You exit to the Telnet application.

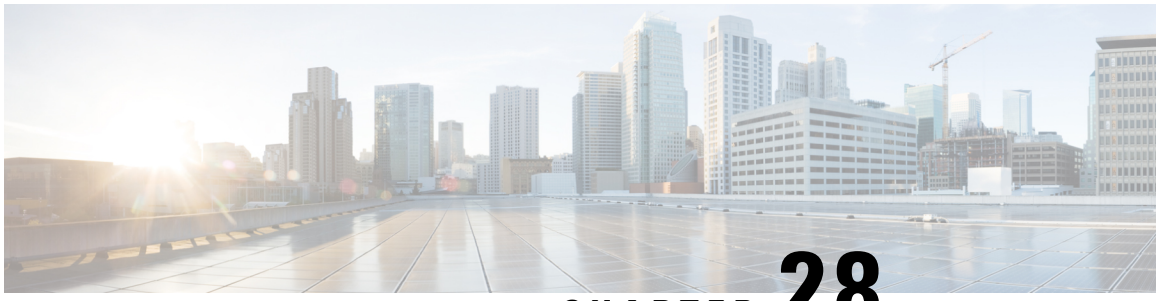
b) To exit the Telnet application, enter:

telnet>**quit**

History for Firepower Threat Defense Logical Devices

Feature	Version	Details
Support for EtherChannels in Firepower Threat Defense inline sets	6.2	You can now use EtherChannels in a Firepower Threat Defense inline set. Supported platforms: Firepower 4100/9300
Inter-chassis clustering for 6 Firepower Threat Defense modules	6.2	You can now enable inter-chassis clustering for the Firepower Threat Defense. You can include up to 6 modules in up to 6 chassis. New/modified Firepower Chassis Manager screens: • **Logical Devices** > **Configuration** Supported platforms: Firepower 4100/9300
Hardware bypass support on the Firepower 4100/9300 for supported network modules	6.1	Hardware Bypass ensures that traffic continues to flow between an inline interface pair during a power outage. This feature can be used to maintain network connectivity in the case of software or hardware failures. New/Modified screens: • **Devices** > **Device Management** > **Interfaces** > **Edit Physical Interface** Supported platforms: Firepower 4100/9300

Feature	Version	Details
Inline set link state propagation support for the Firepower Threat Defense	6.1	When you configure an inline set in the Firepower Threat Defense application and enable link state propagation, the Firepower Threat Defense sends inline set membership to the FXOS chassis. Link state propagation means that the chassis automatically brings down the second interface in the inline interface pair when one of the interfaces in an inline set goes down. New/Modified FXOS commands: **show fault |grep link-down**, **show interface detail** Supported platforms: Firepower 4100/9300
Support for intra-chassis clustering on the Firepower Threat Defense on the Firepower 9300	6.0.1	The Firepower 9300 supports intra-chassis clustering with the Firepower Threat Defense application. New/Modified Firepower Chassis Manager screens: • **Logical Devices** > **Configuration** New/Modified FXOS commands: **enter mgmt-bootstrap ftd, enter bootstrap-key FIREPOWER_MANAGER_IP, enter bootstrap-key FIREWALL_MODE, enter bootstrap-key-secret REGISTRATION_KEY, enter bootstrap-key-secret PASSWORD, enter bootstrap-key FQDN, enter bootstrap-key DNS_SERVERS, enter bootstrap-key SEARCH_DOMAINS, enter ipv4 firepower, enter ipv6 firepower, set value, set gateway, set ip, accept-license-agreement** Supported platforms: Firepower 4100/9300



CHAPTER 28

FlexConfig Policies for FTD

The following topics describe how to configure and deploy FlexConfig policies.

- FlexConfig Policy Overview, on page 575
- Requirements and Prerequisites for FlexConfig Policies, on page 591
- Guidelines and Limitations for FlexConfig, on page 592
- Customizing Device Configuration with FlexConfig Policies, on page 592
- History for FlexConfig, on page 605

FlexConfig Policy Overview

A FlexConfig policy is a container of an ordered list of FlexConfig objects. Each object includes a series of Apache Velocity scripting language commands, ASA software configuration commands, and variables that you define. The contents of each FlexConfig object is essentially a program that generates a sequence of ASA commands that will then be deployed to the assigned devices. This command sequence then configures the related feature on the Firepower Threat Defense device.

Firepower Threat Defense uses ASA configuration commands to implement some features, but not all features. There is no unique set of Firepower Threat Defense configuration commands. Instead, the point of FlexConfig is to allow you to configure features that are not yet directly supported through Firepower Management Center policies and settings.



Caution Cisco **strongly** recommends using FlexConfig policies only if you are an advanced user with a strong ASA background and at your own risk. You may configure any commands that are not prohibited. Enabling features through FlexConfig policies may cause unintended results with other configured features.

You may contact the Cisco Technical Assistance Center for support concerning FlexConfig policies that you have configured. The Cisco Technical Assistance Center does not design or write custom configurations on any customer's behalf. Cisco expresses no guarantees for correct operation or interoperability with other Firepower System features. FlexConfig features may become deprecated at any time. For fully guaranteed feature support, you must wait for Firepower Management Center support. When in doubt, do not use FlexConfig policies.

Recommended Usage for FlexConfig Policies

There are two main recommended uses for FlexConfig:

- You are converting from ASA to Firepower Threat Defense, and there are compatible features you are using (and need to continue using) that Firepower Management Center does not directly support. In this case, use the **show running-config** command on the ASA to see the configuration for the feature and create your FlexConfig objects to implement it. Experiment with the object's deployment settings (once/everytime and append/prepend) to get the right setting. Verify by comparing **show running-config** output on the two devices.
- You are using Firepower Threat Defense but there is a setting or feature that you need to configure, e.g. the Cisco Technical Assistance Center tells you that a particular setting should resolve a specific problem you are encountering. For complicated features, use a lab device to test the FlexConfig and verify that you are getting the expected behavior.

The system includes a set of predefined FlexConfig objects that represent tested configurations. If the feature you need is not represented by these objects, first determine if you can configure an equivalent feature in standard policies. For example, the access control policy includes intrusion detection and prevention, HTTP and other types of protocol inspection, URL filtering, application filtering, and access control, which the ASA implements using separate features. Because many features are not configured using CLI commands, you will not see every policy represented within the output of **show running-config**.



Note At all times, keep in mind that there is not a one-to-one overlap between ASA and Firepower Threat Defense. Do not attempt to completely recreate an ASA configuration on a Firepower Threat Defense device. You must carefully test any feature that you configure using FlexConfig.

CLI Commands in FlexConfig Objects

Firepower Threat Defense uses ASA configuration commands to configure some features. Although not all ASA features are compatible with Firepower Threat Defense, there are some features that can work on Firepower Threat Defense but that you cannot configure in Firepower Management Center policies. You can use FlexConfig objects to specify the CLI required to configure these features.

If you decide to use FlexConfig to manually configure a feature, you are responsible for knowing and implementing the commands according to the proper syntax. FlexConfig policies do not validate CLI command syntax. For more information about proper syntax and configuring CLI commands, use the ASA documentation as a reference:

- ASA CLI configuration guides explain how to configure a feature. Find the guides at http://www.cisco.com/c/en/us/support/security/asa-5500-series-next-generation-firewalls/products-installation-and-configuration-guides-list.html
- ASA command references provide additional information sorted by command name. Find the references at http://www.cisco.com/c/en/us/support/security/asa-5500-series-next-generation-firewalls/products-command-reference-list.html

The following topics explain more about configuration commands.

Determine the ASA Software Version and Current CLI Configuration

Because the system uses ASA software commands to configure some features, you need to determine the current ASA version used in software running on the Firepower Threat Defense device. This version number indicates which ASA CLI configuration guides to use for instructions on configuring a feature. You also

should examine the current CLI-based configuration and compare it to the ASA configuration you want to implement.

Keep in mind that any ASA configuration will be very different from a Firepower Threat Defense configuration. Many Firepower Threat Defense policies are configured outside of the CLI, so you cannot see the configuration by looking at the commands. Do not try to create a one-to-one correspondence between an ASA and Firepower Threat Defense configuration.

To view this information, make an SSH connection to the device's management interface and issue the following commands:

- **show version system** and look for the Cisco Adaptive Security Appliance Software Version number. (If you issue the command through the Firepower Management Center CLI tool, omit the **system** keyword.)
- **show running-config** to view the current CLI configuration.
- **show running-config all** to include all the default commands in the current CLI configuration.

You can also issue these commands from within Firepower Management Center using the following procedure.

Procedure

Step 1 Choose **System** > **Health** > **Monitor**.

Step 2 Click the name of the device targeted by the FlexConfig policy.

You might need to click the open/close arrow in the **Count** column in the Status table to see any devices.

Step 3 Choose **Advanced Troubleshooting**.

Step 4 Choose **Threat Defense CLI**.

Step 5 Choose **show** as the command, and type **version** or one of the other commands as the parameter.

Step 6 Click **Execute**.

For version, search the output for the Cisco Adaptive Security Appliance Software Version number.

You can select the output and press Ctrl+C, then paste it into a text file for later analysis.

Prohibited CLI Commands

The purpose of FlexConfig is to configure features that are available on ASA devices that you cannot configure on Firepower Threat Defense devices using Firepower Management Center.

Thus, you are prevented from configuring ASA features that have equivalents in Firepower Management Center. The following table lists some of these prohibited command areas.

In addition, some **clear** commands are prohibited because they overlap with managed policies, and can delete part of the configuration for a managed policy.

The FlexConfig object editor prevents you from including prohibited commands in the object.

Prohibited CLI Command	Description
AAA	Configuration blocked.

Prohibited CLI Command	Description
AAA-Server	Configuration blocked.
Access-list	Advanced ACL, Extended ACL, and Standard ACL are blocked. Ethertype ACL is allowed. You can use standard and extended ACL objects defined in the object manager inside the template as variables.
ARP Inspection	Configuration blocked.
As-path Object	Configuration blocked.
Banner	Configuration blocked.
BGP	Configuration blocked.
Clock	Configuration blocked.
Community-list Object	Configuration blocked.
Copy	Configuration blocked.
Delete	Configuration blocked.
DHCP	Configuration blocked.
Enable Password	Configuration blocked.
Erase	Configuration blocked.
Fragment Setting	Blocked, except for **fragment reassembly**.
Fsck	Configuration blocked.
HTTP	Configuration blocked.
ICMP	Configuration blocked.
Interface	Only **nameif**, **mode**, **shutdown**, **ip address** and **mac-address** commands are blocked.
Multicast Routing	Configuration blocked.
NAT	Configuration blocked.
Network Object/Object-group	Network object creation in the FlexConfig object is blocked, but you can use network objects and groups defined in the object manager inside the template as variables.
NTP	Configuration blocked.
OSPF/OSPFv3	Configuration blocked.
Password Encryption	Configuration blocked.

Prohibited CLI Command	Description
Policy-list Object	Configuration blocked.
Prefix-list Object	Configuration blocked.
Reload	You cannot schedule reloads. The system does not use the **reload** command to restart the system, it uses the **reboot** command.
RIP	Configuration blocked.
Route-Map Object	Route-map object creation in the FlexConfig object is blocked, but you can use route map objects defined in the object manager inside the template as variables.
Service Object/Object-group	Service object creation in the FlexConfig object is blocked, but you can use port objects defined in the object manager inside the template as variables.
SNMP	Configuration blocked.
SSH	Configuration blocked.
Static Route	Configuration blocked.
Syslog	Configuration blocked.
Time Synchronization	Configuration blocked.
Timeout	Configuration blocked.
VPN	Configuration blocked.

Template Scripts

You can use scripting language to control processing within a FlexConfig object. Scripting language instructions are a subset of commands supported in the Apache Velocity 1.3.1 template engine, a Java-based scripting language that supports looping, if/else statements, and variables.

To learn how to use the scripting language, see the *Velocity Developer Guide* at http://velocity.apache.org/engine/devel/developer-guide.html.

FlexConfig Variables

You can use variables in a FlexConfig object in cases where part of a command or processing instruction depends on runtime information rather than static information. During deployment, the variables are replaced with strings obtained from other configurations for the device based on the type of variable:

- Policy object variables are replaced with strings obtained from objects defined in Firepower Management Center.
- System variables are replaced with information obtained from the device itself or from policies configured for it.

- Processing variables are loaded with the contents of policy object or system variables as scripting commands are processed. For example, in a loop, you iteratively load one value from a policy object or system variable into a processing variable, then use the processing variable to form a command string or perform some other action. These processing variables do not show up in the Variables list within a FlexConfig object. Also, you do not add them using the **Insert** menu in the FlexConfig object editor.

Variables start with the $ character. For example, $ifname is a policy object variable in the following command

```
interface $ifname
```



Note The first time you insert a policy object or system variable, you must do so through the **Insert** menu in the FlexConfig object editor. This action adds the variable to the **Variables** list at the bottom of the FlexConfig object editor. But you must type in the variable string on subsequent uses, even when using system variables. If you are adding a processing variable, which does not have an object or system variable assignment, do not use the **Insert** menu.

Whether a variable is resolved as a single string, a list of strings, or a table of values depends on the type of policy object or system variable you assign to the variable. You must understand what will be returned in order to process the variables correctly.

The following topics explain the various types of variable and how to process them.

How to Process Variables

At runtime, a variable can resolve to a single string, a list of strings of the same type, a list of strings of different types, or a table of named values. In addition, variables that resolve to multiple values can be of determinate or indeterminate length. You must understand what will be returned in order to process the values correctly.

Following are the main possibilities.

Single Value Variables

If a variable always resolves to a single string, use the variable directly without modification in the FlexConfig script.

For example, the predefined text variable tcpMssBytes always resolves to a single value (which must be numeric). The **Sysopt_basic** FlexConfig then uses an if/then/else structure to set the maximum segment size based on the value of another single-value text variable, tcpMssMinimum:

```
#if($tcpMssMinimum == "true")
  sysopt connection tcpmss minimum $tcpMssBytes
#else
  sysopt connection tcpmss $tcpMssBytes
#end
```

In this example, you would use the **Insert** menu in the FlexConfig object editor to add the first use of $tcpMssBytes, but you would type in the variable directly on the #else line.

Note Policy object variables for network objects also equate to a single IP address specification, either a host address, network address, or address range. However, in this case, you must know what type of address to expect, because the ASA commands require specific address types. For example, if a command requires a host address, using a network object variable that points to an object that contains a network address will result in an error during deployment.

Multiple Value Variables, All Values Are the Same Type

Several policy object and system variables resolve to multiple values of the same type. For example, an object variable that points to a network object group resolves to a list of the IP addresses within the group. Similarly, the system variable $SYS_FW_INTERFACE_NAME_LIST resolves to a list of interface names.

You can also create text objects for multiple values of the same type. For example, the predefined text object enableInspectProtocolList can contain more than one protocol name.

Multiple value variables that resolve to a list of items of the same type are frequently of indeterminate length. For example, you cannot know beforehand how many interfaces on a device are named, as users can configure or unconfigure interfaces at any time.

Thus, you would typically use a loop to process multiple value variables of the same type. For example, the predefined FlexConfig **Default_Inspection_Protocol_Enable** uses a #foreach loop to go through the enableInspectProtocolList object and process each value.

```
policy-map global_policy
    class inspection_default
    #foreach ( $protocol in $enableInspectProtocolList)
    inspect $protocol
    #end
```

In this example, the script assigns each value in turn to the $protocol variable, which is then used in an ASA **inspect** command to enable the inspection engine for that protocol. In this case, you simply type in $protocol as a variable name. You do not use the **Insert** menu to add it, because you are not assigning an object or system value to the variable. However, you must use the **Insert** menu to add $enableInspectProtocolList.

The system loops through the code between #foreach and #end until there are no values remaining in $enableInspectProtocolList.

Multiple Value Variables, Values Are Different Types

You can create multiple value text objects, but have each value serve a different purpose. For example, the predefined **netflow_Destination** text object should have 3 values, in order, interface name, destination IP address, and UDP port number.

Objects defined in this way should have a determinate number of values. Otherwise, they would be hard to process.

Use the get method to process these objects. Type **.get(*n*)** at the end of the object name, replacing *n* with an index into the object. Start counting at 0, even though the text object lists its values starting at 1.

For example, the Netflow_Add_Destination object uses the following line to add the 3 values from netflow_Destination to the ASA **flow-export** command.

```
flow-export destination $netflow_Destination.get(0) $netflow_Destination.get(1)
```

```
$netflow_Destination.get(2)
```

In this example, you would use the **Insert** menu in the FlexConfig object editor to add the first use of $netflow_Destination, and then add **.get(0)**. But you would type in the variable directly for the **$netflow_Destination.get(1)** and **$netflow_Destination.get(2)** specifications.

Multiple Value Variables that Resolve to a Table of Values

Some system variables return a table of values. These variables include MAP in their name, for example, $SYS_FTD_ROUTED_INTF_MAP_LIST. The routed interface map returns data that looks like the following (line returns added for clarity):

```
[{intf_hardwarare_id=GigabitEthernet0/0, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=255.255.255.0,
intf_ip_addr_v4=10.100.10.1, intf_ipv6_link_local_address=,
intf_logical_name=outside},

{intf_hardwarare_id=GigabitEthernet0/1, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=255.255.255.0,
intf_ip_addr_v4=10.100.11.1, intf_ipv6_link_local_address=,
intf_logical_name=inside},

{intf_hardwarare_id=GigabitEthernet0/2, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=, intf_ip_addr_v4=,
intf_ipv6_link_local_address=, intf_logical_name=},

{intf_hardwarare_id=Management0/0, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=, intf_ip_addr_v4=,
intf_ipv6_link_local_address=, intf_logical_name=diagnostic}]
```

In the above example, information is returned for 4 interfaces. Each interface includes a table of named values. For example, intf_hardwarare_id is the name of the interface hardware name property, and returns strings such as GigabitEthernet0/0.

This type of variable is typically of indeterminate length, so you need to use looping to process the values. But you also need to add the property name to the variable name to indicate which value to retrieve.

For example, IS-IS configuration requires that you add the ASA **isis** command to an interface that has a logical name in interface configuration mode. However, you enter that mode using the interface's hardware name. Thus, you need to identify which interfaces have logical names, then configure just those interfaces using their hardware names. The ISIS_Interface_Configuration predefined FlexConfig does this using an if/then structure nested in a loop. In the following code, you can see that the #foreach scripting command loads each interface map into the $intf variable, then the #if statement keys off the intf_logical_name value in the map ($intf.intf_logical_name), and if the value is in the list defined in the isisIntfList predefined text variable, enters the interface command using the intf_hardwarare_id value ($intf.intf_hardwarare_id). You would need to edit the isisIntfList variable to add the names of the interfaces on which to configure IS-IS.

```
#foreach ($intf in $SYS_FTD_ROUTED_INTF_MAP_LIST)
 #if ($isIsIntfList.contains($intf.intf_logical_name))
  interface $intf.intf_hardwarare_id
   isis
   #if ($isIsAddressFamily.contains("ipv6"))
   ipv6 router isis
   #end
 #end
#end
```

How to See What a Variable Will Return for a Device

An easy way to evaluate what a variable will return is to create a simple FlexConfig object that does nothing more than process an annotated list of variables. Then, you can assign it to a FlexConfig policy, assign the policy to a device, save the policy, then preview the configuration for that device. The preview will show the resolved values. You can select the preview text, press Ctrl+C, then paste the output into a text file for analysis.



Note Do not deploy this FlexConfig to the device, however, because it will not contain any valid configuration commands. You would get deployment errors. After obtaining the preview, delete the FlexConfig object from the FlexConfig policy and save the policy.

For example, you could construct the following FlexConfig object:

```
Following is a network object group variable for the
IPv4-Private-All-RFC1918 object:

$IPv4_Private_addresses

Following is the system variable SYS_FW_MANAGEMENT_IP:

$SYS_FW_MANAGEMENT_IP

Following is the system variable SYS_FW_ENABLED_INSPECT_PROTOCOL_LIST:

$SYS_FW_ENABLED_INSPECT_PROTOCOL_LIST

Following is the system variable SYS_FTD_ROUTED_INTF_MAP_LIST:

$SYS_FTD_ROUTED_INTF_MAP_LIST

Following is the system variable SYS_FW_INTERFACE_NAME_LIST:

$SYS_FW_INTERFACE_NAME_LIST
```

The preview of this object might look like the following (line returns added for clarity):

```
###Flex-config Prepended CLI ###

###CLI generated from managed features ###

###Flex-config Appended CLI ###
Following is an network object group variable for the
IPv4-Private-All-RFC1918 object:

[10.0.0.0, 172.16.0.0, 192.168.0.0]

Following is the system variable SYS_FW_MANAGEMENT_IP:

192.168.0.171

Following is the system variable SYS_FW_ENABLED_INSPECT_PROTOCOL_LIST:

[dns, ftp, h323 h225, h323 ras, rsh, rtsp, sqlnet, skinny, sunrpc,
xdmcp, sip, netbios, tftp, icmp, icmp error, ip-options]

Following is the system variable SYS_FTD_ROUTED_INTF_MAP_LIST:
```

```
[{intf_hardwarare_id=GigabitEthernet0/0, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=255.255.255.0,
intf_ip_addr_v4=10.100.10.1, intf_ipv6_link_local_address=,
intf_logical_name=outside},

{intf_hardwarare_id=GigabitEthernet0/1, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=255.255.255.0,
intf_ip_addr_v4=10.100.11.1, intf_ipv6_link_local_address=,
intf_logical_name=inside},

{intf_hardwarare_id=GigabitEthernet0/2, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=, intf_ip_addr_v4=,
intf_ipv6_link_local_address=, intf_logical_name=},

{intf_hardwarare_id=Management0/0, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=, intf_ip_addr_v4=,
intf_ipv6_link_local_address=, intf_logical_name=diagnostic}]

Following is the system variable SYS_FW_INTERFACE_NAME_LIST:

[outside, inside, diagnostic]
```

FlexConfig Policy Object Variables

A policy object variable is associated with a specific policy object configured in the Object Manager. When you insert a policy object variable in a FlexConfig object, you give the variable a name and select the object associated with it.

Although you can give the variable the exact same name as the associated object, the variable itself is not the same thing as the associated object. You must use the **Insert** > **Insert Policy Object** > ***Object Type*** menu in the FlexConfig object editor to add the variable for the first time to the script in the FlexConfig to establish the association with the object. Simply typing in the name of the object preceded by a $ sign does not create a policy object variable.

You can create variables to point to the following types of object. Ensure that you create the right type of object for each variable. To create objects, go to the **Objects** > **Object Management** page.

- **Text Objects**—For text strings, which can include IP addresses, numbers, and other free-form text such as interface or zone names. Select **FlexConfig** > **Text Object** from the table of contents, then click **Add Text Object**. You can configure these objects to contain a single value or multiple values. These objects are highly flexible and built specifically for use within FlexConfig objects. For detailed information, see Configure FlexConfig Text Objects, on page 597.
- **Network**—For IP addresses. You can use network objects or groups. Select **Network** from the table of contents, then select **Add Network** > **Add Object** or **Add Group**. If you use a group object, the variable returns a list of each IP address specification within the group. Addresses can be host, network, or address ranges, depending on the object contents. See Network Objects, on page 402.
- **Security Zones**—For interfaces within a security zone or interface group. Select **Interface** from the table of contents, then select **Add** > **Security Zone** or **Interface Group**. A security zone variable returns a list of the interfaces within that zone or group for the device being configured. See Interface Objects: Interface Groups and Security Zones, on page 410.
- **Standard ACL Object**—For standard access control lists. A standard ACL variable returns the name of the standard ACL object. Select **Access List** > **Standard** from the table of contents, then click **Add Standard Access List Object**. See Access List, on page 467.

- **Extended ACL Object**—For extended access control lists. An extended ACL variable returns the name of the extended ACL object. Select **Access List** > **Extended** from the table of contents, then click **Add Extended Access List Object**. See Access List, on page 467.
- **Route Map**—For route map objects. A route map variable returns the name of the route map object. Select **Route Map** from the table of contents, then click **Add Route Map**. See Route Maps, on page 464.

FlexConfig System Variables

System variables are replaced with information obtained from the device itself or from policies configured for it.

You must use the **Insert** > **Insert System Variable** > ***Variable Name*** menu in the FlexConfig object editor to add the variable for the first time to the script in the FlexConfig to establish the association with the system variable. Simply typing in the name of the system variable preceded by a $ sign does not create a system variable within the context of the FlexConfig object.

The following table explains the available system variables. Before using a variable, examine what is typically returned for the variable; see How to See What a Variable Will Return for a Device, on page 583.

Name	Description
SYS_FW_OS_MODE	The operating system mode of the device. Possible values are ROUTED or TRANSPARENT.
SYS_FW_OS_MULTIPLICITY	Whether the device is running in single or multiple context mode. Possible values are SINGLE, MULTI, or NOT_APPLICABLE.
SYS_FW_MANAGEMENT_IP	The management IP address of the device
SYS_FW_HOST_NAME	The device hostname
SYS_FTD_INTF_POLICY_MAP	A map with interface name as key and policy-map as value. This variable returns nothing if there are no interface-based service policies defined on the device.
SYS_FW_ENABLED_INSPECT_PROTOCOL_LIST	The list of protocols for which inspection is enabled.
SYS_FTD_ROUTED_INTF_MAP_LIST	A list of routed interface maps on the device. Each map includes a set of named values related to routed interface configuration.
SYS_FTD_SWITCHED_INTF_MAP_LIST	A list of switched interface maps on the device. Each map includes a set of named values related to switched interface configuration.
SYS_FTD_INLINE_INTF_MAP_LIST	A list of inline interface maps on the device. Each map includes a set of named values related to inline set interface configuration.
SYS_FTD_PASSIVE_INTF_MAP_LIST	A list of passive interface maps on the device. Each map includes a set of named values related to passive interface configuration.
SYS_FTD_INTF_BVI_MAP_LIST	A list of Bridge Virtual Interface maps on the device. Each map includes a set of named values related to BVI configuration.

Name	Description
SYS_FW_INTERFACE_HARDWARE_ID_LIST	A list of the hardware names for interfaces on the device, such as GigabitEthernet0/0.
SYS_FW_INTERFACE_NAME_LIST	A list of logical names for interfaces on the device, such as inside.
SYS_FW_INLINE_INTERFACE_NAME_LIST	A list of logical names for interfaces configured as passive or ERSPAN passive.
SYS_FW_NON_INLINE_INTERFACE_NAME_LIST	A list of logical names for interfaces that are not part of inline sets, such as all routed interfaces.

Predefined FlexConfig Objects

The predefined FlexConfig objects provide tested configurations for select features. Use these objects if you need to configure these features, which otherwise cannot be configured through Firepower Management Center.

The following table lists the available objects. Make note of the associated text objects. You must edit these text objects to customize the behavior of the predefined FlexConfig object. The text objects make it possible for you to customize the configuration using the IP addresses and other attributes required by your network and device.

If you need to modify a predefined FlexConfig object, copy the object, make changes to the copy, and save it with a new name. You cannot directly edit a predefined FlexConfig object.

Although you might be able to configure other ASA-based features using FlexConfig, the configuration of those features has not been tested. If an ASA feature overlaps with something that you can configure in Firepower Management Center policies, do not attempt to configure it through FlexConfig.

For example, Snort inspection includes the HTTP protocol, so do not enable ASA-style HTTP inspection. (In fact, you cannot add **http** to the enableInspectProtocolList object. In this case, you are prevented from misconfiguring your device.) Instead, configure the access control policy to perform application or URL filtering, as needed, to implement your HTTP inspection requirements.

FlexConfig Object Name	Description	Associated Text Objects
Default_Inspection_Protocol_Disable	Disables protocols in the global_policy default policy map.	disableInspectProtocolList
Default_Inspection_Protocol_Enable	Enables protocols in the global_policy default policy map.	enableInspectProtocolList
Eigrp_Configure	Configures EIGRP routing next-hop, auto-summary, router-id, eigrp-stub.	eigrpAS, eigrpNetworks, eigrpDisableAutoSummary, eigrpRouterId, eigrpStubReceiveOnly, eigrpStubRedistributed, eigrpStubConnected, eigrpStubStatic, eigrpStubSummary

FlexConfig Object Name	Description	Associated Text Objects
Eigrp_Interface_Configure	Configures EIGRP interface authentication mode, authentication key, hello interval, hold time, split horizon.	eigrpIntfList, eigrpAS, eigrpAuthKey, eigrpAuthKeyId, eigrpHelloInterval, eigrpHoldTime, eigrpDisableSplitHorizon Also uses the system variable SYS_FTD_ROUTED_INTF_MAP_LIST
Eigrp_Unconfigure	Clears EIGRP configuration for an autonomous system from the device.	—
Eigrp_Unconfigure_all	Clears all EIGRP configurations.	—
ISIS_Configure	Configures global parameters for IS-IS routing.	isIsNet, isIsAddressFamily, isISType
ISIS_Interface_Configuration	Interface level IS-IS configuration.	isIsAddressFamily, IsIsIntfList Also uses the system variable SYS_FTD_ROUTED_INTF_MAP_LIST
ISIS_Unconfigure	Clears the IS-IS router configuration on the device.	—
ISIS_Unconfigure_All	Clears the IS-IS router configuration from the device, including the router assignment from the device interface.	—
Netflow_Add_Destination	Creates and configures a Netflow export destination.	Netflow_Destinations, netflow_Event_Types
Netflow_Clear_Parameters	Restores Netflow export global default settings.	—
Netflow_Delete_Destination	Deletes a Netflow export destination.	Netflow_Destinations, netflow_Event_Types
Netflow_Set_Parameters	Sets global parameters for Netflow export.	netflow_Parameters
NGFW_TCP_NORMALIZATION	Modifies the default TCP normalization configuration.	—
Policy_Based_Routing	To use this example configuration, copy it, modify the interface name, and use the r-map-object text object to identify a route map object in the object manager.	—
Policy_Based_Routing_Clear	Clears Policy Based Routing configurations from the device.	—
Sysopt_AAA_radius	Ignores the authentication key in RADIUS accounting responses.	—

FlexConfig Object Name	Description	Associated Text Objects
Sysopt_AAA_radius_negate	Negates the Sysopt_AAA_radius configuration.	—
Sysopt_basic	Configures sysopt wait time , maximum segment size for TCP packets, and detailed traffic statistics.	tcpMssMinimum, tcpMssBytes
Sysopt_basic_negate	Clears sysopt_basic detailed traffic statistics, wait time, and TCP maximum segment size.	—
Sysopt_clear_all	Clears all sysopt configurations from the device.	—
Sysopt_noproxyarp	Configures noproxy-arp CLIs.	Uses system variable SYS_FW_NON_INLINE_INTF_NAME_LIST
Sysopt_noproxyarp_negate	Clears Sysopt_noproxyarp configurations.	Uses system variable SYS_FW_NON_INLINE_INTF_NAME_LIST
Sysopt_Preserve_Vpn_Flow	Configures syopt preserve VPN flow.	—
Sysopt_Preserve_Vpn_Flow_negate	Clears the Sysopt_Preserve_Vpn_Flow configuration.	—
Sysopt_Reclassify_Vpn	Configures sysopt reclassify vpn.	—
Sysopt_Reclassify_Vpn_Negate	Negates sysopt reclassify vpn.	—
VxLAN_Clear_Nve	Removes the NVE 1 configured when VxLAN_Configure_Port_And_Nve is used from the device.	—
VxLAN_Clear_Nve_Only	Clears the NVE configured on the interface when deployed.	—
VxLAN_Configure_Port_And_Nve	Configures VLAN port and NVE 1.	vxlan_Port_And_Nve
VxLAN_Make_Nve_Only	Sets an interface for NVE only.	vxlan_Nve_Only Also uses system variables SYS_FTD_ROUTED_MAP_LIST and SYS_FTD_SWITCHED_INTF_MAP_LIST
VxLAN_Make_Vni	Creates a VNI interface. After deploying this you have to unregister and re-register the device to properly discover the VNI interface.	vxlan_Vni
Wccp_Configure	This template provides an example for configuring WCCP.	isServiceIdentifier, serviceIdentifier, wccpPassword
Wccp_Configure_Clear	Clears WCCP configurations.	—

Predefined Text Objects

There are several predefined text objects. These objects are associated with variables used in the predefined FlexConfig objects. In most cases, you must edit these objects to add values if you use the associated FlexConfig object, or you will see errors during deployment. Although some of these objects contain default values, others are empty.

For information on editing text objects, see Configure FlexConfig Text Objects, on page 597.

Name	Description	Associated FlexConfig Object
disableInspectProtocolList	Disables protocols in the default policy map (global_policy).	Disable_Default_Inspection_Protocol
eigrpAS	Autonomous system number.	Eigrp_Configure, Eigrp_Interface_Configure, Eigrp_Unconfigure
eigrpAuthKey	EIGRP authentication key.	Eigrp_Interface_Configure
eigrpAuthKeyId	Shared key id that matches the authentication key.	Eigrp_Interface_Configure
eigrpDisableAutoSummary	A flag that, when true, disables auto-summary.	Eigrp_Configure
eigrpDisableSplitHorizon	A flag that, when true, disables split horizon.	Eigrp_Interface_Configure
eigrpHelloInterval	Seconds between hello transmission.	Eigrp_Interface_Configure
eigrpHoldTime	Seconds before neighbor is considered down.	Eigrp_Interface_Configure
eigrpIntfList	List of logical interface names where EIGRP is to be applied.	Eigrp_Interface_Configure
eigrpRouterId	Router-Id, in IP address format.	Eigrp_Configure
eigrpStubConnected	A flag that, when true, allows you to use **connected** in the **eigrp stub** configuration.	Eigrp_Configure
eigrpStubReceiveOnly	A flag that, when true, allows you to use **receive-only** in the **eigrp stub** configuration.	Eigrp_Configure
eigrpStubRedistributed	A flag that, when true, allows you to use **redistributed** in the **eigrp stub** configuration.	Eigrp_Configure
eigrpStubSummary	A flag that, when true, allows you to use **summary** in the **eigrp stub** configuration.	Eigrp_Configure

Name	Description	Associated FlexConfig Object
enableInspectProtocolList	Enables protocols in the default policy map (global_policy). You are prevented from adding protocols whose inspection conflicts with Snort inspection.	Enable_Default_Inspection_Protocol
isIsAddressFamily	The IPv4 or IPv6 address family.	ISIS_Configure ISIS_Interface_Configuration
IsIsIntfList	List of logical interface names.	ISIS_Interface_Configuration
isIsISType	IS Type (level-1, level-2-only or level-1-2).	ISIS_Configure
isIsNet	Network entity.	ISIS_Configure
isServiceIdentifier	When false, uses the standard **web-cache** service identifier.	Wccp_Configure
netflow_Destination	Defines a single Netflow export destination's interface, destination, and UDP port number.	Netflow_Add_Destination
netflow_Event_Types	Defines the types of events to be exported for a destination as any subset of: **all**, **flow-create**, **flow-defined**, **flow-teardown**, **flow-update**.	Netflow_Add_Destination
netflow_Parameters	Provides the Netflow export global settings: active refresh interval (number of minutes between flow update events), delay (flow create delay in seconds; default 0 = command will not appear), and template time-out rate in minutes.	Netflow_Set_Parameters
serviceIdentifier	Dynamic WCCP service identifier number.	Wccp_Configure
tcpMssBytes	Maximum segment size in bytes.	Sysopt_basic
tcpMssMinimum	Checks whether to set maximum segment size (MSS), which is set only if this flag is true.	Sysopt_basic
vxlan_Nve_Only	Parameters for configuring NVE-only on interface: • logical name of interface • IPv4 address (optional for routed interface) • IPv4 netmask (optional for routed interface)	VxLAN_Make_Nve_Only

Name	Description	Associated FlexConfig Object
vxlan_Port_And_Nve	Parameters used for configuring ports and NVE for VXLAN: • vxlan port • source interface (logical name) • type (peer or mcast) • Peer IP Address or default-mcast-grooup	VxLAN_Configure_Port_And_Nve
vxlan_Vni	Parameters used for creating VNI: • Interface number (1-10000) • segment-id (1-16777215) • nameif (Logical Name of the interface) • type (routed or transparent) • IP address (used in case of routed mode device) or bridge-group number (used in case of transparent mode device) • netmask (If device is in routed mode) or unused	VxLAN_Make_Vni
wccpPassword	WCCP password.	Wccp_Configure

Requirements and Prerequisites for FlexConfig Policies

Model Support

FTD

Supported Domains

Any

User Roles

Admin

Guidelines and Limitations for FlexConfig

- If you make a mistake in the FlexConfig policy, the system will roll back all changes included in the deployment attempt that includes the failed FlexConfig. Because rollback due to a failed deployment includes clearing the configuration, this can be disruptive to your network. Consider timing deployments that include FlexConfig changes to non-business hours. Also, consider isolation the deployment so it includes just FlexConfig changes, and no other policy updates.
- When you use the VxLAN_Make_VNI object, you must deploy the same FlexConfig to all units in a cluster or high availability pair before you form the cluster or high availability pair. The Management Center requires the VXLAN interfaces to match on all devices before forming the cluster or high availability pair.
- If you want to configure Equal-Cost-Multi-Path (ECMP) routing using traffic zones, the **zone** command differs for Firepower Threat Defense devices compared to the one used on ASA. Although you can still follow the instructions in the ASA general configuration guide, use **zone** *name* **ecmp** instead of the ASA version of the command. Otherwise, the operation of the traffic zone feature is identical between the ASA and Firepower Threat Defense .

Note The system also configures **zone** *name* **passive** commands to configure passive zones if you define some interfaces as passive. This is handled automatically based on your interface configuration. Do not use FlexConfig to create passive traffic zones.

Customizing Device Configuration with FlexConfig Policies

Use FlexConfig policies to customize the configuration of a Firepower Threat Defense device.

Before using FlexConfig, try to configure all the policies and settings you need using the other features in Firepower Management Center. FlexConfig is a method of last resort to configure ASA-based features that are compatible with Firepower Threat Defense but which are not otherwise configurable in Firepower Management Center.

Following is the end-to-end procedure for configuring and deploying a FlexConfig policy.

Procedure

Step 1 Determine the CLI command sequence that you want to configure.

If you have a functioning configuration on an ASA device, use **show running-config** to get the sequence of commands that you need. Make adjustments to items such as interface names and IP addresses as needed.

If this is for a new feature, it is best to try to implement it on an ASA device in a lab setting to verify that you have the correct command sequence.

For more information, see the following topics:

- Recommended Usage for FlexConfig Policies, on page 575

- CLI Commands in FlexConfig Objects, on page 576

Step 2 Select **Objects** > **Object Management**, then select **FlexConfig** > **FlexConfig Objects** from the table of contents.

Examine the predefined FlexConfig objects to determine if any will be able to generate the commands you need. Click **View** ( ) to see the object contents. If an existing object is close to what you want, start by making a copy of the object, and then edit the copy. See Predefined FlexConfig Objects, on page 586.

Examining the objects will also give you an idea of the structure, command syntax, and expected sequencing for a FlexConfig object.

Note If you find any objects that you will use, either directly or as copies, examine the Variables list at the bottom of the object. Make note of the variable names, except those in all capitals that start with SYS, which are system variables. These variables are text objects that you will probably need to edit and define overrides for, especially if the default value column shows the object has no value.

Step 3 If you need to create your own FlexConfig objects, determine what variables you will need and create the associated objects.

The CLI you need to deploy might contain IP addresses, interface names, port numbers, and other parameters that you might want to adjust over time. These are best isolated into variables, which point to objects that contain the necessary values. You might also need variables for strings that are part of the configuration but which might change over time.

Also, determine if you need different values for each device to which you will assign the policy. For example, you might want to configure the feature on three devices, but you might need to specify a different interface name or IP address on a given command for each of these devices. If you need to customize the object for each device, ensure that you enable overrides when creating the object, and then define the override values per device.

See the following topics for an explanation of the various types of variables and how to configure the related objects when necessary.

- FlexConfig Variables, on page 579
- FlexConfig Policy Object Variables, on page 584
- FlexConfig System Variables, on page 585
- Configure FlexConfig Text Objects, on page 597

Step 4 If you are using the predefined FlexConfig objects, edit the text objects used as variables.

See Configure FlexConfig Text Objects, on page 597.

Step 5 (If necessary.) Configure FlexConfig Objects, on page 594.

You need to create objects only if the predefined objects cannot do the job.

Step 6 Configure the FlexConfig Policy, on page 599.

Step 7 Set Target Devices for a FlexConfig Policy, on page 600.

You can also assign the policy to devices when you create the policy. The policy must have at least one assigned device before you can preview it.

Step 8 Preview the FlexConfig Policy, on page 600.

You must save changes before you can preview the policy.

Verify that the generated commands are the ones intended, and that all variables are resolving correctly.

Step 9 Click **Deploy** in the menu bar, select the devices assigned to the policy, and click the **Deploy** button.

Wait for deployment to complete.

Step 10 Choose **Deploy** > **Deployment** in the menu bar.

Step 11 Verify the Deployed Configuration, on page 601.

Step 12 (If necessary.) Remove Features Configured Using FlexConfig, on page 603.

Unlike other types of policy, simply unassigning a FlexConfig from a device might not remove the related configuration. If you want to remove a FlexConfig-generated configuration, you follow the cited procedure.

If you are removing a Feature because it is now directly supported by the product, see also Convert from FlexConfig to Managed Feature, on page 604.

Configure FlexConfig Objects

Use FlexConfig objects to define a configuration to be deployed to a device. Each FlexConfig policy is composed of a list of FlexConfig objects, so the objects are essentially code modules composed of Apache Velocity scripting commands, ASA software configuration commands, and variables.

There are several predefined FlexConfig objects that you can use directly, or you can make copies if you need to edit them. You can also create your own objects from scratch. A FlexConfig object's content can range from a single simple command string to elaborate CLI command structures that use variables and scripting commands to deploy commands whose content can differ from device to device or deployment to deployment.

You can also create FlexConfig policy objects when defining FlexConfig policies.

Before you begin

Keep the following in mind:

- FlexConfig objects translate into commands that are then deployed to the device. These commands are already issued in global configuration mode. Therefore, do not include the **enable** and **configure terminal** commands as part of the FlexConfig object.
- Determine what types of variables you will need, and create any policy objects that you will require. You cannot create objects for variables while editing a FlexConfig object.
- Ensure that your commands do not conflict in any way with the VPN or access control configuration on the devices.
- If there is more than one set of commands for an interface, only the last set of commands is deployed. Therefore, we recommend you not use beginning and ending commands to configure interfaces. For an example of configuring interfaces, see the ISIS_Interface_Configuration predefined FlexConfig object.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **FlexConfig** > **FlexConfig Object** from the list of object types.

Step 3 Do one of the following:

- Click **Add FlexConfig Object** to create a new object.
- Click **Edit** ( ) to edit an existing object.
- Click **View** ( ) to see the contents of a predefined object.
- If you want to edit a predefined object, click **Copy** ( ) to create a new object with the same contents.

Step 4 Enter a **Name** and optionally, a description for the object.

Step 5 In the object body area, enter the commands and instructions to produce the required configuration.

The object content is a sequence of scripting commands and configuration commands that generate a valid ASA software command sequence. The Firepower Threat Defense device uses ASA software commands to configure some features. For more information on scripting and configuration commands, see:

- Template Scripts, on page 579
- CLI Commands in FlexConfig Objects, on page 576

You can use variables to supply information that can be known only at runtime, or which can differ from device to device. You simply type in processing variables, but you must use the **Insert** menu to add variables that are associated with policy objects or system variables. For a complete discussion of variables, see FlexConfig Variables, on page 579.

- To insert system variables, choose **Insert** > **Insert System Variable** > ***Variable Name***. For a detailed explanation of these variables, see FlexConfig System Variables, on page 585.
- To insert policy object variables, choose **Insert** > **Insert Policy Object** > ***Object Type***, selecting the appropriate type of object. Then, give the variable a name (which can be the same name as the associated policy object), select the object to associate with the variable, and click **Save**. For a detailed explanation of these types, see FlexConfig Policy Object Variables, on page 584. For more detail on the procedure, see Add a Policy Object Variable to a FlexConfig Object, on page 596.

Note You must use the **Insert** menu to create a new policy object or system variable. However, for subsequent uses of that variable, you must type it in, $ included. This is also true for system variables: the first time you use it, add it from the **Insert** menu. Then, type it out for subsequent uses. If you use the **Insert** menu more than once for a system variable, the system variable is added to the Variables list multiple times, and the FlexConfig will not validate, meaning you cannot save your changes. For processing variables (those not associated with a policy object or system variable), simply type in the variable.

Step 6 Choose the deployment frequency and type.

- **Deployment**—Whether to deploy the commands in the object **Once** or **Everytime**. The only way to choose the right option is to test the results of deployment.

Start by selecting **Everytime**. Then, after you attach the object to a FlexConfig policy, deploy the configuration. After a successful deployment, come back to the FlexConfig policy and preview the configuration for one of the assigned devices as described in Preview the FlexConfig Policy, on page 600. If the section labeled `###CLI generated from managed features ###` contains commands that clear or negate the commands in the object, and the `###Flex-config Appended CLI ###` section contains the commands to reconfigure the feature, you know that **Everytime** is the right option.

Even if you do not see negate commands, make some minor change to the device configuration, then run another deployment. If the deployment completes successfully, you can check the deployment transcript (see Verify the Deployed Configuration, on page 601). If you see that the commands were issued again (even when they were already configured) without error, then you can keep **Everytime**.

Change to **Once** only if the system does not first negate the commands in the object before issuing them again, or if the deployment results in errors that are specific to the command. In some cases, the system does not allow you to issue a command that is already configured, but this is the exception.

Some additional tips:

- If the FlexConfig object points to system-managed objects such as network or ACL objects, choose **Everytime**. Otherwise, updates to the objects might not get deployed.
- Use **Once** if the only thing you do in the object is to clear a configuration. Then, remove the object from the FlexConfig policy after the next deployment.

- **Type**—Select one of the following:
  - **Append**—(The default.) Commands in the object are put at the end of the configurations generated from the Firepower Management Center policies. You must use Append if you use policy object variables, which point to objects generated from managed objects. If commands generated for other policies overlap with those specified in the object, you should select this option so your commands are not overwritten. This is the safest option.
  - **Prepend**—Commands in the object are put at the beginning of the configurations generated from the Firepower Management Center policies. You would typically use prepend for commands that clear or negate a configuration.

Step 7 (Optional.) Click **Validate** above the object body to check the integrity of the script.

The object is always validated when you click **Save**. You cannot save an invalid object.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Add a Policy Object Variable to a FlexConfig Object

You can insert variables into a FlexConfig policy object that are associated with other types of policy object. When the FlexConfig is deployed to a device, these variables resolve to the names or content of the associated object.

Use the following procedure for the first use of a policy object variable in a FlexConfig object. If you need to refer to the object again, type in the variable (including the $ sign). To understand how to use these variables, see How to Process Variables, on page 580.

Procedure

Step 1 Choose **Insert** > **Insert Policy Object** > ***Object Type***, selecting the appropriate type of object.

Step 2 Enter a name for the variable, and optionally, a description.

The name must be unique within the context of the FlexConfig object. It cannot include spaces. You are allowed to use the exact same name as the object associated with the variable.

Step 3 Select the object to associate with the variable and click **Add** to move it to the **Selected Object** list.

You can associate a variable with a single object only.

Note For text objects, you can select any of the predefined objects as needed. However, many of these objects have no default values. You must update the objects to add the required values either directly or as overrides for the device to which you will deploy the FlexConfig object. Trying to deploy a FlexConfig without updating these objects typically results in deployment errors.

Step 4 Click **Save**.

The variable appears in the Variables list at the bottom of the FlexConfig object editor.

Configure FlexConfig Text Objects

Use text objects in FlexConfig objects as the target of policy object variables. You can use variables to supply information that can be known only at runtime, or which can differ from device to device. During deployment, variables that point to text objects are replaced by the content of the text object.

Text objects contain free-form strings, which can be keywords, interface names, numbers, IP addresses, and so forth. The content depends on how you will use the information within a FlexConfig script.

Before creating or editing a text object, determine exactly what content you will need. This includes how you intend to process the object, which will help you decide between creating a single string or multiple string object. Read the following topics:

- FlexConfig Variables, on page 579
- How to Process Variables, on page 580

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **FlexConfig** > **Text Object** from the list of object types.

Step 3 Do one of the following:

- Click **Add Text Object** to create a new object.

- Click **Edit** ( ) to edit an existing object. You are allowed to edit the predefined text objects, which is required if you intended to use the predefined FlexConfig objects.

Step 4 Enter a **Name** and optionally, a description for the object.

Step 5 (New objects only.) Choose a **Variable Type** from the drop-down list:

- **Single**—If the object should contain a single text string.
- **Multiple**—If the object should contain a list of text strings.

You cannot change the variable type after you save the object.

Step 6 If the variable type is **Multiple**, use the up and down arrows to specify a **Count**.

Rows are added or removed from the object as you change the number.

Step 7 Add content to the object.

You can either click in the text box next to a variable number and type in a value, or you can set up device overrides for each device that will be assigned a FlexConfig object that uses the text object. You can also do both, in which case the values configured in the base object act as default values in cases where an override does not exist for a given device.

When editing predefined objects, it is a good practice to use device overrides, so that the system defaults remain in place for other users who might need to use the object in different FlexConfig policies. The approach you take depends on the requirements of your organization.

Tip Some predefined objects require multiple values where each value serves a specific purpose. Read the description text carefully to determine the expected values in the object. In some cases, the instructions specify that you must use overrides instead of changing the base values. In the case of enableInspectProtocolList, you are prevented from entering protocols whose inspection is incompatible with Snort inspection.

If you decide to use device overrides, do the following.

a) Select **Allow Overrides**.
b) Expand the Overrides area (if necessary) and click **Add**.

   If an override already exists for the device, click edit for the override to change it.

c) On **Targets** in the Add Object Override dialog box, select the device for which you are defining values and click **Add** to move it to the Selected Devices list.
d) Click **Override**, adjust the **Count** as needed, then click in the variable fields and type in the values for the device.
e) Click **Add**.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configure the FlexConfig Policy

A FlexConfig policy contains two ordered lists of FlexConfig objects, one prepended list and one appended list. For an explanation of prepend/append, see Configure FlexConfig Objects, on page 594.

FlexConfig policies are shared policies that you can assign to multiple devices.

Procedure

Step 1 Choose **Devices** > **FlexConfig**.

Step 2 Do one of the following:

- Click **New Policy** to create a new FlexConfig Policy. You are prompted to enter a name. Optionally, select devices in the Available Devices list and click **Add to Policy** to assign devices. Click **Save**.
- Click **Edit** ( ) to edit an existing Policy. You can change the name or description by clicking them in edit mode.
- Click **Copy** ( ) to create a new policy with the same contents. You are prompted for a name. Device assignments are not retained for the copy.
- Click delete to remove a policy you no longer need.

Step 3 Select the FlexConfig objects required for the policy from the **Available FlexConfig** list and click > to add them to the policy.

Objects are automatically added to the prepended or appended list based on the deployment type specified in the FlexConfig object.

To remove a selected object, click **Delete** ( ) next to an object.

Step 4 For each selected object, click **View** ( ) next to the object to identify the variables used in the object.

Except for system variables, which start with SYS, you need to ensure that the objects associated with the variables are not empty. A blank or brackets with nothing between them, [ ], indicate an empty object. You will need to edit these objects before deploying the policy.

Note If you use object overrides, those values will not show up in this view. Thus, an empty default value does not necessarily mean that you have not updated the object with the required values. Previewing the configuration will show whether the variables resolve correctly for a given device. See Preview the FlexConfig Policy, on page 600.

Step 5 Click **Save**.

What to do next

- Set target devices for the policy; see Set Target Devices for a FlexConfig Policy, on page 600.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Set Target Devices for a FlexConfig Policy

When you create a FlexConfig policy, you can select the devices that use the policy. You can subsequently change device assignments for the policy as described below.

Note Normally, when you unassign a policy from a device, the system automatically removes the associated configuration upon the next deployment. However, because FlexConfig objects are scripts for deploying customized commands, simply unassigning a FlexConfig policy from a device does not remove the commands that were configuring by the FlexConfig objects. If your intention is to remove FlexConfig-generated commands from a device's configuration, see Remove Features Configured Using FlexConfig, on page 603.

Procedure

Step 1 Choose **Devices** > **FlexConfig** and edit a FlexConfig policy.

Step 2 Click **Policy Assignments**.

Step 3 On **Targeted Devices**, build your target list:

- Add—Choose one or more **Available Devices**, then click **Add to Policy** or drag and drop into the list of **Selected Devices**. You can assign the policy to devices, high availability pairs, and clustered devices.
- Delete—Click **Delete** ( ) next to a single device, or select multiple devices, right-click, then choose **Delete Selection**.

Step 4 Click **OK** to save your selection.

Step 5 Click **Save** to save the FlexConfig policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Preview the FlexConfig Policy

Preview a FlexConfig policy to see how the FlexConfig objects get translated into CLI commands. The preview shows the commands that will be generated for a selected device from the scripts and variables used in the FlexConfig objects. The variables are resolved based on the configuration for the device, so you get a clear idea of what will be deployed.

Use the preview to look for potential problems in the FlexConfig objects. Correct the objects until the preview shows the expected results.

You must preview the configuration separately for each device, because the variables can resolve differently based on the device configuration.

Procedure

Step 1 Choose **Devices** > **FlexConfig** and edit a FlexConfig policy.

Step 2 If there are any pending changes, click **Save**.

The preview shows results only for those FlexConfig objects that were in the most recently saved version of the policy. You must save the policy to see a preview of newly-added objects.

Step 3 Click **Preview Config**.

Step 4 Choose a device from the **Select Device** drop-down list.

The system retrieves information from the device and configured policies, and determines what CLI commands will be generated on the next deployment to the device. You can select the output and use Ctrl+C to copy it to the clipboard, where you can paste it into a text file for further analysis.

The preview includes the following sections:

- Flex-config Prepended CLI—These are the commands generated by FlexConfigs that are prepended to the configuration.
- CLI generated from managed features—These are commands generated for policies configured in Firepower Management Center. Commands are generated for new or changed policies since the last successful deployment to the device. These commands do not represent all commands needed to implement the assigned policies. No commands in this section are generated from FlexConfig objects.
- Flex-config Appended CLI—These are the commands generated by FlexConfigs that are appended to the configuration.

Step 5 Click **Close** to close the preview dialog.

Verify the Deployed Configuration

After you deploy a FlexConfig policy to a device, verify that the deployment was successful and that the resulting configuration is what you expected. Also, verify that the device is performing as expected.

Procedure

Step 1 To verify that deployment was successful:

a) Click **System Status** in the menu bar, which is unnamed between **Deploy** and **System**.

The icon looks like one of the following, and it might include a number if there are errors:

- **Indicates No Warnings** — Indicates no warnings or errors are present on the system.
- Indicates **One or More Warnings** — Indicates one or more warnings and no errors are present on the system.
- **Indicates One or More Errors** — Indicates one or more errors and any number of warnings are present on the system.

b) On **Deployments**, verify that the deployment was successful.

c) To see more detailed information, especially for failed deployments, click **Show History**.
d) Select the deployment job in the list of jobs in the left column.

   Jobs are listed in reverse chronological order, with the most recent job at the top of the list.

e) Click download in the **Transcript** column for the device in the right column.

   The deployment transcript includes commands sent to the device, and any responses returned from the device. These response can be informative messages or error messages. For failed deployments, look for messages that indicate errors with the commands that you sent through FlexConfig. These errors can help you correct the script in the FlexConfig object that is trying to configure the commands.

   Note There is no distinction made in the transcript between commands sent for managed features and those generated from FlexConfig policies.

   For example, the following sequence shows that Firepower Management Center (FMC) sent commands to configure GigabitEthernet0/0 with the logical name outside. The device responded that it automatically set the security level to 0. Firepower Threat Defense does not use the security level for anything. Messages relevant to FlexConfig are in the CLI Apply section of the transcript.

```
========= CLI APPLY =========

FMC >> interface GigabitEthernet0/0
FMC >>  nameif outside
FTDv 192.168.0.152 >> [info] : INFO: Security level for "outside" set to 0 by default.
```

Step 2 Verify that the deployed configuration includes the expected commands.

You can do this by making an SSH connection to the device's management IP address. Use the **show running-config** command to view the configuration.

Alternatively, use the CLI tool within Firepower Management Center.

a) Choose **System** > **Health** > **Monitor** and click the name of the device.

   You might need to click the open/close arrow in the **Count** column in the Status table to see any devices.

b) Click **Advanced Troubleshooting**.
c) Click **Threat Defense CLI**.
d) Select **show** as the command, and type **running-config** as the parameter.
e) Click **Execute**.

   The running configuration appears in the text box. You can select the configuration and press Ctrl+C, then paste it into a text file for later analysis.

Step 3 Verify that the device is performing as expected.

Use the **show** commands related to the feature to see detailed information and statistics. For example, if you enabled additional protocol inspections, the **show service-policy** command provides this information. The exact commands to use are feature-dependent and should be mentioned in the ASA configuration guide and command reference you used to learn how to configure the feature.

If commands that show statistics indicate that numbers are not changing (for example, hit counts, connection counts, and so forth), the configuration might be valid but not meaningful. If you know that traffic is going through the device that should show up in statistics, look for what is missing in your configuration. For example, NAT or access rules might be dropping or changing traffic before a feature can act on it.

You can use the **show** commands from an SSH session or through the Firepower Management Center CLI tool.

However, if the **show** command that you need to use is not available directly within the Firepower Threat Defense CLI, you will need make an SSH connection to the device to use the commands. From the CLI, enter the following command sequence to enter Privileged EXEC mode within the diagnostic CLI. From there, you should be able to enter these otherwise unsupported **show** commands.

```
> system support diagnostic-cli
Attaching to Diagnostic CLI ... Press 'Ctrl+a then d' to detach.
Type help or '?' for a list of available commands.
firepower> enable
Password: <press enter, do not enter a password>
firepower#
```

Remove Features Configured Using FlexConfig

If you decide you need to remove a set of configuration commands you configured using FlexConfig, you might need to manually remove that configuration. Unassigning the FlexConfig policy from a device might not remove all of the configuration.

To manually remove the configuration, you create new FlexConfig objects to clear or negate the configuration commands.

Before you begin

To determine if you need to manually remove some or all of the configuration generated by an object:

1. Examine the configuration preview, as described in Preview the FlexConfig Policy, on page 600. If the `###CLI generated from managed features ###` section contains the clear or negate commands to remove all of the commands in the FlexConfig object, then you can simply remove the object from the FlexConfig policy, save, and redeploy.
2. Remove the object from the FlexConfig policy, save the change, then preview the configuration again. If the `###CLI generated from managed features ###` section still does not include the required clear or negate commands, you must follow this procedure to manually remove the configuration.

Procedure

Step 1 Choose **Objects** > **Object Management** and create the FlexConfig Objects to clear or negate the configuration commands.

If a feature has a **clear** command that can remove all configuration settings, then use that command. For example, the predefined Eigrp_Unconfigure_All object contains a single command that removes all EIGRP-related configuration commands:

```
clear configure router eigrp
```

If there is not a **clear** command for the feature, you need to use the **no** form of each command you want to remove. For example, the predefined Sysopt_basic_negate object removes the commands configured through the predefined Sysopt_basic object.

```
no sysopt traffic detailed-statistics

no sysopt connection timewait
```

You would typically configure a FlexConfig object that removes configurations as a prepended, deploy once object.

Step 2 Choose **Devices** > **FlexConfig** and create a new FlexConfig policy or edit the existing policy.

If you want to preserve the FlexConfig policy that deploys the configuration commands, create a new policy specifically for negating the commands, and assign the devices to the policy. Then, add the new FlexConfig objects to the policy.

If you want to completely remove the FlexConfig configuration objects from all devices, you can simply delete those commands from the existing FlexConfig policy and replace them with the objects that negate the configuration.

Step 3 Click **Save** to save the FlexConfig policy.

Step 4 Click **Preview Config** and verify that the clear and negation commands are generating correctly.

Step 5 Click **Deploy** in the menu bar, select the device, and click **Deploy**.

Wait for deployment to complete.

Step 6 Verify that the commands were removed.

View the running configuration on the device to confirm that the commands are removed. For more detailed information, see Verify the Deployed Configuration, on page 601.

Step 7 While editing the FlexConfig policy, click **Policy Assignments** and remove the device. Optionally, remove the FlexConfig Objects from the policy.

Assuming that the FlexConfig policy simply removes the unwanted configuration commands, there is no need to keep the policy assigned to the device after the removal is complete.

However, if the FlexConfig policy retains options that you still want configured on the device, remove the negation objects from the policy. They are no longer needed.

Convert from FlexConfig to Managed Feature

Each software release adds managed features to the product, that is, features that you configure directly through policies that are controlled outside of FlexConfig. This can deprecate FlexConfig commands that you are currently using; your configurations are not automatically converted. After the upgrade, you cannot assign or create FlexConfig objects using the newly deprecated commands. After upgrading software, examine your FlexConfig policies and objects.

When a feature you configured using FlexConfig starts to be supported as a managed feature, you must convert from using FlexConfig to using the managed feature. In most cases, your existing FlexConfig configurations continue to work post-upgrade and you can still deploy. However, in some cases, using deprecated commands can cause deployment issues. Configuring a feature in both the GUI and FlexConfig is not supported.

Procedure

Step 1 Remove the FlexConfig, as explained in Remove Features Configured Using FlexConfig, on page 603.

Step 2 Configure the settings in the newly supported managed feature.

The release notes have a list of new features for the release.

History for FlexConfig

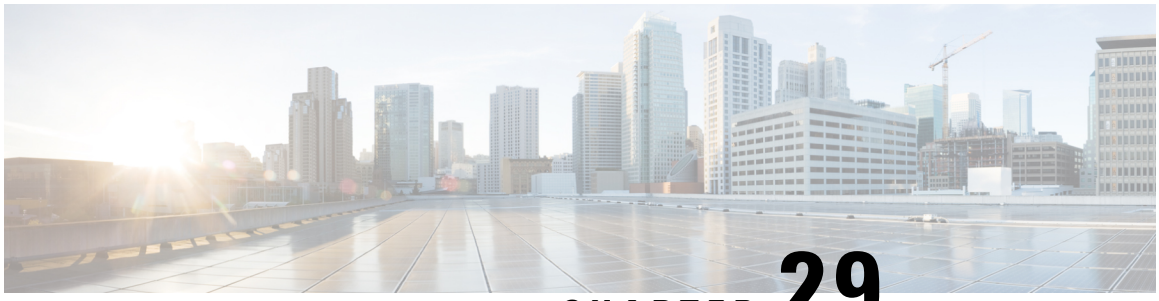
Feature	Version	Description
FlexConfig.	6.2	The FlexConfig feature allows you use the Firepower Management Center to deploy ASA CLI template-based functionality to Firepower Threat Defense devices. This feature allows you to enable some of the most valuable ASA functions that are not currently available on Firepower Threat Defense devices. This functionality is structured as templates and objects that work together in a policy. The default templates are officially supported by Cisco TAC. New screen: **Devices** > **FlexConfig**. Also, under **Objects** > **Object Management**, **FlexConfig** > **FlexConfig Objects** and **FlexConfig** > **Text Object**. Supported platforms: Firepower Threat Defense



PART VIII

Firepower Threat Defense Interfaces and Device Settings

- Interface Overview for Firepower Threat Defense, on page 609
- Regular Firewall Interfaces for Firepower Threat Defense, on page 617
- Inline Sets and Passive Interfaces for Firepower Threat Defense, on page 653
- DHCP and DDNS Services for Threat Defense, on page 663
- Quality of Service (QoS) for Firepower Threat Defense, on page 673



CHAPTER 29

Interface Overview for Firepower Threat Defense

The Firepower Threat Defense device includes data interfaces that you can configure in different modes, as well as a management/diagnostic interface.

- Management/Diagnostic Interface, on page 609
- Interface Mode and Types, on page 610
- Security Zones and Interface Groups, on page 611
- Auto-MDI/MDIX Feature, on page 612
- Default Settings for Interfaces, on page 612
- Enable the Physical Interface and Configure Ethernet Settings, on page 613
- Sync Interface Changes with the Firepower Management Center, on page 614

Management/Diagnostic Interface

he physical management interface is shared between the Diagnostic logical interface and the Management logical interface.

Management Interface

The Management interface is separate from the other interfaces on the device. It is used to set up and register the device to the Firepower Management Center. It uses its own IP address and static routing. You can configure its settings at the CLI using the **configure network** command. If you change the IP address at the CLI after you add it to the Firepower Management Center, you can match the IP address in the Firepower Management Center in the **Devices** > **Device Management** > **Devices** > **Management** area.

Diagnostic Interface

The Diagnostic logical interface can be configured along with the rest of the data interfaces on the **Devices** > **Device Management** > **Interfaces** screen. Using the Diagnostic interface is optional (see the routed and transparent mode deployments for scenarios). The Diagnostic interface only allows management traffic, and does not allow through traffic. It does not support SSH; you can SSH to data interfaces or to the Management interface only. The Diagnostic interface is useful for SNMP or syslog monitoring.

Interface Mode and Types

You can deploy Firepower Threat Defense interfaces in two modes: Regular firewall mode and IPS-only mode. You can include both firewall and IPS-only interfaces on the same device.

Regular Firewall Mode

Firewall mode interfaces subject traffic to firewall functions such as maintaining flows, tracking flow states at both IP and TCP layers, IP defragmentation, and TCP normalization. You can also optionally configure IPS functions for this traffic according to your security policy.

The types of firewall interfaces you can configure depends on the firewall mode set for the device: routed or transparent mode. See Transparent or Routed Firewall Mode for Firepower Threat Defense, on page 545 for more information.

- Routed mode interfaces (routed firewall mode only)—Each interface that you want to route between is on a different subnet.
- Bridge group interfaces (routed and transparent firewall mode)—You can group together multiple interfaces on a network, and the Firepower Threat Defense device uses bridging techniques to pass traffic between the interfaces. Each bridge group includes a Bridge Virtual Interface (BVI) to which you assign an IP address on the network. In routed mode, the Firepower Threat Defense device routes between BVIs and regular routed interfaces. In transparent mode, each bridge group is separate and cannot communicate with each other.

IPS-Only Mode

IPS-only mode interfaces bypass many firewall checks and only support IPS security policy. You might want to implement IPS-only interfaces if you have a separate firewall protecting these interfaces and do not want the overhead of firewall functions.

Note The firewall mode only affects regular firewall interfaces, and not IPS-only interfaces such as inline sets or passive interfaces. IPS-only interfaces can be used in both firewall modes.

IPS-only interfaces can be deployed as the following types:

- Inline Set, with optional Tap mode—An inline set acts like a bump on the wire, and binds two interfaces together to slot into an existing network. This function allows the FTD to be installed in any network environment without the configuration of adjacent network devices. Inline interfaces receive all traffic unconditionally, but all traffic received on these interfaces is retransmitted out of an inline set unless explicitly dropped.

  With tap mode, the FTD is deployed inline, but the network traffic flow is undisturbed. Instead, the FTD makes a copy of each packet so that it can analyze the packets. Note that rules of these types do generate intrusion events when they are triggered, and the table view of intrusion events indicates that the triggering packets would have dropped in an inline deployment. There are benefits to using tap mode with FTDs that are deployed inline. For example, you can set up the cabling between the FTD and the network as if the FTD were inline and analyze the kinds of intrusion events the FTD generates. Based on the results, you can modify your intrusion policy and add the drop rules that best protect your network without impacting its efficiency. When you are ready to deploy the FTD inline, you can disable tap mode and

begin dropping suspicious traffic without having to reconfigure the cabling between the FTD and the network.

Note Tap mode *significantly* impacts FTD performance, depending on the traffic.

Note Inline sets might be familiar to you as "transparent inline sets," but the inline interface type is unrelated to the transparent firewall mode or the firewall-type interfaces.

- Passive or ERSPAN Passive—Passive interfaces monitor traffic flowing across a network using a switch SPAN or mirror port. The SPAN or mirror port allows for traffic to be copied from other ports on the switch. This function provides the system visibility within the network without being in the flow of network traffic. When you configure the FTD in a passive deployment, the FTD cannot take certain actions such as blocking or shaping traffic. Passive interfaces receive all traffic unconditionally. and no traffic received on these interfaces is retransmitted. Encapsulated remote switched port analyzer (ERSPAN) interfaces allow you to monitor traffic from source ports distributed over multiple switches, and uses GRE to encapsulate the traffic. ERSPAN interfaces are only allowed when the FTD is in routed firewall mode.

Note Using SR-IOV interfaces as passive interfaces on NGFWv is not supported on some Intel network adapters (such as Intel X710 or 82599) using SR-IOV drivers due to a promiscuous mode restriction. In such cases, use a network adapter that supports this functionality. See Intel Ethernet Products for more information on Intel network adapters.

Security Zones and Interface Groups

Each interface can be assigned to a *security zone* and/or *interface group*. You then apply your security policy based on zones or groups. For example, you can assign the "inside" interface to the "inside" zone; and the "outside" interface to the "outside" zone. You can configure your access control policy to enable traffic to go from inside to outside, but not from outside to inside, for example. Note that the interface or zone name itself does not provide any default behvior in regards to the security policy, we recommend using names that are self-describing to avoid mistakes in future configuration. A good name signifies a logical segment or traffic specification, for example:

- Names of internal interfaces—InsideV110, InsideV160, InsideV195
- Names of DMZ interfaces—DMZV11, DMZV12, DMZV-TEST
- Names of external interfaces—Outside-ASN78, Outside-ASN91

Some policies only support security zones, while other policies support zones and groups. For specifics, see Interface Objects: Interface Groups and Security Zones, on page 410. You can create security zones and interface groups on the **Objects** page. You can also add a zone when you are configuring the interface. You

can only add interfaces to the correct zone type for your interface, either Passive, Inline, Routed, or Switched zone types.



Note Policies that apply to **any** zone (a global policy) apply to interfaces in zones as well as any interfaces that are not assigned to a zone.

The Diagnostic/Management interface does not belong to a zone or interface group.

Auto-MDI/MDIX Feature

For RJ-45 interfaces, the default auto-negotiation setting also includes the Auto-MDI/MDIX feature. Auto-MDI/MDIX eliminates the need for crossover cabling by performing an internal crossover when a straight cable is detected during the auto-negotiation phase. Either the speed or duplex must be set to auto-negotiate to enable Auto-MDI/MDIX for the interface. If you explicitly set both the speed and duplex to a fixed value, thus disabling auto-negotiation for both settings, then Auto-MDI/MDIX is also disabled. For Gigabit Ethernet, when the speed and duplex are set to 1000 and full, then the interface always auto-negotiates; therefore Auto-MDI/MDIX is always enabled and you cannot disable it.

Default Settings for Interfaces

This section lists default settings for interfaces.

Default State of Interfaces

The default state of an interface depends on the type.

- Physical interfaces—Disabled. The exception is the interface that is enabled for initial setup.
- Redundant Interfaces—Enabled. However, for traffic to pass through the redundant interface, the member physical interfaces must also be enabled.
- VLAN subinterfaces—Enabled. However, for traffic to pass through the subinterface, the physical interface must also be enabled.
- EtherChannel port-channel interfaces (ASA models)—Enabled. However, for traffic to pass through the EtherChannel, the channel group physical interfaces must also be enabled.
- EtherChannel port-channel interfaces (Firepower models)—Disabled.



Note For the Firepower 4100/9300, you can administratively enable and disable interfaces in both the chassis and in the FMC. For an interface to be operational, the interface must be enabled in both operating systems. Because the interface state is controlled independently, you may have a mismatch between the chassis and FMC.

Default Speed and Duplex

By default, the speed and duplex for copper (RJ-45) interfaces are set to auto-negotiate.

By default, the speed and duplex for fiber (SFP) interfaces are set to the maximum speed, with auto-negotiation enabled.

Enable the Physical Interface and Configure Ethernet Settings

This section describes how to:

- Enable the physical interface. By default, physical interfaces are disabled (with the exception of the interface).
- Set a specific speed and duplex. By default, speed and duplex are set to Auto.

This procedure only covers a small subset of Interface settings. Refrain from setting other parameters at this point. For example, you cannot name an interface that you want to use as part of an EtherChannel or redundant interface.

Note For the Firepower 4100/9300, you configure basic interface settings in FXOS. See Configure a Physical Interface, on page 563 for more information.

Before you begin

If you changed the physical interfaces on the device after you added it to the FMC, you need to refresh the interface listing by clicking **Sync Interfaces from device** on the top left of **Interfaces**.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

Step 3 Enable the interface by checking the **Enabled** check box.

Step 4 (Optional) Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

Step 5 (Optional) Set the duplex and speed by clicking **Hardware Configuration**.

- **Duplex**—Choose **Full**, **Half**, or **Auto**. Auto is the default for RJ-45 interfaces. You cannot select Auto for SFP interfaces.
- **Speed**—Choose **Auto** to have the interface negotiate the speed, link status, and flow control (Auto is only available for RJ-45 interfaces), or pick a specific speed: **10**, **100**, **1000** Mbps. For SFP interfaces, setting the speed enables auto-negotiation of link status and flow control. For SFP interfaces, depending on your hardware, you can select **No Negotiate** to set the speed to 1000 and disable link negotiation.

Step 6 In the **Mode** drop-down list, choose one of the following:.

- **None**—Choose this setting for regular firewall interfaces and inline sets. The mode will automatically be changed to Routed, Switched, or Inline based on futher configuration.
- **Passive**—Choose this setting for passive IPS-only interfaces.
- **Erspan**—Choose this setting for ERSPAN passive IPS-only interfaces.

Step 7 Click **OK**.

Step 8 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Step 9 Continue configuring interfaces.

- Regular Firewall Interfaces for Firepower Threat Defense, on page 617
- Inline Sets and Passive Interfaces for Firepower Threat Defense, on page 653

Sync Interface Changes with the Firepower Management Center

Interface configuration changes on the device can cause the FMC and the device to get out of sync. The FMC can detect interface changes by one of the following methods:

- Event sent from the device
- Sync when you deploy from the FMC

  If the FMC detects interface changes when it attempts to deploy, the deploy will fail. You must first accept the interface changes.
- Manual sync

Adding a new interface, or deleting an unused interface has minimal impact on the Firepower Threat Defense configuration. However, deleting an interface that is used in your security policy will impact the configuration. Interfaces can be referenced directly in many places in the Firepower Threat Defense configuration, including access rules, NAT, SSL, identity rules, VPN, DHCP server, and so on. Deleting an interface will delete any configuration associated with that interface. Policies that refer to security zones are not affected. You can also edit the membership of an allocated EtherChannel without affecting the logical device or requiring a sync on the FMC.

This procedure describes how to manually sync device changes if required and how to acknowledge the detected changes. If device changes are temporary, you should not save the changes in the FMC; you should wait until the device is stable, and then re-sync.

Before you begin

- Model Support—FTD
- User Roles:

- Admin
- Access Admin
- Network Admin

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

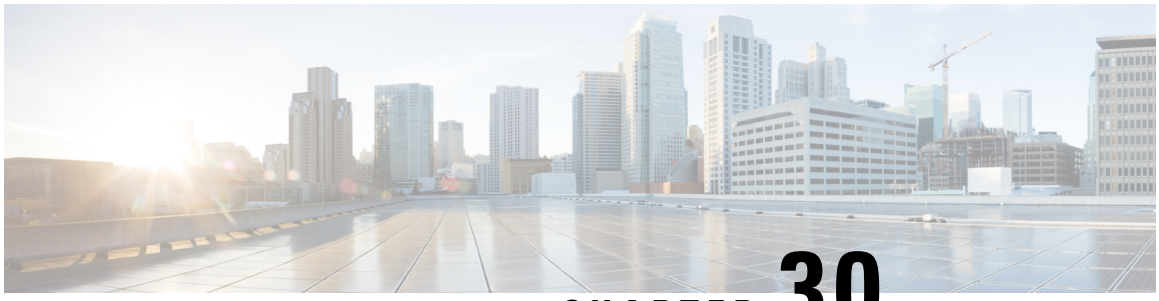
Step 2 If required, click **Sync Device** on the top left of **Interfaces**.

Step 3 After the changes are detected, see the following steps.

a) You will see a red banner on **Interfaces** indicating that the interface configuration has changed. Click the **Click to know more** link to view the interface changes.

b) Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices.



CHAPTER 30

Regular Firewall Interfaces for Firepower Threat Defense

This chapter includes regular firewall Firepower Threat Defense interface configuration including EtherChannels, VLAN subinterfaces, IP addressing, and more.

Note For initial interface configuration on the Firepower 4100/9300, see Configure Interfaces, on page 562.

- Requirements and Prerequisites for Regular Firewall Interfaces, on page 617
- Configure EtherChannel and Redundant Interfaces, on page 617
- Configure VLAN Subinterfaces and 802.1Q Trunking, on page 625
- Configure Routed and Transparent Mode Interfaces, on page 627
- Configure Advanced Interface Settings, on page 641
- History for Regular Firewall Interfaces for Firepower Threat Defense, on page 651

Requirements and Prerequisites for Regular Firewall Interfaces

Model Support

FTD

User Roles

- Admin
- Access Admin
- Network Admin

Configure EtherChannel and Redundant Interfaces

This section tells how to configure EtherChannels and redundant interfaces.

Note For the Firepower 4100/9300, you configure EtherChannels in FXOS. See Add an EtherChannel (Port Channel), on page 564 for more information.

Note Only ASA 5500-X models support redundant interfaces; Firepower models do not support them.

About EtherChannels

This section describes EtherChannels.

About Redundant Interfaces (ASA Platform Only)

A logical redundant interface consists of a pair of physical interfaces: an active and a standby interface. When the active interface fails, the standby interface becomes active and starts passing traffic. You can configure a redundant interface to increase the Firepower Threat Defense device reliability.

You can configure up to 8 redundant interface pairs.

Redundant Interface MAC Address

The redundant interface uses the MAC address of the first physical interface that you add. If you change the order of the member interfaces in the configuration, then the MAC address changes to match the MAC address of the interface that is now listed first. Alternatively, you can assign a manual MAC address to the redundant interface, which is used regardless of the member interface MAC addresses. When the active interface fails over to the standby, the same MAC address is maintained so that traffic is not disrupted.

About EtherChannels

An 802.3ad EtherChannel is a logical interface (called a port-channel interface) consisting of a bundle of individual Ethernet links (a channel group) so that you increase the bandwidth for a single network. A port channel interface is used in the same way as a physical interface when you configure interface-related features.

You can configure up to 48 EtherChannels, depending on how many interfaces your model supports.

Channel Group Interfaces

Each channel group can have up to 8 active interfaces, except for ASA models and the ISA 3000, which supports 16 active interfaces. For switches that support only 8 active interfaces, you can assign up to 16 interfaces to a channel group: while only 8 interfaces can be active, the remaining interfaces can act as standby links in case of interface failure.

All interfaces in the channel group must be the same type and speed. The first interface added to the channel group determines the correct type and speed.

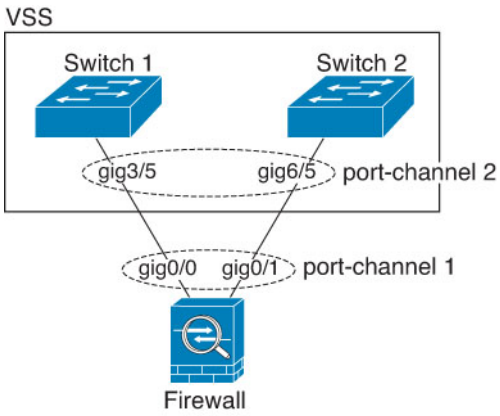
The EtherChannel aggregates the traffic across all the available active interfaces in the channel. The interface is selected using a proprietary hash algorithm, based on source or destination MAC addresses, IP addresses, TCP and UDP port numbers and VLAN numbers.

Connecting to an EtherChannel on Another Device

The device to which you connect the FTD EtherChannel must also support 802.3ad EtherChannels; for example, you can connect to the Catalyst 6500 switch or the Cisco Nexus 7000.

When the switch is part of a Virtual Switching System (VSS) or Virtual Port Channel (vPC), then you can connect FTD interfaces within the same EtherChannel to separate switches in the VSS/vPC. The switch interfaces are members of the same EtherChannel port-channel interface, because the separate switches act like a single switch.

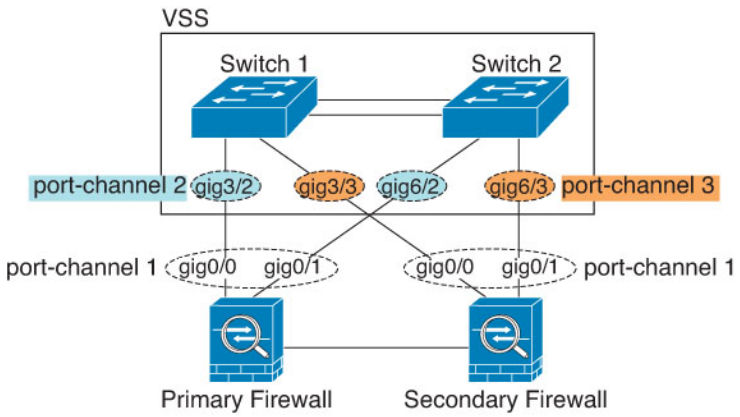
Figure 11: Connecting to a VSS/vPC



Note If the FTD device is in transparent firewall mode, and you place the FTD device between two sets of VSS/vPC switches, then be sure to disable Unidirectional Link Detection (UDLD) on any switch ports connected to the FTD device with an EtherChannel. If you enable UDLD, then a switch port may receive UDLD packets sourced from both switches in the other VSS/vPC pair. The receiving switch will place the receiving interface in a down state with the reason "UDLD Neighbor mismatch".

If you use the FTD device in an Active/Standby failover deployment, then you need to create separate EtherChannels on the switches in the VSS/vPC, one for each FTD device. On each FTD device, a single EtherChannel connects to both switches. Even if you could group all switch interfaces into a single EtherChannel connecting to both FTD devices (in this case, the EtherChannel will not be established because of the separate FTD system IDs), a single EtherChannel would not be desirable because you do not want traffic sent to the standby FTD device.

Figure 12: Active/Standby Failover and VSS/vPC



Link Aggregation Control Protocol

The Link Aggregation Control Protocol (LACP) aggregates interfaces by exchanging the Link Aggregation Control Protocol Data Units (LACPDUs) between two network devices.

You can configure each physical interface in an EtherChannel to be:

- Active—Sends and receives LACP updates. An active EtherChannel can establish connectivity with either an active or a passive EtherChannel. You should use the active mode unless you need to minimize the amount of LACP traffic.
- Passive—Receives LACP updates. A passive EtherChannel can only establish connectivity with an active EtherChannel. Not supported on hardware models.
- On—The EtherChannel is always on, and LACP is not used. An "on" EtherChannel can only establish a connection with another "on" EtherChannel. Not supported on hardware models.

LACP coordinates the automatic addition and deletion of links to the EtherChannel without user intervention. It also handles misconfigurations and checks that both ends of member interfaces are connected to the correct channel group. "On" mode cannot use standby interfaces in the channel group when an interface goes down, and the connectivity and configurations are not checked.

Load Balancing

The FTD device distributes packets to the interfaces in the EtherChannel by hashing the source and destination IP address of the packet (this criteria is configurable). The resulting hash is divided by the number of active links in a modulo operation where the resulting remainder determines which interface owns the flow. All packets with a *hash_value* **mod** *active_links* result of 0 go to the first interface in the EtherChannel, packets with a result of 1 go to the second interface, packets with a result of 2 go to the third interface, and so on. For example, if you have 15 active links, then the modulo operation provides values from 0 to 14. For 6 active links, the values are 0 to 5, and so on.

If an active interface goes down and is not replaced by a standby interface, then traffic is rebalanced between the remaining links. The failure is masked from both Spanning Tree at Layer 2 and the routing table at Layer 3, so the switchover is transparent to other network devices.

EtherChannel MAC Address

All interfaces that are part of the channel group share the same MAC address. This feature makes the EtherChannel transparent to network applications and users, because they only see the one logical connection; they have no knowledge of the individual links.

Firepower Hardware

The port-channel interface uses the MAC address of the internal interface Internal-Data 0/1. Alternatively you can manually configure a MAC address for the port-channel interface. All EtherChannel interfaces on a chassis use the same MAC address, so be aware that if you use SNMP polling, for example, multiple interfaces will have the same MAC address.



Note Member interfaces only use the Internal-Data 0/1 MAC address after a reboot. Prior to rebooting, the member interface uses its own MAC address. If you add a new member interface after a reboot, you will have to perform another reboot to update its MAC address.

ASA Hardware

The port-channel interface uses the lowest numbered channel group interface MAC address as the port-channel MAC address. Alternatively you can manually configure a MAC address for the port-channel interface. We recommend manually configuring a unique MAC address in case the group channel interface membership changes. If you remove the interface that was providing the port-channel MAC address, then the port-channel MAC address changes to the next lowest numbered interface, thus causing traffic disruption.

Guidelines for EtherChannels

Bridge Group

In routed mode, FMC-defined EtherChannels are not supported as bridge group members. EtherChannels on the Firepower 4100/9300 can be bridge group members.

High Availability

- When you use an EtherChannel interface as a High Availability link, it must be pre-configured on both units in the High Availability pair; you cannot configure it on the primary unit and expect it to replicate to the secondary unit because *the High Availability link itself is required for replication*.
- If you use an EtherChannel interface for the state link, no special configuration is required; the configuration can replicate from the primary unit as normal. For the Firepower 4100/9300 chassis, all interfaces, including EtherChannels, need to be pre-configured on both units.
- You can monitor EtherChannel interfaces for High Availability. When an active member interface fails over to a standby interface, this activity does not cause the EtherChannel interface to appear to be failed when being monitored for device-level High Availability. Only when all physical interfaces fail does the EtherChannel interface appear to be failed (for an EtherChannel interface, the number of member interfaces allowed to fail is configurable).
- If you use an EtherChannel interface for a High Availability or state link, then to prevent out-of-order packets, only one interface in the EtherChannel is used. If that interface fails, then the next interface in

the EtherChannel is used. You cannot alter the EtherChannel configuration while it is in use as a High Availability link. To alter the configuration, you need to temporarily disable High Availability, which prevents High Availability from occurring for the duration.

Model Support

- You cannot add EtherChannels in the FMC for the Firepower 4100/9300 or the FTDv. The Firepower 4100/9300 supports EtherChannels, but you must perform all hardware configuration of EtherChannels in FXOS on the chassis.

General EtherChannel Guidelines

- You can configure up to 48 EtherChannels, depending on how many interfaces are available on your model.
- Each channel group can have up to 8 active interfaces, except for ASA models and the ISA 3000, which supports 16 active interfaces. For switches that support only 8 active interfaces, you can assign up to 16 interfaces to a channel group: while only 8 interfaces can be active, the remaining interfaces can act as standby links in case of interface failure.
- All interfaces in the channel group must be the same media type and speed capacity. The media type can be either RJ-45 or SFP; SFPs of different types (copper and fiber) can be mixed. You cannot mix interface capacities (for example 1GB and 10GB interfaces) by setting the speed to be lower on the larger-capacity interface.
- The device to which you connect the FTD EtherChannel must also support 802.3ad EtherChannels.
- The FTD device does not support LACPDUs that are VLAN-tagged. If you enable native VLAN tagging on the neighboring switch using the Cisco IOS **vlan dot1Q tag native** command, then the FTD device will drop the tagged LACPDUs. Be sure to disable native VLAN tagging on the neighboring switch.
- Devices do not support LACP rate fast, except for ASA models and the ISA 3000; LACP always uses the normal rate. This setting is not configurable. Note that the Firepower 4100/9300, which configures EtherChannels in FXOS, has the LACP rate set to fast by default; on these platforms, the rate is configurable.
- In Cisco IOS software versions earlier than 15.1(1)S2, FTD did not support connecting an EtherChannel to a switch stack. With default switch settings, if the FTD EtherChannel is connected cross stack, and if the primary switch is powered down, then the EtherChannel connected to the remaining switch will not come up. To improve compatibility, set the **stack-mac persistent timer** command to a large enough value to account for reload time; for example, 8 minutes or 0 for indefinite. Or, you can upgrade to more a more stable switch software version, such as 15.1(1)S2.
- All the FTD configuration refers to the logical EtherChannel interface instead of the member physical interfaces.

Configure a Redundant Interface (ASA Platform Only)

A logical redundant interface consists of a pair of physical interfaces: an active and a standby interface. When the active interface fails, the standby interface becomes active and starts passing traffic. You can configure a redundant interface to increase the Firepower Threat Defense reliability. By default, redundant interfaces are enabled.

- You can configure up to 8 redundant interface pairs.
- Both member interfaces must be of the same physical type. For example, both must be GigabitEthernet.

Note Redundant interfaces are not supported on the Firepower platform; only ASA 5500-X models support redundant interfaces.

Before you begin

- You cannot add a physical interface to the redundant interface if you configured a name for it. You must first remove the name.

Caution If you are using a physical interface already in your configuration, removing the name will clear any configuration that refers to the interface.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Enable the member interfaces according to Enable the Physical Interface and Configure Ethernet Settings, on page 613.

Step 3 Click **Add Interfaces** > **Redundant Interface**.

Step 4 On the **General** tab, set the following parameters:

a) **Redundant ID**—Set an integer between 1 and 8.

b) **Primary Interface**—Choose an interface from the drop-down list. After you add the interface, any configuration for it (such as an IP address) is removed.

c) **Secondary Interface**—The second interface must be the same physical type as the first interface.

Step 5 Click **OK**.

Step 6 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Step 7 (Optional) Add a VLAN subinterface. See Add a Subinterface, on page 626.

Step 8 Configure the routed or transparent mode interface parameters. See Configure Routed Mode Interfaces, on page 629 or Configure Bridge Group Interfaces, on page 632.

Configure an EtherChannel

This section describes how to create an EtherChannel port-channel interface, assign interfaces to the EtherChannel, and customize the EtherChannel.

Guidelines

- You can configure up to 48 EtherChannels, depending on the number of interfaces for your model.
- Each channel group can have up to 16 active interfaces. For switches that support only 8 active interfaces, you can assign up to 16 interfaces to a channel group: while only 8 interfaces can be active, the remaining interfaces can act as standby links in case of interface failure.
- All interfaces in the channel group must be the same type, speed, and duplex. Half duplex is not supported.

Note For the Firepower 4100/9300, you configure EtherChannels in FXOS. See Add an EtherChannel (Port Channel), on page 564 for more information.

Before you begin

- You cannot add a physical interface to the channel group if you configured a name for it. You must first remove the name.

Note If you are using a physical interface already in your configuration, removing the name will clear any configuration that refers to the interface.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Enable the member interfaces according to Enable the Physical Interface and Configure Ethernet Settings, on page 613.

Step 3 Click **Add Interfaces** > **Ether Channel Interface**.

Step 4 On the **General** tab, set the **Ether Channel ID** to a number between 1 and 48.

Step 5 In the **Available Interfaces** area, click an interface and then click **Add** to move it to the **Selected Interfaces** area. Repeat for all interfaces that you want to make members.

Make sure all interfaces are the same type and speed. The first interface you add determines the type and speed of the EtherChannel. Any non-matching interfaces you add will be put into a suspended state. The FMC does not prevent you from adding non-matching interfaces.

Step 6 (Optional) Click the **Advanced** tab to customize the EtherChannel. Set the following parameters on the **Information** sub-tab:

- **Load Balancing**—Select the criteria used to load balance the packets across the group channel interfaces. By default, the Firepower Threat Defense device balances the packet load on interfaces according to the source and destination IP address of the packet. If you want to change the properties on which the packet is categorized, choose a different set of criteria. For example, if your traffic is biased heavily towards the same source and destination IP addresses, then the traffic assignment to interfaces in the EtherChannel

will be unbalanced. Changing to a different algorithm can result in more evenly distributed traffic. For more information about load balancing, see Load Balancing, on page 620.

- **LACP Mode**—Choose Active, Passive, or On. We recommend using Active mode (the default).
- **Active Physical Interface: Range**—From the left drop-down list, choose the minimum number of active interfaces required for the EtherChannel to be active, between 1 and 16. The default is 1. From the right drop-down list, choose the maximum number of active interfaces allowed in the EtherChannel, between 1 and 16. The default is 16. If your switch does not support 16 active interfaces, be sure to set this command to 8 or fewer.
- **Active Mac Address**—Set a manual MAC address if desired. The mac_address is in H.H.H format, where H is a 16-bit hexadecimal digit. For example, the MAC address 00-0C-F1-42-4C-DE is entered as 000C.F142.4CDE.

Step 7 (Optional) Click the **Hardware Configuration** tab and set the Duplex and Speed to override these settings for all member interfaces. This method provides a shortcut to set these parameters because these parameters must match for all interfaces in the channel group.

Step 8 Click **OK**.

Step 9 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Step 10 (Optional) Add a VLAN subinterface. See Add a Subinterface, on page 626.

Step 11 Configure the routed or transparent mode interface parameters. See Configure Routed Mode Interfaces, on page 629 or Configure Bridge Group Interfaces, on page 632.

Configure VLAN Subinterfaces and 802.1Q Trunking

VLAN subinterfaces let you divide a physical, redundant, or EtherChannel interface into multiple logical interfaces that are tagged with different VLAN IDs. An interface with one or more VLAN subinterfaces is automatically configured as an 802.1Q trunk. Because VLANs let you keep traffic separate on a given physical interface, you can increase the number of interfaces available to your network without adding additional physical interfaces or devices.

Guidelines and Limitations for VLAN Subinterfaces

High Availability

You cannot use a subinterface for the failover or state link.

Additional Guidelines

- Preventing untagged packets on the physical interface—If you use subinterfaces, you typically do not also want the physical interface to pass traffic, because the physical interface passes untagged packets. This property is also true for the active physical interface in a redundant interface pair and for EtherChannel links. Because the physical, redundant, or EtherChannel interface must be enabled for the subinterface to pass traffic, ensure that the physical, redundant, or EtherChannel interface does not pass traffic by not

configuring a name for the interface. If you want to let the physical, redundant, or EtherChannel interface pass untagged packets, you can configure the name as usual.

- You cannot configure subinterfaces on the interface.
- All subinterfaces on the same parent interface must be either bridge group members or routed interfaces; you cannot mix and match.
- The FTD does not support the Dynamic Trunking Protocol (DTP), so you must configure the connected switch port to trunk unconditionally.
- You might want to assign unique MAC addresses to subinterfaces defined on the FTD, because they use the same burned-in MAC address of the parent interface. For example, your service provider might perform access control based on the MAC address. Also, because IPv6 link-local addresses are generated based on the MAC address, assigning unique MAC addresses to subinterfaces allows for unique IPv6 link-local addresses, which can avoid traffic disruption in certain instances on the FTD.

Maximum Number of VLAN Subinterfaces by Device Model

The device model limits the maximum number of VLAN subinterfaces that you can configure. Note that you can configure subinterfaces on data interfaces only, you cannot configure them on the management interface.

The following table explains the limits for each device model.

Model	Maximum VLAN Subinterfaces
Firepower 4100	1024
Firepower 9300	1024
FTDv	50
ASA 5506-X ASA 5506W-X ASA 5506H-X	30
ASA 5508-X	50
ASA 5512-X	100
ASA 5515-X	100
ASA 5516-X	100
ASA 5525-X	200
ASA 5545-X	300
ASA 5555-X	500

Add a Subinterface

Add one or more subinterfaces to a physical, redundant, or port-channel interface.

Note The parent physical interface passes untagged packets. You may not want to pass untagged packets, so be sure not to include the parent interface in your security policy.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Enable the parent interface according to Enable the Physical Interface and Configure Ethernet Settings, on page 613.

Step 3 Click **Add Interfaces** > **Sub Interface**.

Step 4 On **General**, set the following parameters:

a) **Interface**—Choose the physical, redundant, or port-channel interface to which you want to add the subinterface.

b) **Sub-Interface ID**—Enter the subinterface ID as an integer between 1 and 4294967295. The number of subinterfaces allowed depends on your platform. You cannot change the ID after you set it.

c) **VLAN ID**—Enter the VLAN ID between 1 and 4094 that will be used to tag the packets on this subinterface.

This VLAN ID must be unique.

Step 5 Click **OK**.

Step 6 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Step 7 Configure the routed or transparent mode interface parameters. See Configure Routed Mode Interfaces, on page 629 or Configure Bridge Group Interfaces, on page 632.

Configure Routed and Transparent Mode Interfaces

This section includes tasks to complete the regular interface configuration for all models in routed or transparent firewall mode.

About Routed and Transparent Mode Interfaces

Firewall mode interfaces subject traffic to firewall functions such as maintaining flows, tracking flow states at both IP and TCP layers, IP defragmentation, and TCP normalization. You can also optionally configure IPS functions for this traffic according to your security policy.

The types of firewall interfaces you can configure depends on the firewall mode set for the device: routed or transparent mode. See Transparent or Routed Firewall Mode for Firepower Threat Defense, on page 545 for more information.

- Routed mode interfaces (routed firewall mode only)—Each interface that you want to route between is on a different subnet.
- Bridge group interfaces (routed and transparent firewall mode)—You can group together multiple interfaces on a network, and the Firepower Threat Defense device uses bridging techniques to pass traffic between the interfaces. Each bridge group includes a Bridge Virtual Interface (BVI) to which you assign an IP address on the network. In routed mode, the Firepower Threat Defense device routes between BVIs and regular routed interfaces. In transparent mode, each bridge group is separate and cannot communicate with each other.

Dual IP Stack (IPv4 and IPv6)

The FTD device supports both IPv6 and IPv4 addresses on an interface. Make sure you configure a default route for both IPv4 and IPv6.

Guidelines and Limitations for Routed and Transparent Mode Interfaces

High Availability, Clustering

- Do not configure failover links with the procedures in this chapter. See the High Availability chapter for more information.
- For cluster interfaces, see the clustering chapter for requirements.
- When you use High Availability, you must set the IP address and standby address for data interfaces manually; DHCP and PPPoE are not supported. Set the standby IP addresses on the **Devices** > **Device Management** > **High Availability** tab in the **Monitored Interfaces** area. See the High Availability chapter for more information.

IPv6

- IPv6 is supported on all interfaces.
- You can only configure IPv6 addresses manually in transparent mode.
- The FTD device does not support IPv6 anycast addresses.

Model Guidelines

- For the FTDv on VMware with bridged ixgbevf interfaces, bridge groups are not supported.

Transparent Mode and Bridge Group Guidelines

- You can create up to 250 bridge groups, with 64 interfaces per bridge group.
- Each directly-connected network must be on the same subnet.
- The FTD device does not support traffic on secondary networks; only traffic on the same network as the BVI IP address is supported.

- An IP address for the BVI is required for each bridge group for to-the-device and from-the-device management traffic, as well as for data traffic to pass through the FTD device. For IPv4 traffic, specify an IPv4 address. For IPv6 traffic, specify an IPv6 address.
- You can only configure IPv6 addresses manually.
- The BVI IP address must be on the same subnet as the connected network. You cannot set the subnet to a host subnet (255.255.255.255).
- Management interfaces are not supported as bridge group members.
- For the FTDv on VMware with bridged ixgbevf interfaces, transparent mode is not supported, and bridge groups are not supported in routed mode.
- In transparent mode, you must use at least 1 bridge group; data interfaces must belong to a bridge group.
- In transparent mode, do not specify the BVI IP address as the default gateway for connected devices; devices need to specify the router on the other side of the FTD as the default gateway.
- In transparent mode, the *default* route, which is required to provide a return path for management traffic, is only applied to management traffic from one bridge group network. This is because the default route specifies an interface in the bridge group as well as the router IP address on the bridge group network, and you can only define one default route. If you have management traffic from more than one bridge group network, you need to specify a regular static route that identifies the network from which you expect management traffic.
- In transparent mode, PPPoE is not supported for the interface.
- Transparent mode is not supported on threat defense virtual instances deployed on Amazon Web Services, Microsoft Azure, Google Cloud Platform, and Oracle Cloud Infrastructure.
- In routed mode, to route between bridge groups and other routed interfaces, you must name the BVI.
- In routed mode, FTD-defined EtherChannel interfaces are not supported as bridge group members. EtherChannels on the Firepower 4100/9300 can be bridge group members.
- Bidirectional Forwarding Detection (BFD) echo packets are not allowed through the FTD when using bridge group members. If there are two neighbors on either side of the FTD running BFD, then the FTD will drop BFD echo packets because they have the same source and destination IP address and appear to be part of a LAND attack.

Additional Guidelines and Requirements

- The FTD supports only one 802.1Q header in a packet and does not support multiple headers (known as Q-in-Q support) for firewall interfaces.**Note:** For inline sets and passive interfaces, the FTD supports Q-in-Q up to two 802.1Q headers in a packet, with the exception of the Firepower 4100/9300, which only supports one 802.1Q header.

Configure Routed Mode Interfaces

This procedure describes how to set the name, security zone, and IPv4 address.

Before you begin

- **Firepower 4100/9300**

1. Configure a Physical Interface, on page 563
2. (Optional) Configure any special interfaces.
   - Add an EtherChannel (Port Channel), on page 564
   - Add a Subinterface, on page 626 in FMC

- (Optional) **All other models:**
  - Configure a Redundant Interface (ASA Platform Only), on page 622
  - Configure an EtherChannel, on page 623
  - Add a Subinterface, on page 626

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

Step 3 In the **Name** field, enter a name up to 48 characters in length.

Step 4 Enable the interface by checking the **Enabled** check box.

Step 5 (Optional) Set this interface to **Management Only** to limit traffic to management traffic; through-the-box traffic is not allowed.

Step 6 (Optional) Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

Step 7 In the **Mode** drop-down list, choose **None**.

Regular firewall interfaces have the mode set to None. The other modes are for IPS-only interface types.

Step 8 From the **Security Zone** drop-down list, choose a security zone or add a new one by clicking **New**.

The routed interface is a Routed-type interface, and can only belong to Routed-type zones.

Step 9 See Configure the MTU, on page 646 for information about the **MTU**.

Step 10 Click the **IPv4** tab. To set the IP address, use one of the following options from the **IP Type** drop-down list.

High Availability and clustering interfaces only support static IP address configuration; DHCP and PPPoE are not supported.

- **Use Static IP**—Enter the IP address and subnet mask. For High Availabilty, you can only use a static IP address. Set the standby IP address on the **Devices** > **Device Management** > **High Availability** tab in the **Monitored Interfaces** area. If you do not set the standby IP address, the active unit cannot monitor the standby interface using network tests; it can only track the link state.
- **Use DHCP**—Configure the following optional parameters:
  - **Obtain default route using DHCP**—Obtains the default route from the DHCP server.

- **DHCP route metric**—Assigns an administrative distance to the learned route, between 1 and 255. The default administrative distance for the learned routes is 1.

- **Use PPPoE**—If the interface is connected to a DSL, cable modem, or other connection to your ISP, and your ISP uses PPPoE to provide your IP address, configure the following parameters:
  - **VPDN Group Name**—Specify a group name of your choice to represent this connection.
  - **PPPoE User Name**—Specify the username provided by your ISP.
  - **PPPoE Password/Confirm Password**—Specify and confirm the password provided by your ISP.
  - **PPP Authentication**—Choose **PAP**, **CHAP**, or **MSCHAP**.

    PAP passes a cleartext username and password during authentication and is not secure. With CHAP, the client returns the encrypted [challenge plus password], with a cleartext username in response to the server challenge. CHAP is more secure than PAP, but it does not encrypt data. MSCHAP is similar to CHAP but is more secure because the server stores and compares only encrypted passwords rather than cleartext passwords as in CHAP. MSCHAP also generates a key for data encryption by MPPE.
  - **PPPoE route metric**—Assign an administrative distance to the learned route. Valid values are from 1 to 255. By default, the administrative distance for the learned routes is 1.
  - **Store Username and Password in Flash**—Stores the username and password in flash memory.

    The Firepower Threat Defense device stores the username and password in a special location of NVRAM.

Step 11 (Optional) See Configure IPv6 Addressing, on page 636 to configure IPv6 addressing on the **IPv6** tab.

Step 12 (Optional) See Configure the MAC Address, on page 646 to manually configure the MAC address on the **Advanced** tab.

Step 13 (Optional) Set the duplex and speed by clicking **Hardware Configuration**.

- **Duplex**—Choose **Full**, **Half**, or **Auto**. Auto is the default for RJ-45 interfaces. You cannot select Auto for SFP interfaces.
- **Speed**—Choose **Auto** to have the interface negotiate the speed, link status, and flow control (Auto is only available for RJ-45 interfaces), or pick a specific speed: **10**, **100**, **1000** Mbps. For SFP interfaces, setting the speed enables auto-negotiation of link status and flow control. For SFP interfaces, depending on your hardware, you can select **No Negotiate** to set the speed to 1000 and disable link negotiation.

Step 14 Click **OK**.

Step 15 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure Bridge Group Interfaces

A bridge group is a group of interfaces that the Firepower Threat Defense device bridges instead of routes. Bridge groups are supported in both transparent and routed firewall mode. For more information about bridge groups, see About Bridge Groups, on page 547.

To configure bridge groups and associated interfaces, perform these steps.

Configure General Bridge Group Member Interface Parameters

This procedure describes how to set the name and security zone for each bridge group member interface. The same bridge group can include different types of interfaces: physical interfaces, VLAN subinterfaces, EtherChannels, and redundant interfaces. The interface is not supported. In routed mode, EtherChannels are not supported.

Before you begin

- **Firepower 4100/9300**
  1. Configure a Physical Interface, on page 563
  2. (Optional) Configure any special interfaces.
     - Add an EtherChannel (Port Channel), on page 564
     - Add a Subinterface, on page 626 in FMC

- (Optional) **All other models:**
  - Configure a Redundant Interface (ASA Platform Only), on page 622
  - Configure an EtherChannel, on page 623
  - Add a Subinterface, on page 626

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

Step 3 In the **Name** field, enter a name up to 48 characters in length.

Step 4 Enable the interface by checking the **Enabled** check box.

Step 5 (Optional) Set this interface to **Management Only** to limit traffic to management traffic; through-the-box traffic is not allowed.

Step 6 (Optional) Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

Step 7 In the **Mode** drop-down list, choose **None**.

Regular firewall interfaces have the mode set to None. The other modes are for IPS-only interface types. After you assign this interface to a bridge group, the mode will show as **Switched**.

Step 8 From the **Security Zone** drop-down list, choose a security zone or add a new one by clicking **New**.

The bridge group member interface is a Switched-type interface, and can only belong to Switched-type zones. Do not configure any IP address settings for this interface. You will set the IP address for the Bridge Virtual Interface (BVI) only. Note that the BVI does not belong to a zone, and you cannot apply access control policies to the BVI.

Step 9 See Configure the MTU, on page 646 for information about the **MTU**.

Step 10 (Optional) Set the duplex and speed by clicking **Hardware Configuration**.

- **Duplex**—Choose **Full**, **Half**, or **Auto**. Auto is the default for RJ-45 interfaces. You cannot select Auto for SFP interfaces.
- **Speed**—Choose **Auto** to have the interface negotiate the speed, link status, and flow control (Auto is only available for RJ-45 interfaces), or pick a specific speed: **10**, **100**, **1000** Mbps. For SFP interfaces, setting the speed enables auto-negotiation of link status and flow control. For SFP interfaces, depending on your hardware, you can select **No Negotiate** to set the speed to 1000 and disable link negotiation.

Step 11 (Optional) See Configure IPv6 Addressing, on page 636 to configure IPv6 addressing on the **IPv6** tab.

Step 12 (Optional) See Configure the MAC Address, on page 646 to manually configure the MAC address on the **Advanced** tab.

Step 13 Click **OK**.

Step 14 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure the Bridge Virtual Interface (BVI)

Each bridge group requires a BVI for which you configure an IP address. The Firepower Threat Defense uses this IP address as the source address for packets originating from the bridge group. The BVI IP address must be on the same subnet as the connected network. For IPv4 traffic, the BVI IP address is required to pass any traffic. For IPv6 traffic, you must, at a minimum, configure the link-local addresses to pass traffic, but a global management address is recommended for full functionality, including remote management and other management operations.

For routed mode, if you provide a name for the BVI, then the BVI participates in routing. Without a name, the bridge group remains isolated as in transparent firewall mode.

Note For a separate interface, a non-configurable bridge group (ID 301) is automatically added to your configuration. This bridge group is not included in the bridge group limit.

Before you begin

You cannot add the BVI to a security zone; therefore, you cannot apply Access Control policies to the BVI. You must apply your policy to the bridge group member interfaces based on their zones.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Choose **Add Interfaces** > **Bridge Group Interface**.

Step 3 (Routed Mode) In the **Name** field, enter a name up to 48 characters in length.

You must name the BVI if you want to route traffic outside the bridge group members, for example, to the outside interface or to members of other bridge groups. The name is not case-sensitive.

Step 4 In the **Bridge Group ID** field, enter the bridge group ID between 1 and 250.

Step 5 In the **Description** field, enter a description for this bridge group.

Step 6 On the **Interfaces** tab, click an interface and then click **Add** to move it to the **Selected Interfaces** area. Repeat for all interfaces that you want to make members of the bridge group.

Step 7 (Transparent Mode) Click the **IPv4** tab. In the **IP Address** field, enter the IPv4 address and subnet mask.

Do not assign a host address (/32 or 255.255.255.255) to the BVI. Also, do not use other subnets that contain fewer than 3 host addresses (one each for the upstream router, downstream router, and transparent firewall) such as a /30 subnet (255.255.255.252). The Firepower Threat Defense device drops all ARP packets to or from the first and last addresses in a subnet. For example, if you use a /30 subnet and assign a reserved address from that subnet to the upstream router, then the Firepower Threat Defense device drops the ARP request from the downstream router to the upstream router.

For High Availabilty, set the standby IP address on the **Devices** > **Device Management** > **High Availability** tab in the **Monitored Interfaces** area. If you do not set the standby IP address, the active unit cannot monitor the standby interface using network tests; it can only track the link state.

Step 8 (Routed Mode) Click the **IPv4** tab. To set the IP address, use one of the following options from the **IP Type** drop-down list.

High Availability and clustering interfaces only support static IP address configuration; DHCP is not supported.

- **Use Static IP**—Enter the IP address and subnet mask. For High Availabilty, you can only use a static IP address. Set the standby IP address on the **Devices** > **Device Management** > **High Availability** tab in the **Monitored Interfaces** area. If you do not set the standby IP address, the active unit cannot monitor the standby interface using network tests; it can only track the link state.
- **Use DHCP**—Configure the following optional parameters:
  - **Obtain default route using DHCP**—Obtains the default route from the DHCP server.
  - **DHCP route metric**—Assigns an administrative distance to the learned route, between 1 and 255. The default administrative distance for the learned routes is 1.

Step 9 (Optional) See Configure IPv6 Addressing, on page 636 to configure IPv6 addressing.

Step 10 (Optional) See Add a Static ARP Entry, on page 647 and Add a Static MAC Address and Disable MAC Learning for a Bridge Group, on page 648 (for transparent mode only) to configure the **ARP** and **MAC** settings.

Step 11 Click **OK**.

Step 12 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure a Diagnostic (Management) Interface for Transparent Mode

In transparent firewall mode, all interfaces must belong to a bridge group. The only exception is the Diagnostic *slot/port* interface. For the Firepower 4100/9300 chassis, the diagnostic interface ID depends on the mgmt-type interface that you assigned to the Firepower Threat Defense logical device. You cannot use any other interface types as diagnostic interfaces. You can configure one diagnostic interface.

Before you begin

Do not assign this interface to a bridge group; a non-configurable bridge group (ID 301) is automatically added to your configuration. This bridge group is not included in the bridge group limit.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** () for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** () for the interface.

Step 3 In the **Name** field, enter a name up to 48 characters in length.

Step 4 Click the **IPv4** tab. To set the IP address, use one of the following options from the **IP Type** drop-down list.

- **Use Static IP**—Enter the IP address and subnet mask.
- **Use DHCP**—Configure the following optional parameters:
  - **Obtain default route using DHCP**—Obtains the default route from the DHCP server.
  - **DHCP route metric**—Assigns an administrative distance to the learned route, between 1 and 255. The default administrative distance for the learned routes is 1.
- **Use PPPoE**—Configure the following parameters:
  - **VPDN Group Name**—Specify a group name.
  - **PPPoE User Name**—Specify the username provided by your ISP.
  - **PPPoE Password/Confirm Password**—Specify and confirm the password provided by your ISP.
  - **PPP Authentication**—Choose **PAP**, **CHAP**, or **MSCHAP**.

    PAP passes a cleartext username and password during authentication and is not secure. With CHAP, the client returns the encrypted [challenge plus password], with a cleartext username in response to the server challenge. CHAP is more secure than PAP, but it does not encrypt data. MSCHAP is similar to CHAP but is more secure because the server stores and compares only encrypted passwords rather than cleartext passwords as in CHAP. MSCHAP also generates a key for data encryption by MPPE.
  - **PPPoE route metric**—Assign an administrative distance to the learned route. Valid values are from 1 to 255. By default, the administrative distance for the learned routes is 1.

- **Enable Route Settings**—To manually configure the PPPoE IP address, check this box and then enter the **IP Address**.
- **Store Username and Password in Flash**—Stores the username and password in flash memory.

  The Firepower Threat Defense device stores the username and password in a special location of NVRAM.

Step 5 (Optional) See Configure IPv6 Addressing, on page 636 to configure **IPv6** addressing.

Step 6 (Optional) On the **Advanced** tab, configure optional settings.

- See Configure the MAC Address, on page 646.
- See Add a Static ARP Entry, on page 647.
- See Set Security Configuration Parameters, on page 649.

Step 7 Click **OK**.

Step 8 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure IPv6 Addressing

This section describes how to configure IPv6 addressing in routed and transparent mode.

About IPv6

This section includes information about IPv6.

IPv6 Addressing

You can configure two types of unicast addresses for IPv6:

- Global—The global address is a public address that you can use on the public network. For a bridge group, this address needs to be configured for the BVI, and not per member interface. You can also configure a global IPv6 address for the management interface in transparent mode.
- Link-local—The link-local address is a private address that you can only use on the directly-connected network. Routers do not forward packets using link-local addresses; they are only for communication on a particular physical network segment. They can be used for address configuration or for the Neighbor Discovery functions such as address resolution. In a bridge group, only member interfaces have link-local addresses; the BVI does not have a link-local address.

At a minimum, you need to configure a link-local address for IPv6 to operate. If you configure a global address, a link-local address is automatically configured on the interface, so you do not also need to specifically configure a link-local address. For bridge group member interfaces, when you configure the global address on the BVI, the FTD device automatically generates link-local addresses for member interfaces. If you do not configure a global address, then you need to configure the link-local address, either automatically or manually.

Modified EUI-64 Interface IDs

RFC 3513: Internet Protocol Version 6 (IPv6) Addressing Architecture requires that the interface identifier portion of all unicast IPv6 addresses, except those that start with binary value 000, be 64 bits long and be constructed in Modified EUI-64 format. The FTD device can enforce this requirement for hosts attached to the local link.

When this feature is enabled on an interface, the source addresses of IPv6 packets received on that interface are verified against the source MAC addresses to ensure that the interface identifiers use the Modified EUI-64 format. If the IPv6 packets do not use the Modified EUI-64 format for the interface identifier, the packets are dropped and the following system log message is generated:

```
325003: EUI-64 source address check failed.
```

The address format verification is only performed when a flow is created. Packets from an existing flow are not checked. Additionally, the address verification can only be performed for hosts on the local link.

Configure a Global IPv6 Address

To configure a global IPv6 address for any routed mode interface and for the transparent or routed mode BVI, perform the following steps.

Note Configuring the global address automatically configures the link-local address, so you do not need to configure it separately. For bridge groups, configuring the global address on the BVI automatically configures link-local addresses on all member interfaces.

For subinterfaces defined on the FTD, we recommend that you also set the MAC address manually, because they use the same burned-in MAC address of the parent interface. IPv6 link-local addresses are generated based on the MAC address, so assigning unique MAC addresses to subinterfaces allows for unique IPv6 link-local addresses, which can avoid traffic disruption in certain instances on the FTD. See Configure the MAC Address, on page 646.

Before you begin

For IPv6 neighbor discovery for bridge groups, you must explicitly allow Neighbor Solicitation (ICMPv6 type 135) and Neighbor Advertisement (ICMPv6 type 136) packets through the FTD bridge group member interfaces using a bidirectional access rule.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

Step 3 Click the **IPv6** page.

For routed mode, the **Basic** page is selected by default. For transparent mode, the **Address** page is selected by default.

Step 4 On the **Basic** page, check **Enable IPv6**.

Step 5 Configure the global IPv6 address using one of the following methods.

- (Routed interface) Stateless autoconfiguration—Check the **Autoconfiguration** check box.

  Enabling stateless autconfiguration on the interface configures IPv6 addresses based upon prefixes received in Router Advertisement messages. A link-local address, based on the Modified EUI-64 interface ID, is automatically generated for the interface when stateless autoconfiguration is enabled.

  Although RFC 4862 specifies that hosts configured for stateless autoconfiguration do not send Router Advertisement messages, the Firepower Threat Defense device does send Router Advertisement messages in this case. Uncheck the **IPv6** > **Settings** > **Enable RA** check box to suppress messages.

- Manual configuration—To manually configure a global IPv6 address:

  a. Click the **Address** page, and click **Add Address**.

  The **Add Address** dialog box appears.

  b. In the **Address** field, enter either a full global IPv6 address, including the interface ID, or enter the IPv6 prefix, along with the IPv6 prefix length. (Routed Mode) If you only enter the prefix, then be sure to check the **Enforce EUI 64** check box to generate the interface ID using the Modified EUI-64 format. For example, 2001:0DB8::BA98:0:3210/48 (full address) or 2001:0DB8::/48 (prefix, with EUI 64 checked).

  For High Availabilty (if you did not set **Enforce EUI 64**), set the standby IP address on the **Devices** > **Device Management** > **High Availability** page in the **Monitored Interfaces** area. If you do not set the standby IP address, the active unit cannot monitor the standby interface using network tests; it can only track the link state.

Step 6 For Routed interfaces, you can optionally set the following values on the **Basic** page:

- To enforce the use of Modified EUI-64 format interface identifiers in IPv6 addresses on a local link, check the **Enforce EUI-64** check box.
- To manually set the link-local address, enter an address in the **Link-Local address** field.

  A link-local address should start with FE8, FE9, FEA, or FEB, for example fe80::20d:88ff:feee:6a82. If you do not want to configure a global address, and only need to configure a link-local address, you have the option of manually defining the link-local address. Note that we recommend automatically assigning the link-local address based on the Modified EUI-64 format. For example, if other devices enforce the use of the Modified EUI-64 format, then a manually-assigned link-local address may cause packets to be dropped.

- Check the **Enable DHCP for address config** check box to set the Managed Address Config flag in the IPv6 router advertisement packet.

  This flag in IPv6 router advertisements informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain addresses, in addition to the derived stateless autoconfiguration address.

- Check the **Enable DHCP for non-address config** check box to set the Other Address Config flag in the IPv6 router advertisement packet.

  This flag in IPv6 router advertisements informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain additional information from DHCPv6, such as the DNS server address.

Step 7 For Routed interfaces, see Configure IPv6 Neighbor Discovery, on page 639 to configure settings on the **Prefixes** and **Settings** pages. For BVI interfaces, see the following parameters on the **Settings** page:

- **DAD attempts**—The maximum number of DAD attempts, between 1 and 600. Set the value to 0 to disable duplicate address detection (DAD) processing. This setting configures the number of consecutive neighbor solicitation messages that are sent on an interface while DAD is performed on IPv6 addresses. 1 attempt is the default.
- **NS Interval**—The interval between IPv6 neighbor solicitation retransmissions on an interface, between 1000 and 3600000 ms. The default value is 1000 ms.
- **Reachable Time**—The amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event has occurred, between 0 and 3600000 ms. The default value is 0 ms. When 0 is used for the value, the reachable time is sent as undetermined. It is up to the receiving devices to set and track the reachable time value. The neighbor reachable time enables detecting unavailable neighbors. Shorter configured times enable detecting unavailable neighbors more quickly, however, shorter times consume more IPv6 network bandwidth and processing resources in all IPv6 network devices. Very short configured times are not recommended in normal IPv6 operation.

Step 8 Click **OK**.

Step 9 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure IPv6 Neighbor Discovery

The IPv6 neighbor discovery process uses ICMPv6 messages and solicited-node multicast addresses to determine the link-layer address of a neighbor on the same network (local link), verify the readability of a neighbor, and keep track of neighboring routers.

Nodes (hosts) use neighbor discovery to determine the link-layer addresses for neighbors known to reside on attached links and to quickly purge cached values that become invalid. Hosts also use neighbor discovery to find neighboring routers that are willing to forward packets on their behalf. In addition, nodes use the protocol to actively keep track of which neighbors are reachable and which are not, and to detect changed link-layer addresses. When a router or the path to a router fails, a host actively searches for functioning alternates.

Before you begin

Supported in Routed mode only. For IPv6 neighbor settings supported in transparent mode, see Configure a Global IPv6 Address, on page 637.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

Step 3 Click **IPv6**, and then **Prefixes**.

Step 4 (Optional) To configure which IPv6 prefixes are included in IPv6 router advertisements, perform the following steps:

a) Click **Add Prefix**.
b) In the **Address** field, enter the IPv6 address with the prefix length or check the **Default** check box to use the default prefix.
c) (Optional) Uncheck the **Advertisement** check box to indicate that the IPv6 prefix is not advertised.
d) Check the **Off Link** check box to indicate that the specified prefix is assigned to the link. Nodes sending traffic to addresses that contain the specified prefix consider the destination to be locally reachable on the link. This prefix should not be used for on-link determination.
e) To use the specified prefix for autoconfiguration, check the **Autoconfiguration** check box.
f) For the **Prefix Lifetime**, click **Duration** or **Expiration Date**.

- **Duration**—Enter a **Preferred Lifetime** for the prefix in seconds. This setting is the amount of time that the specified IPv6 prefix is advertised as being valid. The maximum value represents infinity. Valid values are from 0 to 4294967295. The default is 2592000 (30 days). Enter a **Valid Lifetime** for the prefix in seconds. This setting is the amount of time that the specified IPv6 prefix is advertised as being preferred. The maximum value represents infinity. Valid values are from 0 to 4294967295. The default setting is 604800 (seven days). Alternatively, check the **Infinite** checkbox to set an unlimited duration.
- **Expiration Date**—Choose a **Valid** and **Preferred** date and time.

g) Click **OK**.

Step 5 Click **Settings**.

Step 6 (Optional) Set the maximum number of **DAD attempts**, between 1 and 600. 1 attempt is the default. Set the value to 0 to disable duplicate address detection (DAD) processing.

This setting configures the number of consecutive neighbor solicitation messages that are sent on an interface while DAD is performed on IPv6 addresses.

During the stateless autoconfiguration process, Duplicate Address Detection verifies the uniqueness of new unicast IPv6 addresses before the addresses are assigned to interfaces.

When a duplicate address is identified, the state of the address is set to DUPLICATE, the address is not used, and the following error message is generated:

```
325002: Duplicate address ipv6_address/MAC_address on interface
```

If the duplicate address is the link-local address of the interface, the processing of IPv6 packets is disabled on the interface. If the duplicate address is a global address, the address is not used.

Step 7 (Optional) Configure the interval between IPv6 neighbor solicitation retransmissions in the **NS Interval** field, between 1000 and 3600000 ms.

The default value is 1000 ms.

Neighbor solicitation messages (ICMPv6 Type 135) are sent on the local link by nodes attempting to discover the link-layer addresses of other nodes on the local link. After receiving a neighbor solicitation message, the destination node replies by sending a neighbor advertisement message (ICPMv6 Type 136) on the local link.

After the source node receives the neighbor advertisement, the source node and destination node can communicate. Neighbor solicitation messages are also used to verify the reachability of a neighbor after the

link-layer address of a neighbor is identified. When a node wants to verifying the reachability of a neighbor, the destination address in a neighbor solicitation message is the unicast address of the neighbor.

Neighbor advertisement messages are also sent when there is a change in the link-layer address of a node on a local link.

Step 8 (Optional) Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event has occurred in the **Reachable Time** field, between 0 and 3600000 ms.

The default value is 0 ms. When 0 is used for the value, the reachable time is sent as undetermined. It is up to the receiving devices to set and track the reachable time value.

The neighbor reachable time enables detecting unavailable neighbors. Shorter configured times enable detecting unavailable neighbors more quickly, however, shorter times consume more IPv6 network bandwidth and processing resources in all IPv6 network devices. Very short configured times are not recommended in normal IPv6 operation.

Step 9 (Optional) To suppress the router advertisement transmissions, uncheck the **Enable RA** check box. If you enable router advertisement transmissions, you can set the RA lifetime and interval.

Router advertisement messages (ICMPv6 Type 134) are automatically sent in response to router solicitation messages (ICMPv6 Type 133). Router solicitation messages are sent by hosts at system startup so that the host can immediately autoconfigure without needing to wait for the next scheduled router advertisement message.

You may want to disable these messages on any interface for which you do not want the Firepower Threat Defense device to supply the IPv6 prefix (for example, the outside interface).

- **RA Lifetime**—Configure the router lifetime value in IPv6 router advertisements, between 0 and 9000 seconds.

  The default is 1800 seconds.

- **RA Interval**—Configure the interval between IPv6 router advertisement transmissions, between 3 and 1800 seconds.

  The default is 200 seconds.

Step 10 Click **OK**.

Step 11 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure Advanced Interface Settings

This section describes how to configure MAC addresses for regular firewall mode interfaces, how to set the maximum transmission unit (MTU), and how to set other advanced parameters.

About Advanced Interface Configuration

This section describes advanced interface settings.

About MAC Addresses

You can manually assign MAC addresses to override the default.



Note You might want to assign unique MAC addresses to subinterfaces defined on the FTD, because they use the same burned-in MAC address of the parent interface. For example, your service provider might perform access control based on the MAC address. Also, because IPv6 link-local addresses are generated based on the MAC address, assigning unique MAC addresses to subinterfaces allows for unique IPv6 link-local addresses, which can avoid traffic disruption in certain instances on the FTD device.

Default MAC Addresses

Default MAC address assignments depend on the type of interface.

- Physical interfaces—The physical interface uses the burned-in MAC address.
- EtherChannels (Firepower Models)—For an EtherChannel, all interfaces that are part of the channel group share the same MAC address. This feature makes the EtherChannel transparent to network applications and users, because they only see the one logical connection; they have no knowledge of the individual links. The port-channel interface uses a unique MAC address from a pool; interface membership does not affect the MAC address.
- EtherChannels (ASA Models)—The port-channel interface uses the lowest-numbered channel group interface MAC address as the port-channel MAC address. Alternatively you can configure a MAC address for the port-channel interface. We recommend configuring a unique MAC address in case the group channel interface membership changes. If you remove the interface that was providing the port-channel MAC address, then the port-channel MAC address changes to the next lowest numbered interface, thus causing traffic disruption.
- Subinterfaces—All subinterfaces of a physical interface use the same burned-in MAC address. You might want to assign unique MAC addresses to subinterfaces. For example, your service provider might perform access control based on the MAC address. Also, because IPv6 link-local addresses are generated based on the MAC address, assigning unique MAC addresses to subinterfaces allows for unique IPv6 link-local addresses, which can avoid traffic disruption in certain instances on the FTD.

About the MTU

The MTU specifies the maximum frame *payload* size that the FTD device can transmit on a given Ethernet interface. The MTU value is the frame size *without* Ethernet headers, VLAN tagging, or other overhead. For example, when you set the MTU to 1500, the expected frame size is 1518 bytes including the headers, or 1522 when using VLAN. Do not set the MTU value higher to accommodate these headers.

Path MTU Discovery

The FTD device supports Path MTU Discovery (as defined in RFC 1191), which lets all devices in a network path between two hosts coordinate the MTU so they can standardize on the lowest MTU in the path.

Default MTU

The default MTU on the FTD device is 1500 bytes. This value does not include the 18-22 bytes for the Ethernet header, VLAN tagging, or other overhead.

MTU and Fragmentation

For IPv4, if an outgoing IP packet is larger than the specified MTU, it is fragmented into 2 or more frames. Fragments are reassembled at the destination (and sometimes at intermediate hops), and fragmentation can cause performance degradation. For IPv6, packets are typically not allowed to be fragmented at all. Therefore, your IP packets should fit within the MTU size to avoid fragmentation.

For TCP packets, the endpoints typically use their MTU to determine the TCP maximum segment size (MTU - 40, for example). If additional TCP headers are added along the way, for example for site-to-site VPN tunnels, then the TCP MSS might need to be adjusted down by the tunneling entity. See About the TCP MSS, on page 643.

For UDP or ICMP, the application should take the MTU into account to avoid fragmentation.

Note The FTD device can receive frames larger than the configured MTU as long as there is room in memory.

MTU and Jumbo Frames

A larger MTU lets you send larger packets. Larger packets might be more efficient for your network. See the following guidelines:

- Matching MTUs on the traffic path—We recommend that you set the MTU on all FTD interfaces and other device interfaces along the traffic path to be the same. Matching MTUs prevents intermediate devices from fragmenting the packets.
- Accommodating jumbo frames—You can set the MTU 9000 bytes or higher when you enable jumbo frames. The maximum depends on the model.

About the TCP MSS

The TCP maximum segment size (MSS) is the size of the TCP payload *before* any TCP and IP headers are added. UDP packets are not affected. The client and the server exchange TCP MSS values during the three-way handshake when establishing the connection.

You can set the TCP MSS on the FTD device for through traffic using the Sysopt_Basic object in FlexConfig; see FlexConfig Policies for FTD, on page 575; by default, the maximum TCP MSS is set to 1380 bytes. This setting is useful when the FTD device needs to add to the size of the packet for IPsec VPN encapsulation. However, for non-IPsec endpoints, you should disable the maximum TCP MSS on the FTD device.

If you set a maximum TCP MSS, if either endpoint of a connection requests a TCP MSS that is larger than the value set on the FTD device, then the FTD device overwrites the TCP MSS in the request packet with the FTD device maximum. If the host or server does not request a TCP MSS, then the FTD device assumes the RFC 793-default value of 536 bytes (IPv4) or 1220 bytes (IPv6), but does not modify the packet. For example, you leave the default MTU as 1500 bytes. A host requests an MSS of 1500 minus the TCP and IP header length, which sets the MSS to 1460. If the FTD device maximum TCP MSS is 1380 (the default), then the FTD device changes the MSS value in the TCP request packet to 1380. The server then sends packets with 1380-byte payloads. The FTD device can then add up to 120 bytes of headers to the packet and still fit in the MTU size of 1500.

You can also configure the minimum TCP MSS; if a host or server requests a very small TCP MSS, the FTD device can adjust the value up. By default, the minimum TCP MSS is not enabled.

For to-the-box traffic, including for SSL VPN connections, this setting does not apply. The FTD device uses the MTU to derive the TCP MSS: MTU - 40 (IPv4) or MTU - 60 (IPv6).

Default TCP MSS

By default, the maximum TCP MSS on the FTD device is 1380 bytes. This default accommodates IPv4 IPsec VPN connections where the headers can equal up to 120 bytes; this value fits within the default MTU of 1500 bytes.

Suggested Maximum TCP MSS Setting

The default TCP MSS assumes the FTD device acts as an IPv4 IPsec VPN endpoint and has an MTU of 1500. When the FTD device acts as an IPv4 IPsec VPN endpoint, it needs to accommodate up to 120 bytes for TCP and IP headers.

If you change the MTU value, use IPv6, or do not use the FTD device as an IPsec VPN endpoint, then you should change the TCP MSS setting using the Sysopt_Basic object in FlexConfig.



Note Even if you explicitly set an MSS, if a component such as TLS/SSL decryption or server discovery needs a particular MSS, it will set that MSS based on the interface MTU and ignore your MSS setting.

See the following guidelines:

- Normal traffic—Disable the TCP MSS limit and accept the value established between connection endpoints. Because connection endpoints typically derive the TCP MSS from the MTU, non-IPsec packets usually fit this TCP MSS.
- IPv4 IPsec endpoint traffic—Set the maximum TCP MSS to the MTU - 120. For example, if you use jumbo frames and set the MTU to 9000, then you need to set the TCP MSS to 8880 to take advantage of the new MTU.
- IPv6 IPsec endpoint traffic—Set the maximum TCP MSS to the MTU - 140.

ARP Inspection for Bridge Group Traffic

By default, all ARP packets are allowed between bridge group members. You can control the flow of ARP packets by enabling ARP inspection.

ARP inspection prevents malicious users from impersonating other hosts or routers (known as ARP spoofing). ARP spoofing can enable a “man-in-the-middle” attack. For example, a host sends an ARP request to the gateway router; the gateway router responds with the gateway router MAC address. The attacker, however, sends another ARP response to the host with the attacker MAC address instead of the router MAC address. The attacker can now intercept all the host traffic before forwarding it on to the router.

ARP inspection ensures that an attacker cannot send an ARP response with the attacker MAC address, so long as the correct MAC address and the associated IP address are in the static ARP table.

When you enable ARP inspection, the FTD device compares the MAC address, IP address, and source interface in all ARP packets to static entries in the ARP table, and takes the following actions:

- If the IP address, MAC address, and source interface match an ARP entry, the packet is passed through.
- If there is a mismatch between the MAC address, the IP address, or the interface, then the FTD device drops the packet.

- If the ARP packet does not match any entries in the static ARP table, then you can set the FTD device to either forward the packet out all interfaces (flood), or to drop the packet.

Note The dedicated interface never floods packets even if this parameter is set to flood.

MAC Address Table

When you use bridge groups, the FTD learns and builds a MAC address table in a similar way as a normal bridge or switch: when a device sends a packet through the bridge group, the FTD adds the MAC address to its table. The table associates the MAC address with the source interface so that the FTD knows to send any packets addressed to the device out the correct interface. Because traffic between bridge group members is subject to the FTD security policy, if the destination MAC address of a packet is not in the table, the FTD does not flood the original packet on all interfaces as a normal bridge does. Instead, it generates the following packets for directly-connected devices or for remote devices:

- Packets for directly-connected devices—The FTD generates an ARP request for the destination IP address, so that it can learn which interface receives the ARP response.
- Packets for remote devices—The FTD generates a ping to the destination IP address so that it can learn which interface receives the ping reply.

The original packet is dropped.

Default Settings

- If you enable ARP inspection, the default setting is to flood non-matching packets.
- The default timeout value for dynamic MAC address table entries is 5 minutes.
- By default, each interface automatically learns the MAC addresses of entering traffic, and the FTD device adds corresponding entries to the MAC address table.

Note Firepower Threat Defense device generates a reset packet to reset a connection that is denied by a stateful inspection engine. Here, the destination MAC address of the packet is not determined based on the ARP table lookup but instead it is taken directly from the packets (connections) that are being denied.

Guidelines for ARP Inspection and the MAC Address Table

- ARP inspection is only supported for bridge groups.
- MAC address table configuration is only supported for bridge groups.

Configure the MTU

Customize the MTU on the interface, for example, to allow jumbo frames.

Caution Changing the highest MTU value on the device for a data interface restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all data interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. This caution does not apply to the Diagnostic interface or management-only interfaces. See Snort® Restart Traffic Behavior, on page 350 for more information.

Before you begin

- Changing the MTU above 1500 bytes automatically enables jumbo frames; for ASA models, you must reload the system before you can use jumbo frames.
- If you use an interface in an inline set, the MTU setting is not used. However, the jumbo frame setting *is* relevant to inline sets; jumbo frames enable the inline interfaces to receive packets up to 9000 bytes. To enable jumbo frames, you must set the MTU of *any* interface above 1500 bytes.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

Step 3 On the **General** tab, set the **MTU** between 64 and 9198 bytes; the maximum is 9000 for the Firepower Threat Defense Virtual and 9184 for the Firepower Threat Defense on the Firepower 4100/9300 chassis.

The default is 1500 bytes.

Step 4 Click **OK**.

Step 5 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Step 6 For ASA models, if you set the MTU above 1500 bytes, reload the system to enable jumbo frames.

Configure the MAC Address

You might need to manually assign a MAC address. You can also set the Active and Standby MAC addresses on the **Devices** > **Device Management** > **High Availability** tab. If you set the MAC address for an interface on both screens, the addresses on the **Interfaces** > **Advanced** tab take precedence.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

Step 3 Click the **Advanced** tab.
The **Information** tab is selected.

Step 4 In the **Active MAC Address** field, enter a MAC address in H.H.H format, where H is a 16-bit hexadecimal digit.

For example, the MAC address 00-0C-F1-42-4C-DE would be entered as 000C.F142.4CDE. The MAC address must not have the multicast bit set, that is, the second hexadecimal digit from the left cannot be an odd number.

Step 5 In the **Standby MAC Address** field, enter a MAC address for use with High Availability.

If the active unit fails over and the standby unit becomes active, the new active unit starts using the active MAC addresses to minimize network disruption, while the old active unit uses the standby address.

Step 6 Click **OK**.

Step 7 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Add a Static ARP Entry

By default, all ARP packets are allowed between bridge group members. You can control the flow of ARP packets by enabling ARP inspection (see Configure ARP Inspection, on page 949). ARP inspection compares ARP packets with *static* ARP entries in the ARP table.

For routed interfaces, you can enter static ARP entries, but normally dynamic entries are sufficient. For routed interfaces, the ARP table is used to deliver packets to directly-connected hosts. Although senders identify a packet destination by an IP address, the actual delivery of the packet on Ethernet relies on the Ethernet MAC address. When a router or host wants to deliver a packet on a directly connected network, it sends an ARP request asking for the MAC address associated with the IP address, and then delivers the packet to the MAC address according to the ARP response. The host or router keeps an ARP table so it does not have to send ARP requests for every packet it needs to deliver. The ARP table is dynamically updated whenever ARP responses are sent on the network, and if an entry is not used for a period of time, it times out. If an entry is incorrect (for example, the MAC address changes for a given IP address), the entry needs to time out before it can be updated with the new information.

For transparent mode, the Firepower Threat Defense only uses dynamic ARP entries in the ARP table for traffic to and from the Firepower Threat Defense device, such as management traffic.

Before you begin

This screen is only available for named interfaces.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

Step 3 Click the **Advanced** tab, and then click the **ARP** tab (called **ARP and MAC** for transparent mode).

Step 4 Click **Add ARP Config**.
The **Add ARP Config** dialog box appears.

Step 5 In the **IP Address** field, enter the IP address of the host.

Step 6 In the **MAC Address** field, enter the MAC address of the host; for example, 00e0.1e4e.3d8b.

Step 7 To perform proxy ARP for this address, check the **Enable Alias** check box.

If the Firepower Threat Defense device receives an ARP request for the specified IP address, then it responds with the specified MAC address.

Step 8 Click **OK**, and then click **OK** again to exit the Advanced settings.

Step 9 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Add a Static MAC Address and Disable MAC Learning for a Bridge Group

Normally, MAC addresses are added to the MAC address table dynamically as traffic from a particular MAC address enters an interface. You can disable MAC address learning; however, unless you statically add MAC addresses to the table, no traffic can pass through the Firepower Threat Defense device. You can also add static MAC addresses to the MAC address table. One benefit to adding static entries is to guard against MAC spoofing. If a client with the same MAC address as a static entry attempts to send traffic to an interface that does not match the static entry, then the Firepower Threat Defense device drops the traffic and generates a system message. When you add a static ARP entry (see Add a Static ARP Entry, on page 647), a static MAC address entry is automatically added to the MAC address table.

Before you begin

This screen is only available for named interfaces.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

Step 3 Click the **Advanced** tab, and then click the **ARP and MAC** tab.

Step 4 (Optional) Disable MAC learning by unchecking the **Enable MAC Learning** check box.

Step 5 To add a static MAC address, click **Add MAC Config**.
The **Add MAC Config** dialog box appears.

Step 6 In the **MAC Address** field, enter the MAC address of the host; for example, 00e0.1e4e.3d8b. Click **OK**.

Step 7 Click **OK** to exit the Advanced settings.

Step 8 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Set Security Configuration Parameters

This section describes how to prevent IP spoofing, allow full fragment reassembly, and override the default fragment setting set for at the device level in **Platform Settings** .

Anti-Spoofing

This section lets you enable Unicast Reverse Path Forwarding on an interface. Unicast RPF guards against IP spoofing (a packet uses an incorrect source IP address to obscure its true source) by ensuring that all packets have a source IP address that matches the correct source interface according to the routing table.

Normally, the Firepower Threat Defense device only looks at the destination address when determining where to forward the packet. Unicast RPF instructs the device to also look at the source address; this is why it is called Reverse Path Forwarding. For any traffic that you want to allow through the Firepower Threat Defense device, the device routing table must include a route back to the source address. See RFC 2267 for more information.

For outside traffic, for example, the Firepower Threat Defense device can use the default route to satisfy the Unicast RPF protection. If traffic enters from an outside interface, and the source address is not known to the routing table, the device uses the default route to correctly identify the outside interface as the source interface.

If traffic enters the outside interface from an address that is known to the routing table, but is associated with the inside interface, then the Firepower Threat Defense device drops the packet. Similarly, if traffic enters the inside interface from an unknown source address, the device drops the packet because the matching route (the default route) indicates the outside interface.

Unicast RPF is implemented as follows:

- ICMP packets have no session, so each packet is checked.
- UDP and TCP have sessions, so the initial packet requires a reverse route lookup. Subsequent packets arriving during the session are checked using an existing state maintained as part of the session. Non-initial packets are checked to ensure they arrived on the same interface used by the initial packet.

Fragment per Packet

By default, the Firepower Threat Defense device allows up to 24 fragments per IP packet, and up to 200 fragments awaiting reassembly. You might need to let fragments on your network if you have an application that routinely fragments packets, such as NFS over UDP. However, if you do not have an application that fragments traffic, we recommend that you do not allow fragments through the Firepower Threat Defense device. Fragmented packets are often used as DoS attacks.

Fragment Reassembly

The Firepower Threat Defense device performs the following fragment reassembly processes:

- IP fragments are collected until a fragment set is formed or until a timeout interval has elapsed.
- If a fragment set is formed, integrity checks are performed on the set. These checks include no overlapping, no tail overflow, and no chain overflow.
- IP fragments that terminate at the Firepower Threat Defense device are always fully reassembled.
- If **Full Fragment Reassembly** is disabled (the default), the fragment set is forwarded to the transport layer for further processing.
- If **Full Fragment Reassembly** is enabled, the fragment set is first coalesced into a single IP packet. The single IP packet is then forwarded to the transport layer for further processing.

Before you begin

This screen is only available for named interfaces.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

Step 3 Click the **Advanced** tab, and then click the **Security Configuration** tab.

Step 4 To enable Unicast Reverse Path Forwarding, check the **Anti-Spoofing** check box.

Step 5 To enable full fragment reassembly, check the **Full Fragment Reassembly** check box.

Step 6 To change the number of fragments allowed per packet, check the **Override Default Fragment Setting** check box, and set the following values:

- **Size**—Set the maximum number of packets that can be in the IP reassembly database waiting for reassembly. The default is 200. Set this value to 1 to disable fragments.
- **Chain**—Set the maximum number of packets into which a full IP packet can be fragmented. The default is 24 packets.
- **Timeout**—Set the maximum number of seconds to wait for an entire fragmented packet to arrive. The timer starts after the first fragment of a packet arrives. If all fragments of the packet do not arrive by the number of seconds specified, all fragments of the packet that were already received will be discarded. The default is 5 seconds.

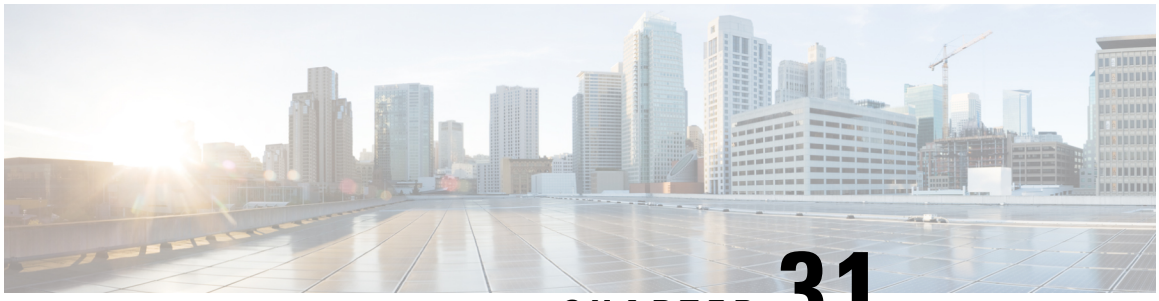
Step 7 Click **OK**.

Step 8 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

History for Regular Firewall Interfaces for Firepower Threat Defense

Feature	Version	Details
Integrated Routing and Bridging	6.2.0	Integrated Routing and Bridging provides the ability to route between a bridge group and a routed interface. A bridge group is a group of interfaces that the Firepower Threat Defense bridges instead of routes. The Firepower Threat Defense is not a true bridge in that the Firepower Threat Defense continues to act as a firewall: access control between interfaces is controlled, and all of the usual firewall checks are in place. Previously, you could only configure bridge groups in transparent firewall mode, where you cannot route between bridge groups. This feature lets you configure bridge groups in routed firewall mode, and to route between bridge groups and between a bridge group and a routed interface. The bridge group participates in routing by using a Bridge Virtual Interface (BVI) to act as a gateway for the bridge group. Integrated Routing and Bridging provides an alternative to using an external Layer 2 switch if you have extra interfaces on the Firepower Threat Defense to assign to the bridge group. In routed mode, the BVI can be a named interface and can participate separately from member interfaces in some features, such as access rules and DHCP server. The following features that are supported in transparent mode are not supported in routed mode: clustering. The following features are also not supported on BVIs: dynamic routing and multicast routing. New/Modified screens: • **Devices** > **Device Management** > **Interfaces** > **Edit Physical Interface** • **Devices** > **Device Management** > **Interfaces** > **Add Interfaces** > **Bridge Group Interface** Supported platforms: All except for the Firepower 2100 and the Firepower Threat Defense Virtual



CHAPTER 31

Inline Sets and Passive Interfaces for Firepower Threat Defense

You can configure IPS-only passive interfaces, passive ERSPAN interfaces, and inline sets. IPS-only mode interfaces bypass many firewall checks and only support IPS security policy. You might want to implement IPS-only interfaces if you have a separate firewall protecting these interfaces and do not want the overhead of firewall functions.

- About IPS Interfaces, on page 653
- Requirements and Prerequisites for Inline Sets, on page 655
- Guidelines for Inline Sets and Passive Interfaces, on page 656
- Configure a Passive Interface, on page 657
- Configure an Inline Set, on page 659
- History for Inline Sets and Passive Interfaces for Firepower Threat Defense, on page 662

About IPS Interfaces

This section describes IPS interfaces.

IPS Interface Types

IPS-only mode interfaces bypass many firewall checks and only support IPS security policy. You might want to implement IPS-only interfaces if you have a separate firewall protecting these interfaces and do not want the overhead of firewall functions.

Note The firewall mode only affects regular firewall interfaces, and not IPS-only interfaces such as inline sets or passive interfaces. IPS-only interfaces can be used in both firewall modes.

IPS-only interfaces can be deployed as the following types:

- Inline Set, with optional Tap mode—An inline set acts like a bump on the wire, and binds two interfaces together to slot into an existing network. This function allows the FTD to be installed in any network environment without the configuration of adjacent network devices. Inline interfaces receive all traffic unconditionally, but all traffic received on these interfaces is retransmitted out of an inline set unless explicitly dropped.

With tap mode, the FTD is deployed inline, but the network traffic flow is undisturbed. Instead, the FTD makes a copy of each packet so that it can analyze the packets. Note that rules of these types do generate intrusion events when they are triggered, and the table view of intrusion events indicates that the triggering packets would have dropped in an inline deployment. There are benefits to using tap mode with FTDs that are deployed inline. For example, you can set up the cabling between the FTD and the network as if the FTD were inline and analyze the kinds of intrusion events the FTD generates. Based on the results, you can modify your intrusion policy and add the drop rules that best protect your network without impacting its efficiency. When you are ready to deploy the FTD inline, you can disable tap mode and begin dropping suspicious traffic without having to reconfigure the cabling between the FTD and the network.

Note Tap mode *significantly* impacts FTD performance, depending on the traffic.

Note Inline sets might be familiar to you as "transparent inline sets," but the inline interface type is unrelated to the transparent firewall mode or the firewall-type interfaces.

- Passive or ERSPAN Passive—Passive interfaces monitor traffic flowing across a network using a switch SPAN or mirror port. The SPAN or mirror port allows for traffic to be copied from other ports on the switch. This function provides the system visibility within the network without being in the flow of network traffic. When you configure the FTD in a passive deployment, the FTD cannot take certain actions such as blocking or shaping traffic. Passive interfaces receive all traffic unconditionally. and no traffic received on these interfaces is retransmitted. Encapsulated remote switched port analyzer (ERSPAN) interfaces allow you to monitor traffic from source ports distributed over multiple switches, and uses GRE to encapsulate the traffic. ERSPAN interfaces are only allowed when the FTD is in routed firewall mode.

Note Using SR-IOV interfaces as passive interfaces on NGFWv is not supported on some Intel network adapters (such as Intel X710 or 82599) using SR-IOV drivers due to a promiscuous mode restriction. In such cases, use a network adapter that supports this functionality. See Intel Ethernet Products for more information on Intel network adapters.

About Hardware Bypass for Inline Sets

For certain interface modules on the Firepower 9300 and 4100 series (see Requirements and Prerequisites for Inline Sets, on page 655), you can enable the Hardware Bypass feature. Hardware Bypass ensures that traffic continues to flow between an inline interface pair during a power outage. This feature can be used to maintain network connectivity in the case of software or hardware failures.

Hardware Bypass Triggers

Hardware Bypass can be triggered in the following scenarios:

- Firepower Threat Defense application crash
- Firepower Threat Defense application reboot
- Security Module reboot
- Firepower chassis crash
- Firepower chassis reboot or upgrade
- Manual trigger
- Firepower chassis power loss
- Security Module power loss

Note Hardware bypass is intended for unplanned/unexpected failure scenarios, and is not automatically triggered during planned software upgrades. Hardware bypass only engages at the end of a planned upgrade process, when the FTD application reboots.

Hardware Bypass Switchover

When switching from normal operation to hardware bypass or from hardware bypass back to normal operation, traffic may be interrupted for several seconds. A number of factors can affect the length of the interruption; for example, copper port auto-negotiation; behavior of the optical link partner such as how it handles link faults and de-bounce timing; spanning tree protocol convergence; dynamic routing protocol convergence; and so on. During this time, you may experience dropped connections.

You may also experience dropped connections due to application identification errors when analyzing connections midstream after the return to normal operations.

Snort Fail Open vs. Hardware Bypass

For inline sets other than those in tap mode, you can use the Snort Fail Open option to either drop traffic or allow traffic to pass without inspection when the Snort process is busy or down. Snort Fail Open is supported on all inline sets except those in tap mode, not just on interfaces that support Hardware Bypass.

The Hardware Bypass functionality allows traffic to flow during a hardware failure, including a complete power outage, and certain limited software failures. A software failure that triggers Snort Fail Open does not trigger a Hardware Bypass.

Hardware Bypass Status

If the system has power, then the Bypass LED indicates the Hardware Bypass status. See the Firepower chassis hardware installation guide for LED descriptions.

Requirements and Prerequisites for Inline Sets

Model Support

FTD

User Roles

- Admin
- Access Admin
- Network Admin

Hardware Bypass Support

The Firepower Threat Defense supports Hardware Bypass for interface pairs on specific network modules on the following models:

- Firepower 9300
- Firepower 4100 series

The supported Hardware Bypass network modules for these models include:

- Firepower 6-port 1G SX FTW Network Module single-wide (FPR-NM-6X1SX-F)
- Firepower 6-port 10G SR FTW Network Module single-wide (FPR-NM-6X10SR-F)
- Firepower 6-port 10G LR FTW Network Module single-wide (FPR-NM-6X10LR-F)
- Firepower 2-port 40G SR FTW Network Module single-wide (FPR-NM-2X40G-F)
- Firepower 8-port 1G Copper FTW Network Module single-wide (FPR-NM-8X1G-F)

Hardware Bypass can only use the following port pairs:

- 1 & 2
- 3 & 4
- 5 & 6
- 7 & 8

Guidelines for Inline Sets and Passive Interfaces

Firewall Mode

- ERSPAN interfaces are only allowed when the device is in routed firewall mode.
- Firepower Threat Defense cannot load balance GRE traffic to multiple Rx rings. Flow-offload feature does not work for ERSPAN traffic. Hence, high rate of ERSPAN traffic can cause packet drop and impact the Firepower Threat Defense performance.

General Guidelines

- Inline sets and passive interfaces support physical interfaces and EtherChannels only, and cannot use redundant interfaces, VLANs, and so on.
- Inline sets and passive interfaces are supported in intra-chassis and inter-chassis clustering.

- Bidirectional Forwarding Detection (BFD) echo packets are not allowed through the Firepower Threat Defense when using inline sets. If there are two neighbors on either side of the Firepower Threat Defense running BFD, then the Firepower Threat Defense will drop BFD echo packets because they have the same source and destination IP address and appear to be part of a LAND attack.
- For inline sets and passive interfaces, the FTD supports up to two 802.1Q headers in a packet (also known as Q-in-Q support), with the exception of the Firepower 4100/9300, which only supports one 802.1Q header. **Note:** Firewall-type interfaces do not support Q-in-Q, and only support one 802.1Q header.

Hardware Bypass Guidelines

- Hardware Bypass ports are supported only for inline sets.
- Hardware Bypass ports cannot be part of an EtherChannel.
- Supported with intra-chassis clustering. Ports are placed in Hardware Bypass mode when the last unit in the chassis fails. Inter-chassis clustering is not supported.
- If all units in the cluster fail, then Hardware Bypass is triggered on the final unit, and traffic continues to pass. When units come back up, Hardware Bypass returns to standby mode. However, when you use rules that match application traffic, those connections may be dropped and need to be reestablished. Connections are dropped because state information is not retained on the cluster unit, and the unit cannot identify the traffic as belonging to an allowed application. To avoid a traffic drop, use a port-based rule instead of an application-based rule, if appropriate for your deployment.
- Hardware Bypass is not supported in high availability mode.

Unsupported Firewall Features on IPS Interfaces

- DHCP server
- DHCP relay
- DHCP client
- TCP Intercept
- Routing
- NAT
- VPN
- Application inspection
- QoS
- NetFlow
- VXLAN

Configure a Passive Interface

This section describes how to:

- Enable the interface. By default, interfaces are disabled.
- Set the interface mode to Passive or ERSPAN. For ERSPAN interfaces, you will set the ERSPAN parameters and the IP address.
- Change the MTU. By default, the MTU is set to 1500 bytes. For more information about the MTU, see About the MTU, on page 642.
- Set a specific speed and duplex (if available). By default, speed and duplex are set to Auto.

Note For the Firepower Threat Defense on the FXOS chassis, you configure basic interface settings on the Firepower 4100/9300 chassis. See Configure a Physical Interface, on page 563 for more information.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

Step 3 In the **Mode** drop-down list, choose **Passive** or **Erspan**.

Step 4 Enable the interface by checking the **Enabled** check box.

Step 5 In the **Name** field, enter a name up to 48 characters in length.

Step 6 From the **Security Zone** drop-down list, choose a security zone or add a new one by clicking **New**.

Step 7 (Optional) Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

Step 8 (Optional) On **General**, set the **MTU** between 64 and 9198 bytes; for the Firepower Threat Defense Virtual and Firepower Threat Defense on the FXOS chassis, the maximum is 9000 bytes.

The default is 1500 bytes.

Step 9 For ERSPAN interfaces, set the following parameters:

- **Flow Id**—Configure the ID used by the source and destination sessions to identify the ERSPAN traffic, between 1 and 1023. This ID must also be entered in the ERSPAN destination session configuration.
- **Source IP**—Configure the IP address used as the source of the ERSPAN traffic.

Step 10 For ERSPAN interfaces, set the IPv4 address and mask on **IPv4**.

Step 11 (Optional) Set the duplex and speed by clicking **Hardware Configuration**.

The exact speed and duplex options depend on your hardware.

- **Duplex**—Choose **Full**, **Half**, or **Auto**. Auto is the default.
- **Speed**—Choose **10**, **100**, **1000**, or **Auto**. Auto is the default.

Step 12 Click **OK**.

Step 13 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure an Inline Set

This section enables and names two physical interfaces that you can add to an inline set. You can also optionally enable Hardware Bypass for supported interface pairs.

Note For the FTD on the FXOS chassis, you configure basic interface settings on the Firepower 4100/9300 chassis. See Configure a Physical Interface, on page 563 for more information.

Before you begin

- We recommend that you set STP PortFast for STP-enabled switches that connect to the FTD inline pair interfaces. This setting is especially useful for Hardware Bypass configurations and can reduce bypass times.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

Step 3 In the **Mode** drop-down list, choose **None**.

After you add this interface to an inline set, this field will show Inline for the mode.

Step 4 Enable the interface by checking the **Enabled** check box.

Step 5 In the **Name** field, enter a name up to 48 characters in length.

Do not set the security zone yet; you must set it after you create the inline set later in this procedure.

Step 6 (Optional) Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

Step 7 (Optional) Set the duplex and speed by clicking **Hardware Configuration**.

The exact speed and duplex options depend on your hardware.

- **Duplex**—Choose **Full**, **Half**, or **Auto**. Auto is the default.
- **Speed**—Choose **10**, **100**, **1000**, or **Auto**. Auto is the default.

Step 8 Click **OK**.

Do not set any other settings for this interface.

Step 9 Click **Edit** ( ) for the second interface you want to add to the inline set.

Step 10 Configure the settings as for the first interface.

Step 11 Click **Inline Sets**.

Step 12 Click **Add Inline Set**.
The **Add Inline Set** dialog box appears with **General** selected.

Step 13 In the **Name** field, enter a name for the set.

Step 14 (Optional) Change the **MTU** to enable jumbo frames.

For inline sets, the MTU setting is not used. However, the jumbo frame setting *is* relevant to inline sets; jumbo frames enable the inline interfaces to receive packets up to 9000 bytes. To enable jumbo frames, you must set the MTU of *any* interface on the device above 1500 bytes.

Step 15 (Optional) For the **Bypass** mode, choose one of the following options:

- **Disabled**—Set Hardware Bypass to disabled for interfaces where Hardware Bypass is supported, or use interfaces where Hardware Bypass is not supported.
- **Standby**—Set Hardware Bypass to the standby state on supported interfaces. Only pairs of Hardware Bypass interfaces are shown. In the standby state, the interfaces remain in normal operation until there is a trigger event.
- **Bypass-Force**—Manually forces the interface pair to go into a bypass state. **Inline Sets** shows **Yes** for any interface pairs that are in Bypass-Force mode.

Step 16 In the **Available Interfaces Pairs** area, click a pair and then click **Add** to move it to the **Selected Interface Pair** area.

All possible pairings between named and enabled interfaces with the mode set to None show in this area.

Step 17 (Optional) Click **Advanced** to set the following optional parameters:

- **Tap Mode**—Set to inline tap mode.

  Note that you cannot enable this option and strict TCP enforcement on the same inline set.

  Note Tap mode *significantly* impacts the FTD performance, depending on the traffic.

- **Propagate Link State**—Configure link state propagation.

  Link state propagation automatically brings down the second interface in the inline interface pair when one of the interfaces in an inline set goes down. When the downed interface comes back up, the second interface automatically comes back up, also. In other words, if the link state of one interface changes, the device senses the change and updates the link state of the other interface to match it. Note that devices require up to 4 seconds to propagate link state changes. Link state propagation is especially useful in resilient network environments where routers are configured to reroute traffic automatically around network devices that are in a failure state.

- **Strict TCP Enforcement**—To maximize TCP security, you can enable strict enforcement, which blocks connections where the three-way handshake was not completed.

  Strict enforcement also blocks:

  - Non-SYN TCP packets for connections where the three-way handshake was not completed

- Non-SYN/RST packets from the initiator on a TCP connection before the responder sends the SYN-ACK
- Non-SYN-ACK/RST packets from the responder on a TCP connection after the SYN but before the session is established
- SYN packets on an established TCP connection from either the initiator or the responder

- **Snort Fail Open**—Enable or disable either or both of the **Busy** and **Down** options if you want new and existing traffic to pass without inspection (enabled) or drop (disabled) when the Snort process is busy or down.

  By default, traffic passes without inspection when the Snort process is down, and drops when it is busy.

  When the Snort process is:

  - Busy—It cannot process traffic fast enough because traffic buffers are full, indicating that there is more traffic than the device can handle, or because of other software resource issues.
  - Down—It is restarting because you deployed a configuration that requires it to restart. See Configurations that Restart the Snort Process When Deployed or Activated, on page 352.

    When the Snort process is down and comes back up, it inspects new connections. To prevent false positives and false negatives, it does not inspect existing connections on inline, routed, or transparent interfaces because initial session information might have been lost while it was down.

Note When Snort fails open, features that rely on the Snort process do not function. These include application control and deep inspection. The system performs only basic access control using simple, easily determined transport and network layer characteristics.

Step 18 Click **Interfaces**.

Step 19 Click **Edit** () for one of the member interfaces.

Step 20 From the **Security Zone** drop-down list, choose a security zone or add a new one by clicking **New**.

You can only set the zone after you add the interface to the inline set; adding it to an inline set configures the mode to Inline and lets you choose inline-type security zones.

Step 21 Click **OK**.

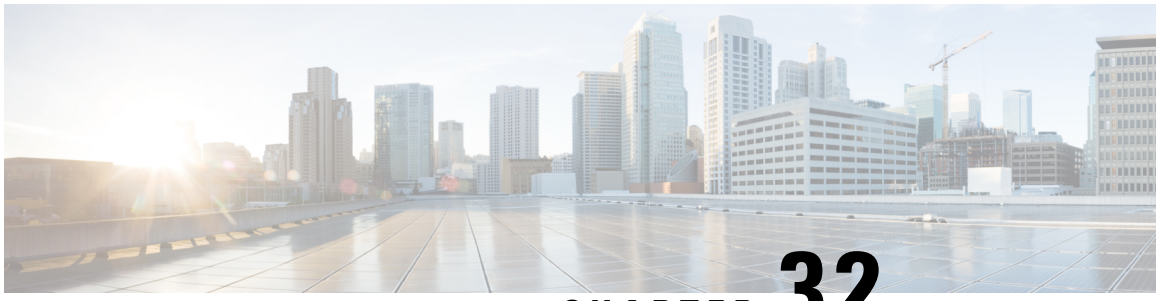
Step 22 Set the security zone for the second interface.

Step 23 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

History for Inline Sets and Passive Interfaces for Firepower Threat Defense

Feature	Version	Details
Support for EtherChannels in Firepower Threat Defense inline sets	6.2.0	You can now use EtherChannels in a Firepower Threat Defense inline set. Supported platforms: Firepower 4100/9300, Firepower 2100 (6.2.1 and later)
Hardware bypass support on the Firepower 4100/9300 for supported network modules	6.1.0	Hardware Bypass ensures that traffic continues to flow between an inline interface pair during a power outage. This feature can be used to maintain network connectivity in the case of software or hardware failures. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Edit Physical Interface** Supported platforms: Firepower 4100/9300
Inline set link state propagation support for the Firepower Threat Defense	6.1.0	When you configure an inline set in the Firepower Threat Defense application and enable link state propagation, the Firepower Threat Defense sends inline set membership to the FXOS chassis. Link state propagation means that the chassis automatically brings down the second interface in the inline interface pair when one of the interfaces in an inline set goes down. New/Modified FXOS commands: **show fault |grep link-down**, **show interface detail** Supported platforms: Firepower 4100/9300, Firepower 2100 (6.2.1 and later)



CHAPTER 32

DHCP and DDNS Services for Threat Defense

The following topics explain DHCP and DDNS services and how to configure them on Threat Defense devices.

- About DHCP and DDNS Services, on page 663
- Requirements and Prerequisites for DHCP and DDNS, on page 664
- Guidelines for DHCP and DDNS Services, on page 664
- Configure the DHCP Server, on page 666
- Configure the DHCP Relay Agent, on page 667
- Configure Dynamic DNS, on page 668

About DHCP and DDNS Services

The following topics describe the DHCP server, DHCP relay agent, and DDNS update.

About the DHCPv4 Server

DHCP provides network configuration parameters, such as IP addresses, to DHCP clients. The FTD device can provide a DHCP server to DHCP clients attached to FTD device interfaces. The DHCP server provides network configuration parameters directly to DHCP clients.

An IPv4 DHCP client uses a broadcast rather than a multicast address to reach the server. The DHCP client listens for messages on UDP port 68; the DHCP server listens for messages on UDP port 67.

The DHCP server for IPv6 is not supported; you can, however, enable DHCP relay for IPv6 traffic.

DHCP Options

DHCP provides a framework for passing configuration information to hosts on a TCP/IP network. The configuration parameters are carried in tagged items that are stored in the Options field of the DHCP message and the data are also called options. Vendor information is also stored in Options, and all of the vendor information extensions can be used as DHCP options.

For example, Cisco IP Phones download their configuration from a TFTP server. When a Cisco IP Phone starts, if it does not have both the IP address and TFTP server IP address preconfigured, it sends a request with option 150 or 66 to the DHCP server to obtain this information.

- DHCP option 150 provides the IP addresses of a list of TFTP servers.
- DHCP option 66 gives the IP address or the hostname of a single TFTP server.

- DHCP option 3 sets the default route.

A single request might include both options 150 and 66. In this case, the ASA DHCP server provides values for both options in the response if they are already configured on the ASA.

You can use advanced DHCP options to provide DNS, WINS, and domain name parameters to DHCP clients; DHCP option 15 is used for the DNS domain suffix. You can also use the DHCP automatic configuration setting to obtain these values or define them manually. When you use more than one method to define this information, it is passed to DHCP clients in the following sequence:

1. Manually configured settings.
2. Advanced DHCP options settings.
3. DHCP automatic configuration settings.

For example, you can manually define the domain name that you want the DHCP clients to receive and then enable DHCP automatic configuration. Although DHCP automatic configuration discovers the domain together with the DNS and WINS servers, the manually defined domain name is passed to DHCP clients with the discovered DNS and WINS server names, because the domain name discovered by the DHCP automatic configuration process is superseded by the manually defined domain name.

About the DHCP Relay Agent

You can configure a DHCP relay agent to forward DHCP requests received on an interface to one or more DHCP servers. DHCP clients use UDP broadcasts to send their initial DHCPDISCOVER messages because they do not have information about the network to which they are attached. If the client is on a network segment that does not include a server, UDP broadcasts normally are not forwarded by the FTD device because it does not forward broadcast traffic. The DHCP relay agent lets you configure the interface of the FTD device that is receiving the broadcasts to forward DHCP requests to a DHCP server on another interface.

Requirements and Prerequisites for DHCP and DDNS

Model Support

FTD

User Roles

- Admin
- Access Admin
- Network Admin

Guidelines for DHCP and DDNS Services

This section includes guidelines and limitations that you should check before configuring DHCP and DDNS services.

Firewall Mode

- DHCP Relay is not supported in transparent firewall mode or in routed mode on the BVI or bridge group member interface.
- DHCP Server is supported in transparent firewall mode on a bridge group member interface. In routed mode, the DHCP server is supported on the BVI interface, not the bridge group member interface. The BVI must have a name for the DHCP server to operate.
- DDNS is not supported in transparent firewall mode or in routed mode on the BVI or bridge group member interface.

IPv6

Does not support IPv6 for DHCP server; IPv6 for DHCP relay is supported.

DHCPv4 Server

- The maximum available DHCP pool is 256 addresses.
- You can configure only one DHCP server on each interface. Each interface can have its own pool of addresses to use. However the other DHCP settings, such as DNS servers, domain name, options, ping timeout, and WINS servers, are configured globally and used by the DHCP server on all interfaces.
- You cannot configure an interface as a DHCP client if that interface also has DHCP server enabled; you must use a static IP address.
- You cannot configure both a DHCP server and DHCP relay on the same device, even if you want to enable them on different interfaces; you can only configure one type of service.
- FTD device does not support QIP DHCP servers for use with the DHCP proxy service.
- The DHCP server does not support BOOTP requests.

DHCP Relay

- You can configure a maximum of 10 DHCPv4 relay servers, global and interface-specific servers combined, with a maximum of 4 servers per interface.
- You can configure a maximum of 10 DHCPv6 relay servers. Interface-specific servers for IPv6 are not supported.
- You cannot configure both a DHCP server and DHCP relay on the same device, even if you want to enable them on different interfaces; you can only configure one type of service.
- DHCP relay services are not available in transparent firewall mode. You can, however, allow DHCP traffic through using an access rule. To allow DHCP requests and replies through the FTD device, you need to configure two access rules, one that allows DCHP requests from the inside interface to the outside (UDP destination port 67), and one that allows the replies from the server in the other direction (UDP destination port 68).
- For IPv4, clients must be directly-connected to the FTD device and cannot send requests through another relay agent or a router. For IPv6, the FTD device supports packets from another relay server.
- The DHCP clients must be on different interfaces from the DHCP servers to which the FTD device relays requests.

- You cannot enable DHCP Relay on an interface in a traffic zone.
- DHCP relay is not supported on Virtual Tunnel Interfaces (VTIs).

Configure the DHCP Server

See the following steps to configure a DHCP server.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Select **DHCP > DHCP Server**.

Step 3 Configure the following DHCP server options:

- **Ping Timeout**—The amount of time in milliseconds that Firepower Threat Defense device waits to time out a DHCP ping attempt. Valid values range from 10 to 10000 milliseconds. The default value is 50 milliseconds.

  To avoid address conflicts, the Firepower Threat Defense device sends two ICMP ping packets to an address before assigning that address to a DHCP client.
- **Lease Length**—The amount of time in seconds that the client may use its allocated IP address before the lease expires. Valid values range from 300 to 1048575 seconds. The default value is 3600 seconds (1 hour).
- (Routed mode) **Auto-configuration**—Enables DHCP auto configuration on the Firepower Threat Defense device. Auto-configuration enables the DHCP server to provide the DHCP clients with the DNS server, domain name, and WINS server information obtained from a DHCP client running on the specified interface. Otherwise, you can disable auto configuration and add the values yourself in Step 4.
- (Routed mode) **Interface**—Specifies the interface to be used for auto configuration.

Step 4 To override auto-configured settings, do the following:

- Enter the domain name of the interface. For example, your device may be in the Your_Company domain.
- From the drop-down list, choose the DNS servers (primary and secondary) configured for the interface. To add a new DNS server, see Creating Network Objects, on page 403.
- From the drop-down list, choose the WINS servers (primary and secondary) configured for the interface. To add a new WINS server, see Creating Network Objects, on page 403.

Step 5 Select **Server**, click **Add**, and configure the following options:

- **Interface**—Choose the interface from the drop-down list. In transparent mode, specify a named bridge group member interface. In routed mode, specify a named routed interface or a named BVI; do not specify the bridge group member interface. Note that each bridge group member interface for the BVI must also be named for the DHCP server to operate.
- **Address Pool**—The range of IP addresses from lowest to highest that is used by the DHCP server. The range of IP addresses must be on the same subnet as the selected interface and cannot include the IP address of the interface itself.

- **Enable DHCP Server**—Enables the DHCP server on the selected interface.

Step 6 Click **OK** to save the DHCP server configuration.

Step 7 (Optional) Select **Advanced**, click **Add**, and specify the type of information you want the option to return to the DHCP client:

- **Option Code**—The Firepower Threat Defense device supports the DHCP options listed in RFC 2132, RFC 2562, and RFC 5510 to send information. All DHCP options (1 through 255) are supported except for 1, 12, 50–54, 58–59, 61, 67, and 82. See About the DHCPv4 Server, on page 663for more information on DHCP option codes.

  Note The Firepower Threat Defense device does not verify that the option type and value that you provide match the expected type and value for the option code, as defined in RFC 2132. For more information about option codes and their associated types and expected values, see RFC 2132.

- **Type**—DHCP option type. Available options include **IP**, **ASCII**, and **HEX**. If you chose IP, you must add IP addresses in the IP Address fields. If you chose ASCII, you must add the ASCII value in the ASCII field. If you chose HEX, you must add the HEX value in the HEX field.
- **IP Address 1** and **IP Address 2**—The IP address(es) to be returned with this option code. To add a new IP address, see Creating Network Objects, on page 403.
- **ASCII**—The ASCII value that is returned to the DHCP client. The string cannot include spaces.
- **HEX**—The HEX value that is returned to the DHCP client. The string must have an even number of digits and no spaces. You do not need to use a 0x prefix.

Step 8 Click **OK** to save the option code configuration.

Step 9 Click **Save** on the DHCP page to save your changes.

Configure the DHCP Relay Agent

You can configure a DHCP relay agent to forward DHCP requests received on an interface to one or more DHCP servers. DHCP clients use UDP broadcasts to send their initial DHCPDISCOVER messages because they do not have information about the network to which they are attached. If the client is on a network segment that does not include a server, UDP broadcasts normally are not forwarded by the Firepower Threat Defense device because it does not forward broadcast traffic.

You can remedy this situation by configuring the interface of the Firepower Threat Defense device that is receiving the broadcasts to forward DHCP requests to a DHCP server on another interface.



Note DHCP Relay is not supported in transparent firewall mode.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Select **DHCP** > **DHCP Relay**.

Step 3 In the **Timeout** field, enter the amount of time in seconds that the Firepower Threat Defense device waits to time out the DHCP relay agent. Valid values range from 1 to 3600 seconds. The default value is 60 seconds.

The timeout is for address negotiation through the local DHCP Relay agent.

Step 4 On **DHCP Relay Agent**, click **Add**, and configure the following options:

- **Interface**—The interface connected to the DHCP clients.
- **Enable IPv4 Relay**—Enables IPv4 DHCP Relay for this interface.
- **Set Route**—(For IPv4) Changes the default gateway address in the DHCP message from the server to that of the Firepower Threat Defense device interface that is closest to the DHCP client, which relayed the original DHCP request. This action allows the client to set its default route to point to the Firepower Threat Defense device even if the DHCP server specifies a different router. If there is no default router option in the packet, the Firepower Threat Defense device adds one containing the interface address.
- **Enable IPv6 Relay**—Enables IPv6 DHCP Relay for this interface.

Step 5 Click **OK** to save the DHCP relay agent changes.

Step 6 On **DHCP Servers**, click **Add**, and configure the following options:

Add the IPv4 and IPv6 server addresses as separate entries, even if they belong to the same server.

- **Server**—The IP address of the DHCP server. Chose an IP address from the drop-down list. To add a new one, see Creating Network Objects, on page 403
- **Interface**—The interface to which the specified DHCP server is attached. The DHCP Relay agent and the DHCP server cannot be configured on the same interface.

Step 7 Click **OK** to save the DHCP server changes.

Step 8 Click **Save** on the DHCP page to save your changes.

Configure Dynamic DNS

When an interface uses DHCP IP addressing, the assigned IP address can change when the DHCP lease is renewed. When the interface needs to be reachable using a fully qualified domain name (FQDN), the IP address change can cause the DNS server resource records (RRs) to become stale. Dynamic DNS (DDNS) provides a mechanism to update DNS RRs whenever the IP address or hostname changes. You can also use DDNS for static or PPPoE IP addressing.

DDNS updates the following RRs on the DNS server: the A RR includes the name-to-IP address mapping, while the PTR RR maps addresses to names.

The FTD supports the DDNS update method is defined by RFC 2136. It does not support the Web update method. With this method, the FTD and the DHCP server use DNS requests to update the DNS RRs. The FTD or DHCP server sends a DNS request to its local DNS server for information about the hostname and,

based on the response, determines the main DNS server that owns the RRs. The FTD or DHCP server then sends an update request directly to the main DNS server. See the following typical scenarios.

- The FTD updates the A RR, and the DHCP server updates the PTR RR.

  Typically, the FTD "owns" the A RR, while the DHCP server "owns" the PTR RR, so both entities need to request updates separately. When the IP address or hostname changes, the FTD sends a DHCP request to the DHCP server to inform it that it needs to request a PTR RR update.

- The DHCP server updates both the A and PTR RR.

  Use this scenario if the FTD does not have the authority to update the A RR. When the IP address or hostname changes, the FTD sends a DHCP request to the DHCP server to inform it that it needs to request an A and PTR RR update.

You can configure different ownership depending on your security needs and the requirements of the main DNS server. For example, for a static address, the FTD should own the updates for both records.

The **DDNS** page also supports setting DHCP server settings relating to DDNS.

Note DDNS is not supported on the BVI or bridge group member interfaces.

Before you begin

- Configure a DNS server using FlexConfig. See FlexConfig Policies for FTD, on page 575. You must configure a DNS server unless you enable DHCP server-only updates.
- Configure the device hostname. You can configure the hostname when you perform the FTD initial setup, or by using the **configure network hostname** command. If you do not specify the hostname per interface, then the device hostname is used.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Choose **DHCP** > **DDNS**.

Step 3 Configure a DDNS update method to enable DNS requests from the FTD.

You do not need to configure a DDNS update method if the DHCP server will perform all requests.

a) On **DDNS Update Methods**, click **Add**.
b) Set the **Method Name**.
c) (Optional) Configure the **Update Interval** between DNS requests. By default when all values are set to 0, update requests are sent whenever the IP address or hostname changes. To send requests regularly, set the **Days** (0-364), **Hours**, **Minutes**, and **Seconds**.
d) Set the **Update Records** you want the FTD to update.

   This setting only affects the records you want to update directly from the FTD; to determine the records you want the DHCP server to update, configure the DHCP client settings per interface or globally. See Step Step 4, on page 670.

   - **Not Defined**—Disables DNS updates from the FTD.

- **Both A and PTR Records**—Sets the FTD to update both A and PTR RRs. Use this option for static or PPPoE IP addressing.
- **A Records**—Sets the FTD to update the A RR only. Use this option if you want the DHCP server to update the PTR RR.

e) Click **OK**.
f) Assign this method to the interface in Step Step 4, on page 670.

Step 4 Configure interface settings for DDNS, including setting the update method, DHCP client settings, and the hostname for this interface.

a) On **DDNS Interface Settings**, click **Add**.
b) Choose the **Interface** from the drop-down list.
c) Choose the **Method Name** that you created on the **DDNS Update Methods** page.

You do not need to assign a method if you want the DHCP server to perform all updates.

d) Set the **Host Name** for this interface.

If you do not set the hostname, the device hostname is used. If you do not specify an FQDN, then the default domain from the DNS server group is appended (for static or PPPoE IP addressing) or the domain name from the DHCP server is appended (for DHCP IP addressing).

e) Configure the **DHCP Client requests DHCP server to update requests** to determine which records you want the DHCP server to update.

The FTD sends DHCP client requests to the DHCP server. Note that the DHCP server must also be configured to support DDNS. The server can be configured to honor the client requests, or it can override the client (in which case, it will reply to the client so the client does not also try to perform updates that the server is performing).

For static or PPPoE IP addressing, these settings are ignored.

Note You can also set these values globally for all interfaces on the **DDNS** page. The per-interface settings take precedence over the global settings.

- **Not Selected**—Disables DDNS requests to the DHCP server. Even if the client does not request DDNS updates, the DHCP server can be configured to send updates anyway.
- **No Update**—Requests the DHCP server not to perform updates. This setting works in conjunction with a DDNS update method with **Both A and PTR Records** enabled.
- **Only PTR**—Requests that the DHCP server perform the PTR RR update. This setting works in conjunction with a DDNS update method with **A Records** enabled.
- **Both A and PTR Records**—Requests that the DHCP server perform both A and PTR RR updates. This setting does not require a DDNS update method to be associated with the interface.

f) Click **OK**.

Note The **Dynamic DNS Update** settings relate to DHCP server settings when you enable a DHCP server on the FTD. See Step Step 5, on page 670 for more information.

Step 5 If you enable the DHCP server on an FTD, you can configure DHCP server settings for DDNS.

To enable the DHCP server, see Configure the DHCP Server, on page 666). You can configure the server behavior when DHCP clients use the standard DDNS update method. If the server performs any updates, then if the client lease expires (and is not renewed), the server will request that the DNS server remove the RRs for which it was responsible.

a) You can configure server settings globally or per interface. For global settings, see the main **DDNS** page. For per-interface settings, see the **DDNS Interface Settings** page. Interface settings take precedence over global settings.
b) Configure which DNS RRs you want the DHCP server to update under **Dynamic DNS Update**.
   - **Not Selected**—DDNS updates are disabled, even if the client requests them.
   - **Only PTR**—Enables DDNS updates. If you enable the **Override DHCP Client Requests** setting, then the server will only update the PTR RR. Otherwise, the server will update RRs that the client requests. If the client does not send an update request with the FQDN option, the server will request an update for both A and PTR RRs using the hostname discovered in DHCP option 12.
   - **Both A and PTR Records**—Enables DDNS updates. If you enable the **Override DHCP Client Requests** setting, then the server will update both the A and PTR RRs. Otherwise, the server will update RRs that the client requests. If the client does not send an update request with the FQDN option, the server will request an update for both A and PTR RRs using the hostname discovered in DHCP option 12.
c) To override the update actions requested by the DHCP client, check **Override DHCP Client Requests**.

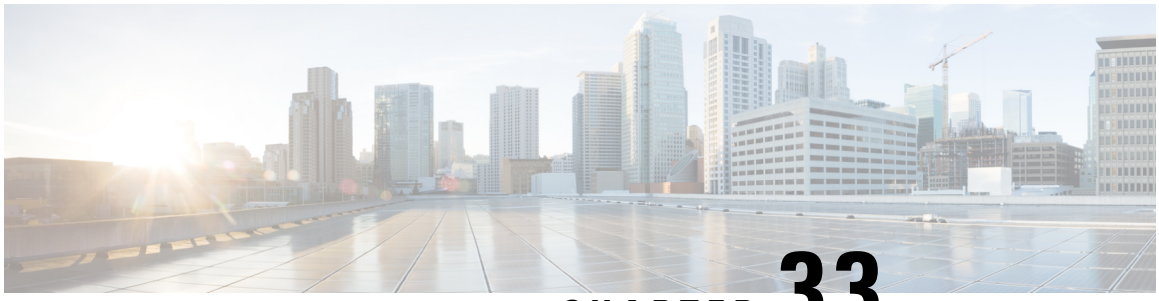
   The server will reply to the client that the request was overridden, so the client does not also try to perform updates that the server is performing.

Step 6 (Optional) Configure general DHCP client settings. These settings are not related to DDNS, but are related to how the DHCP client behaves.

a) On the **DDNS** page, check **Enable DHCP Client Broadcast** to request that the DHCP server broadcast the DHCP reply (DHCP option 1).
b) To force a MAC address to be stored inside a DHCP request packet for option 61 instead of the default internally generated string, on **DDNS** > **DHCP Client ID Interface**, choose the interface from the **Available Interfaces** list, and then click **Add** to move it to the **Selected Interfaces** list.

   Some ISPs expect option 61 to be the interface MAC address. If the MAC address is not included in the DHCP request packet, then an IP address will not be assigned. This setting does not directly relate to DDNS, but is a general DHCP client setting.

Step 7 Click **Save** on the Device page to save your changes.



CHAPTER 33

Quality of Service (QoS) for Firepower Threat Defense

The following topics describe how to use the Quality of Service (QoS) feature to police network traffic using Firepower Threat Defense devices:

- Introduction to QoS, on page 673
- About QoS Policies, on page 673
- Requirements and Prerequisites for QoS, on page 674
- Rate Limiting with QoS Policies, on page 675

Introduction to QoS

Quality of Service, or QoS, rate limits (polices) network traffic that is allowed or trusted by access control. The system does not rate limit traffic that was fastpathed.

QoS is supported for routed interfaces on Firepower Threat Defense devices only.

Logging Rate-Limited Connections

There are no logging configurations for QoS. A connection can be rate limited without being logged, and you cannot log a connection simply because it was rate limited. To view QoS information in connection events, you must independently log the ends of the appropriate connections to the Firepower Management Center database; see Other Connections You Can Log, on page 2209.

Connection events for rate-limited connections contain information on how much traffic was dropped, and which QoS configurations limited the traffic. You can view this information in event views (workflows), dashboards, and reports.

About QoS Policies

QoS policies deployed to managed devices govern rate limiting. Each QoS policy can target multiple devices; each device can have one deployed QoS policy at a time.

The system matches traffic to QoS rules in the order you specify. The system rate limits traffic according to the first rule where all rule conditions match the traffic. Traffic that does not match any of the rules is not rate limited.

Note The total number of rules including QoS rules on the device cannot exceed 255. When this threshold is reached, a deployment warning message is displayed. You need to reduce the number of rules for a successful deployment.

You must constrain QoS rules by source or destination (routed) interfaces. The system enforces rate limiting *independently* on *each* of those interfaces; you cannot specify an aggregate rate limit for a set of interfaces.

QoS rules can also rate limit traffic by other network characteristics, as well as contextual information such as application, URL, and user identity.

You can rate limit download and upload traffic independently. The system determines download and upload directions based on the connection initiator.

Note QoS is not subordinate to a main access control configuration; you configure QoS independently. However, the access control and QoS policies deployed to the same device share identity configurations; see Associating Other Policies with Access Control, on page 1213.

QoS Policies and Multitenancy

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Administrators in ancestor domains can deploy the same QoS policy to devices in different descendant domains. Administrators in those descendant domains can use this read-only ancestor-deployed QoS policy, or replace it with a local policy.

Requirements and Prerequisites for QoS

Model Support

FTD

Supported Domains

Any

User Roles

Admin

Access Admin

Network Admin

Rate Limiting with QoS Policies

To perform policy-based rate limiting, configure and deploy QoS policies to managed devices. Each QoS policy can target mutiple devices; each device can have one deployed QoS policy at a time.

Only one person should edit a policy at a time, using a single browser window. If multiple users save the same policy, the last saved changes are retained. For your convenience, the system displays information on who (if anyone) is currently editing each policy. To protect the privacy of your session, a warning appears after 30 minutes of inactivity on the policy editor. After 60 minutes, the system discards your changes.

Procedure

Step 1 Choose **Devices** > **QoS**.

Step 2 Click **New Policy** to create a new QoS policy and, optionally, assign target devices; see Creating a QoS Policy, on page 676.

You can also **Copy** () or **Edit** () an exisiting policy.

Step 3 Configure QoS rules; see Configuring QoS Rules, on page 677 and Rule Management: Common Characteristics, on page 361.

The Rules in the QoS policy editor lists each rule in evaluation order, and displays a summary of the rule conditions and rate limiting configurations. A right-click menu provides rule management options, including moving, enabling, and disabling.

Helpful in larger deployments, you can **Filter by Device** to display only the rules that affect a specfic device or group of devices. You can also search for and within rules; the system matches text you enter in the **Search Rules** field to rule names and condition values, including objects and object groups.

Note Properly creating and ordering rules is a complex task, but one that is essential to building an effective deployment. If you do not plan carefully, rules can preempt other rules, require additional licenses, or contain invalid configurations. Icons represent comments, warnings, and errors. If issues exist, click **Show Warnings** to display a list. For more information, see Best Practices for Access Control Rules, on page 1194.

Step 4 Click **Policy Assignments** to identify the managed devices targeted by the policy; see Setting Target Devices for a QoS Policy, on page 676.

If you identified target devices during policy creation, verify your choices.

Step 5 Save the QoS policy.

Step 6 Because this feature must allow some packets to pass, you must configure your system to examine those packets. See Best Practices for Handling Packets That Pass Before Traffic Identification, on page 1670 and Specify a Policy to Handle Packets That Pass Before Traffic Identification, on page 1670.

Step 7 Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Creating a QoS Policy

A new QoS policy with no rules performs no rate limiting.

Procedure

Step 1 Choose **Devices** > **QoS**.

Step 2 Click **New Policy**.

Step 3 Enter a **Name** and, optionally, a **Description**.

Step 4 (Optional) Choose the **Available Devices** where you want to deploy the policy, then click **Add to Policy**, or drag and drop to the **Selected Devices**. To narrow the devices that appear, type a search string in the **Search** field.

You must assign devices before you deploy the policy.

Step 5 Click **Save**.

What to do next

- Configure and deploy the QoS policy; see Rate Limiting with QoS Policies, on page 675.

Setting Target Devices for a QoS Policy

Each QoS policy can target mutiple devices; each device can have one deployed QoS policy at a time.

Procedure

Step 1 In the QoS policy editor, click **Policy Assignments**.

Step 2 Build your target list:

- Add—Choose one or more **Available Devices**, then click **Add to Policy** or drag and drop into the list of **Selected Devices**.
- Delete—Click **Delete** ( ) next to a single device, or choose multiple devices, right-click, then choose **Delete Selected**.
- Search—Enter a search string in the search field. Click **Clear** ( ) to clear the search.

Step 3 Click **OK** to save policy assignments.

Step 4 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configuring QoS Rules

When you create or edit a rule, use the upper portion of the rule editor to configure general rule properties. Use the lower portion of the rule editor to configure rule conditions and comments.

Procedure

Step 1 On Rules of the QoS policy editor:

- Add Rule—Click **Add Rule**.
- Edit Rule—Click **Edit** ( ).

Step 2 Enter a **Name**.

Step 3 Configure rule components:

- Enabled—Specify whether the rule is **Enabled**.
- Apply QoS On—Choose the interfaces you want to rate limit, either **Interfaces in Destination Interface Objects** or **Interfaces in Source Interface Objects**. Your choice must correspond with a populated interface constraint (not **any**).
- Traffic Limit Per Interface—Enter a **Download/Upload Limit** in Mbits/sec. Click **Advanced** to enter different limits for download and upload traffic. The default value of **Unlimited** prevents matching traffic from being rate limited.
- Conditions—Click the corresponding condition you want to add. You must configure a source or destination interface condition, corresponding to your choice for **Apply QoS On**.
- Comments—Click **Comments**. To add a comment click **New Comment**, enter a comment, and click **OK**. You can edit or delete this comment until you save the rule.

For detailed information on rule components, see QoS Rule Components, on page 678.

Step 4 Save the rule.

Step 5 In the policy editor, set the rule position. Click and drag or use the right-click menu to cut and paste.

Rules are numbered starting at 1. The system matches traffic to rules in top-down order by ascending rule number. The first rule that traffic matches is the rule that handles that traffic. Proper rule order reduces the resources required to process network traffic and prevents rule preemption.

Step 6 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Best Practices for Access Control Rules, on page 1194

QoS Rule Components

State (Enabled/Disabled)

By default, rules are enabled. If you disable a rule, the system does not use it and stops generating warnings and errors for that rule.

Interfaces (Apply QoS On)

You cannot save a QoS rule that rate limits all traffic. For each QoS rule, you must apply QoS on either:

- Interfaces in Source Interface Objects—Rate limits traffic through the rule's source interfaces. If you choose this option, you must add at least one source interface constraint (cannot be **any**).
- Interfaces in Destination Interface Objects—Rate limits traffic through the rule's destination interfaces. If you choose this option, you must add at least one destination interface constraint (cannot be **any**).

Traffic Limit Per Interface

A QoS rule enforces rate limiting *independently* on *each* of the interfaces you specify with the Apply QoS On option. You cannot specify an aggregate rate limit for a set of interfaces.

You can rate limit traffic by Mbits per second. The default value of **Unlimited** prevents matching traffic from being rate limited.

Advanced settings allow you to rate limit download and upload traffic independently. The system determines download and upload directions based on the connection initiator.

If you specify a limit greater than the maximum throughput of an interface, the system does not rate limit matching traffic. Maximum throughput may be affected by an interface's hardware configuration, which you specify in each device's properties (**Devices** > **Device Management**).

Conditions

Conditions specify the specific traffic the rule handles. You can configure each rule with multiple conditions. Traffic must match all conditions to match the rule. Each condition type has its own tab in the rule editor. You can rate limit traffic using:

- Interface Conditions, on page 366 (routed only; required)
- Network Conditions, on page 368
- Port and ICMP Code Conditions, on page 373
- Application Conditions (Application Control), on page 375
- URL Filtering, on page 1229
- User, Realm, and ISE Attribute Conditions (User Control), on page 384

Comments

Each time you save changes to a rule you can add comments. For example, you might summarize the overall configuration for the benefit of other users, or note when you change a rule and the reason for the change.

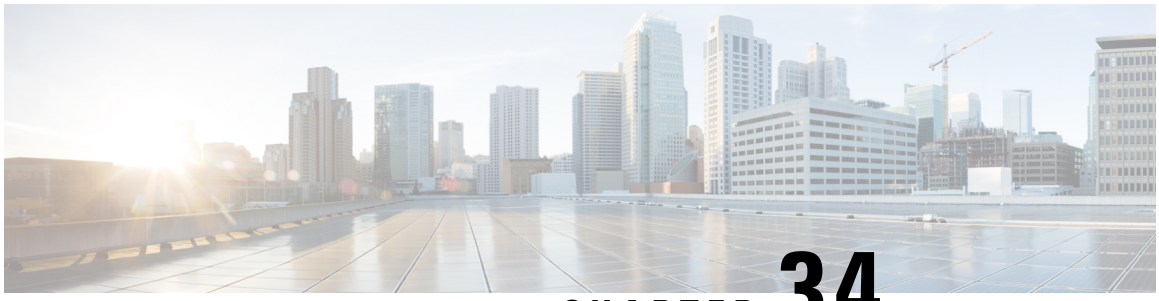
In the policy editor, the system displays how many comments a rule has. In the rule editor, use the Comments tab to view existing comments and add new ones.



PART IX

Firepower Threat Defense High Availability and Scalability

- High Availability for FTD, on page 683
- Clustering for the Firepower Threat Defense, on page 709



CHAPTER 34

High Availability for FTD

The following topics describe how to configure Active/Standby failover to accomplish high availability of the Firepower Threat Defense.

- About Firepower Threat Defense High Availability, on page 683
- Requirements and Prerequisites for High Availability, on page 696
- Guidelines for High Availability, on page 697
- Add a Firepower Threat Defense High Availability Pair, on page 699
- Configure Optional High Availability Parameters, on page 700
- Manage High Availability, on page 702
- Monitoring High Availability, on page 706

About Firepower Threat Defense High Availability

Configuring high availability, also called failover, requires two identical Firepower Threat Defense devices connected to each other through a dedicated failover link and, optionally, a state link. Firepower Threat Defense supports Active/Standby failover, where one unit is the active unit and passes traffic. The standby unit does not actively pass traffic, but synchronizes configuration and other state information from the active unit. When a failover occurs, the active unit fails over to the standby unit, which then becomes active.

The health of the active unit (hardware, interfaces, software, and environmental status) is monitored to determine if specific failover conditions are met. If those conditions are met, failover occurs.

Note High availability is not supported on Firepower Threat Defense Virtual running in the public cloud.

High Availability System Requirements

This section describes the hardware, software, and license requirements for Firepower Threat Defense devices in a High Availability configuration.

Hardware Requirements

The two units in a High Availability configuration must:

- Be the same model.

- Have the same number and types of interfaces.

  For the Firepower 4100/9300 chassis, all interfaces must be preconfigured in FXOS identically before you enable High Availability. If you change the interfaces after you enable High Availability, make the interface changes in FXOS on the Standby unit, and then make the same changes on the Active unit.

If you are using units with different flash memory sizes in your High Availability configuration, make sure the unit with the smaller flash memory has enough space to accommodate the software image files and the configuration files. If it does not, configuration synchronization from the unit with the larger flash memory to the unit with the smaller flash memory will fail.

Software Requirements

The two units in a High Availability configuration must:

- Be in the same firewall mode (routed or transparent).
- Have the same software version.
- Be in the same domain or group on the FMC.
- Have the same NTP configuration. See Configure NTP Time Synchronization for Threat Defense, on page 976.
- Be fully deployed on the FMC with no uncommitted changes.
- Not have DHCP or PPPoE configured in any of their interfaces.
- (Firepower 9300) Have the same flow offload mode, either both enabled or both disabled.

License Requirements for FTD Devices in a High Availability Pair

Firepower Threat Defense devices in a high availability configuration must have the same licenses.

High availability configurations require two Smart License entitlements; one for each device in the pair.

Before high availability is established, it does not matter which licenses are assigned to the secondary/standby device. During high availability configuration, the Firepower Management Center releases any unnecessary licenses assigned to the standby device and replaces them with identical licenses assigned to the primary/active device. For example, if the active device has a Base license and a Threat license, and the standby device has only a Base license, the Firepower Management Center communicates with the Cisco Smart Software Manager to obtain an available Threat license from your account for the standby device. If your Smart Licenses account does not include enough purchased entitlements, your account becomes Out-of-Compliance until you purchase the correct number of licenses.

In a virtual Firepower Management Center high availability configuration, each FTD to be registered requires an additional Firepower MCv Device license.

Failover and Stateful Failover Links

The failover link and the optional stateful failover link are dedicated connections between the two units. Cisco recommends to use the same interface between two devices in a failover link or a stateful failover link. For example, in a failover link, if you have used eth0 in device 1, use the same interface (eth0) in device 2 as well.

Failover Link

The two units in a failover pair constantly communicate over a failover link to determine the operating status of each unit.

Failover Link Data

The following information is communicated over the failover link:

- The unit state (active or standby)
- Hello messages (keep-alives)
- Network link status
- MAC address exchange
- Configuration replication and synchronization

Interface for the Failover Link

You can use an unused data interface (physical, redundant, or EtherChannel) as the failover link; however, you cannot specify an interface that is currently configured with a name. You also cannot use a subinterface. The failover link interface is not configured as a normal networking interface; it exists for failover communication only. This interface can only be used for the failover link (and also for the state link).

The FTD does not support sharing interfaces between user data and the failover link.



Note When using an EtherChannel or redundant interface as the failover or state link, you must confirm that the same EtherChannel or redundant interface with the same member interfaces exists on both devices before establishing high availability.

See the following guidelines for the failover link:

- Firepower 4100/9300—We recommend that you use a 10 GB data interface for the combined failover and state link.
- All other models—1 GB interface is large enough for a combined failover and state link.

For a redundant interface used as the failover link, see the following benefits for added redundancy:

- When a failover unit boots up, it alternates between the member interfaces to detect an active unit.
- If a failover unit stops receiving keepalive messages from its peer on one of the member interfaces, it switches to the other member interface.

The alternation frequency is equal to the unit hold time.



Note If you have a large configuration and a low unit hold time, alternating between the member interfaces can prevent the secondary unit from joining/re-joining. In this case, disable one of the member interfaces until after the secondary unit joins.

For an EtherChannel used as the failover link, to prevent out-of-order packets, only one interface in the EtherChannel is used. If that interface fails, then the next interface in the EtherChannel is used. You cannot alter the EtherChannel configuration while it is in use as a failover link.

Connecting the Failover Link

Connect the failover link in one of the following two ways:

- Using a switch, with no other device on the same network segment (broadcast domain or VLAN) as the failover interfaces of the Firepower Threat Defense device.
- Using an Ethernet cable to connect the units directly, without the need for an external switch.

If you do not use a switch between the units, if the interface fails, the link is brought down on both peers. This condition may hamper troubleshooting efforts because you cannot easily determine which unit has the failed interface and caused the link to come down.

Stateful Failover Link

To use Stateful Failover, you must configure a Stateful Failover link (also known as the state link) to pass connection state information.

Shared with the Failover Link

Sharing a failover link is the best way to conserve interfaces. However, you must consider a dedicated interface for the state link and failover link, if you have a large configuration and a high traffic network.

Dedicated Interface for the Stateful Failover Link

You can use a dedicated data interface (physical, redundant, or EtherChannel) for the state link. See Interface for the Failover Link, on page 685 for requirements for a dedicated state link, and Connecting the Failover Link, on page 686 for information about connecting the state link as well.

For optimum performance when using long distance failover, the latency for the state link should be less than 10 milliseconds and no more than 250 milliseconds. If latency is more than 10 milliseconds, some performance degradation occurs due to retransmission of failover messages.

Avoiding Interrupted Failover and Data Links

We recommend that failover links and data interfaces travel through different paths to decrease the chance that all interfaces fail at the same time. If the failover link is down, the Firepower Threat Defense device can use the data interfaces to determine if a failover is required. Subsequently, the failover operation is suspended until the health of the failover link is restored.

See the following connection scenarios to design a resilient failover network.

Scenario 1—Not Recommended

If a single switch or a set of switches are used to connect both failover and data interfaces between two Firepower Threat Defense devices, then when a switch or inter-switch-link is down, both Firepower Threat Defense devices become active. Therefore, the two connection methods shown in the following figures are **not** recommended.

Figure 13: Connecting with a Single Switch—Not Recommended

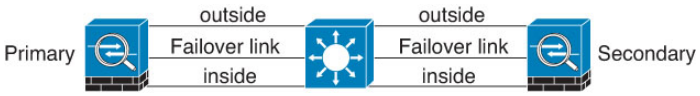
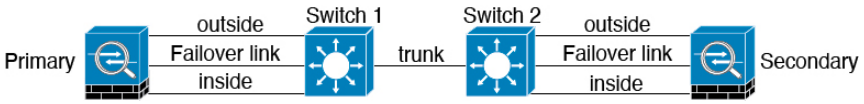


Figure 14: Connecting with a Double-Switch—Not Recommended



Scenario 2—Recommended

We recommend that failover links not use the same switch as the data interfaces. Instead, use a different switch or use a direct cable to connect the failover link, as shown in the following figures.

Figure 15: Connecting with a Different Switch

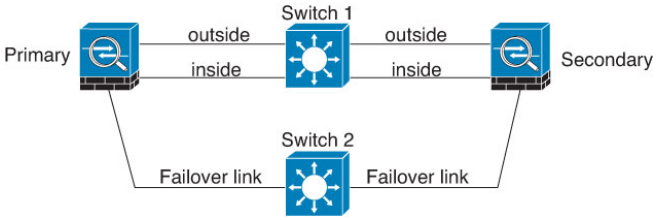
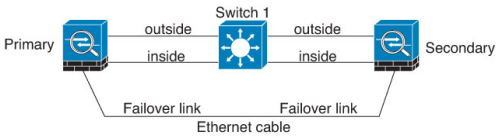


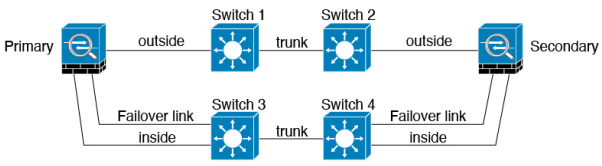
Figure 16: Connecting with a Cable



Scenario 3—Recommended

If the Firepower Threat Defense data interfaces are connected to more than one set of switches, then a failover link can be connected to one of the switches, preferably the switch on the secure (inside) side of network, as shown in the following figure.

Figure 17: Connecting with a Secure Switch



Scenario 4—Recommended

The most reliable failover configurations use a redundant interface on the failover link, as shown in the following figures.

Figure 18: Connecting with Redundant Interfaces

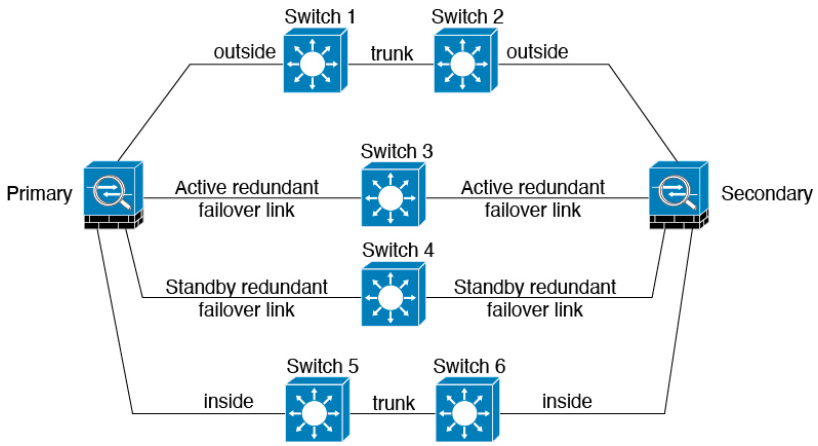
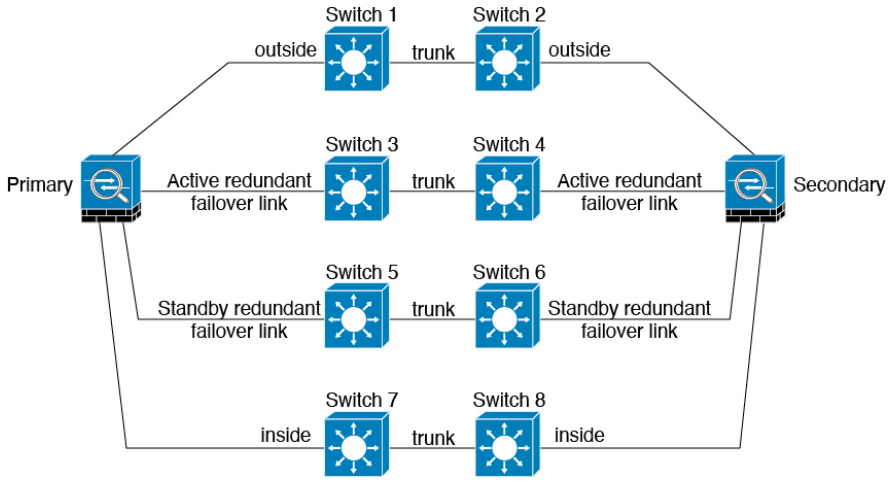


Figure 19: Connecting with Inter-switch Links



MAC Addresses and IP Addresses in High Availability

When you configure your interfaces, you can specify an active IP address and a standby IP address on the same network. Generally, when a failover occurs, the new active unit takes over the active IP addresses and MAC addresses. Because network devices see no change in the MAC to IP address pairing, no ARP entries change or time out anywhere on the network.

Note Although recommended, the standby address is not required. Without a standby IP address, the active unit cannot perform network tests to check the standby interface health; it can only track the link state. You also cannot connect to the standby unit on that interface for management purposes.

The IP address and MAC address for the state link do not change at failover.

Active/Standby IP Addresses and MAC Addresses

For Active/Standby High Availability, see the following for IP address and MAC address usage during a failover event:

1. The active unit always uses the primary unit's IP addresses and MAC addresses.
2. When the active unit fails over, the standby unit assumes the IP addresses and MAC addresses of the failed unit and begins passing traffic.
3. When the failed unit comes back online, it is now in a standby state and takes over the standby IP addresses and MAC addresses.

However, if the secondary unit boots without detecting the primary unit, then the secondary unit becomes the active unit and uses its own MAC addresses, because it does not know the primary unit MAC addresses. When the primary unit becomes available, the secondary (active) unit changes the MAC addresses to those of the primary unit, which can cause an interruption in your network traffic. Similarly, if you swap out the primary unit with new hardware, a new MAC address is used.

If you reload the standby unit with the failover configuration disabled, the standby unit boots up as the active unit and uses the primary unit's IP addresses and MAC addresses. This leads to duplicate IP addresses and causes network traffic disruptions. Use the command **configure high-availability resume** to enable failover and restore the traffic flow.

Virtual MAC addresses guard against this disruption, because the active MAC addresses are known to the secondary unit at startup, and remain the same in the case of new primary unit hardware. We recommend that you configure the virtual MAC address on both the primary and secondary units to ensure that the secondary unit uses the correct MAC addresses when it is the active unit, even if it comes online before the primary unit. If you do not configure virtual MAC addresses, you might need to clear the ARP tables on connected routers to restore traffic flow. The FTD device does not send gratuitous ARPs for static NAT addresses when the MAC address changes, so connected routers do not learn of the MAC address change for these addresses.

Virtual MAC Addresses

The FTD device has multiple methods to configure virtual MAC addresses. We recommend using only one method. If you set the MAC address using multiple methods, the MAC address used depends on many variables, and might not be predictable.

Stateful Failover

During Stateful Failover, the active unit continually passes per-connection state information to the standby unit. After a failover occurs, the same connection information is available at the new active unit. Supported end-user applications are not required to reconnect to keep the same communication session.

Supported Features

For Stateful Failover, the following state information is passed to the standby FTD device:

- NAT translation table.
- TCP and UDP connections and states, including HTTP connection states. Other types of IP protocols, and ICMP, are not parsed by the active unit, because they get established on the new active unit when a new packet arrives.
- Snort connection states, inspection results, and pin hole information, including strict TCP enforcement.

- The ARP table
- The Layer 2 bridge table (for bridge groups)
- The ISAKMP and IPsec SA table
- GTP PDP connection database
- SIP signaling sessions and pin holes.
- Static and dynamic routing tables—Stateful Failover participates in dynamic routing protocols, like OSPF and EIGRP, so routes that are learned through dynamic routing protocols on the active unit are maintained in a Routing Information Base (RIB) table on the standby unit. Upon a failover event, packets travel normally with minimal disruption to traffic because the active secondary unit initially has rules that mirror the primary unit. Immediately after failover, the re-convergence timer starts on the newly active unit. Then the epoch number for the RIB table increments. During re-convergence, OSPF and EIGRP routes become updated with a new epoch number. Once the timer is expired, stale route entries (determined by the epoch number) are removed from the table. The RIB then contains the newest routing protocol forwarding information on the newly active unit.



Note Routes are synchronized only for link-up or link-down events on an active unit. If the link goes up or down on the standby unit, dynamic routes sent from the active unit may be lost. This is normal, expected behavior.

- DHCP Server—DHCP address leases are not replicated. However, a DHCP server configured on an interface will send a ping to make sure an address is not being used before granting the address to a DHCP client, so there is no impact to the service. State information is not relevant for DHCP relay or DDNS.
- Access control policy decisions—Decisions related to traffic matching (including URL, URL category, geolocation, and so forth), intrusion detection, malware, and file type are preserved during failover. However, for connections being evaluated at the moment of failover, there are the following caveats:
  - AVC—App-ID verdicts are replicated, but not detection states. Proper synchronization occurs as long as the App-ID verdicts are complete and synchronized before failover occurs.
  - Intrusion detection state—Upon failover, once mid-flow pickup occurs, new inspections are completed, but old states are lost.
  - File malware blocking—The file disposition must become available before failover.
  - File type detection and blocking—The file type must be identified before failover. If failover occurs while the original active device is identifying the file, the file type is not synchronized. Even if your file policy blocks that file type, the new active device downloads the file.
- User identity decisions from the identity policy, including the user-to-IP address mappings gathered passively through the User Agent and ISE Session Directory, and active authentication through captive portal. Users who are actively authenticating at the moment of failover might be prompted to authenticate again.
- Network AMP—Cloud lookups are independent from each device, so failover does not affect this feature in general. Specifically:

- Signature Lookup—If failover occurs in the middle of a file transmission, no file event is generated and no detection occurs.
- File Storage—If failover occurs when the file is being stored, it is stored on the original active device. If the original active device went down while the file was being stored, the file does not get stored.
- File Pre-classification (Local Analysis)—If failover occurs in the middle of pre-classification, detection fails.
- File Dynamic Analysis (Connectivity to the cloud)—If failover occurs, the system might submit the file to the cloud.
- Archive File Support—If failover occurs in the middle of an analysis, the system loses visibility into the file/archive.
- Custom Blocking—If failover occurs, no events are generated.

- Security Intelligence decisions. However, DNS-based decisions that are in process at the moment of failover are not completed.
- From all the connections, only established ones will be replicated on the Standby ASA.

Unsupported Features

For Stateful Failover, the following state information is not passed to the standby FTD device:

- Sessions in plaintext tunnels such as GRE or IP-in-IP. Sessions inside tunnels are not replicated and the new active node will not be able to reuse existing inspection verdicts to match the correct policy rules.
- Decrypted TLS/SSL connections—The decryption states are not synchronized, and if the active unit fails, then decrypted connections will be reset. New connections will need to be established to the new active unit. Connections that are not decrypted (in other words, those that match a TLS/SSL Do Not Decrypt rule action) are not affected and are replicated correctly.
- Multicast routing.

Bridge Group Requirements for High Availability

There are special considerations for high availability when using bridge groups.

When the active unit fails over to the standby unit, the switch port running Spanning Tree Protocol (STP) can go into a blocking state for 30 to 50 seconds when it senses the topology change. To avoid traffic loss on the bridge group member interfaces while the port is in a blocking state, you can configure one of the following workarounds:

- Switch port is in Access mode—Enable the STP PortFast feature on the switch:

```
interface interface_id
  spanning-tree portfast
```

The PortFast feature immediately transitions the port into STP forwarding mode upon linkup. The port still participates in STP. So if the port is to be a part of the loop, the port eventually transitions into STP blocking mode.

- If the switch port is in Trunk mode, or you cannot enable STP PortFast, then you can use one of the following less desirable workarounds that impacts failover functionality or STP stability:
  - Disable interface monitoring on the bridge group and member interfaces.
  - Increase the interface hold time in the failover criteria to a high value that will allow STP to converge before the unit fails over.
  - Decrease the STP timers on the switch to allow STP to converge faster than the interface hold time.

Failover Health Monitoring

The Firepower Threat Defense device monitors each unit for overall health and for interface health. This section includes information about how the Firepower Threat Defense device performs tests to determine the state of each unit.

Unit Health Monitoring

The FTD device determines the health of the other unit by monitoring the failover link with hello messages. When a unit does not receive three consecutive hello messages on the failover link, the unit sends LANTEST messages on each data interface, including the failover link, to validate whether or not the peer is responsive. The action that the FTD device takes depends on the response from the other unit. See the following possible actions:

- If the FTD device receives a response on the failover link, then it does not fail over.
- If the FTD device does not receive a response on the failover link, but it does receive a response on a data interface, then the unit does not failover. The failover link is marked as failed. You should restore the failover link as soon as possible because the unit cannot fail over to the standby while the failover link is down.
- If the FTD device does not receive a response on any interface, then the standby unit switches to active mode and classifies the other unit as failed.

Interface Monitoring

When a unit does not receive hello messages on a monitored interface for 15 seconds, it runs interface tests. If one of the interface tests fails for an interface, but this same interface on the other unit continues to successfully pass traffic, then the interface is considered to be failed, and the device stops running tests.

If the threshold you define for the number of failed interfaces is met (see **Devices** > **Device Management** > **High Availability** > **Failover Trigger Criteria**), and the active unit has more failed interfaces than the standby unit, then a failover occurs. If an interface fails on both units, then both interfaces go into the "Unknown" state and do not count towards the failover limit defined by failover interface policy.

An interface becomes operational again if it receives any traffic. A failed device returns to standby mode if the interface failure threshold is no longer met.

If an interface has IPv4 and IPv6 addresses configured on it, the device uses the IPv4 addresses to perform the health monitoring. If an interface has only IPv6 addresses configured on it, then the device uses IPv6 neighbor discovery instead of ARP to perform the health monitoring tests. For the broadcast ping test, the device uses the IPv6 all nodes address (FE02::1).

Interface Tests

The Firepower Threat Defense device uses the following interface tests. The duration of each test is approximately 1.5 seconds.

1. Link Up/Down test—A test of the interface status. If the Link Up/Down test indicates that the interface is down, then the device considers it failed, and testing stops. If the status is Up, then the device performs the Network Activity test.
2. Network Activity test—A received network activity test. At the start of the test, each unit clears its received packet count for its interfaces. As soon as a unit receives any eligible packets during the test, then the interface is considered operational. If both units receive traffic, then testing stops. If one unit receives traffic and the other unit does not, then the interface on the unit that does not receive traffic is considered failed, and testing stops. If neither unit receives traffic, then the device starts the ARP test.
3. ARP test—A test for successful ARP replies. Each unit sends a single ARP request for the IP address in the most recent entry in its ARP table. If the unit receives an ARP reply or other network traffic during the test, then the interface is considered operational. If the unit does not receive an ARP reply, then the device sends a single ARP request for the IP address in the *next* entry in the ARP table. If the unit receives an ARP reply or other network traffic during the test, then the interface is considered operational. If both units receive traffic, then testing stops. If one unit receives traffic, and the other unit does not, then the interface on the unit that does not receive traffic is considered failed, and testing stops. If neither unit receives traffic, then the device starts the Broadcast Ping test.
4. Broadcast Ping test—A test for successful ping replies. Each unit sends a broadcast ping, and then counts all received packets. If the unit receives any packets during the test, then the interface is considered operational. If both units receive traffic, then testing stops. If one unit receives traffic, and the other unit does not, then the interface on the unit that does not receive traffic is considered failed, and testing stops. If neither unit receives traffic, then testing starts over again with the ARP test. If both units continue to receive no traffic from the ARP and Broadcast Ping tests, then these tests will continue running in perpetuity.

Interface Status

Monitored interfaces can have the following status:

- Unknown—Initial status. This status can also mean the status cannot be determined.
- Normal—The interface is receiving traffic.
- Normal (Waiting)—The interface is up, but has not yet received a hello packet from the corresponding interface on the peer unit.
- Normal (Not-Monitored)—The interface is up, but is not monitored by the failover process.
- Testing—Hello messages are not heard on the interface for five poll times.
- Link Down—The interface or VLAN is administratively down.
- Link Down (Waiting)—The interface or VLAN is administratively down and has not yet received a hello packet from the corresponding interface on the peer unit.
- Link Down (Not-Monitored)—The interface or VLAN is administratively down, but is not monitored by the failover process.
- No Link—The physical link for the interface is down.

- No Link (Waiting)—The physical link for the interface is down and has not yet received a hello packet from the corresponding interface on the peer unit.
- No Link (Not-Monitored)—The physical link for the interface is down, but is not monitored by the failover process.
- Failed—No traffic is received on the interface, yet traffic is heard on the peer interface.

Failover Triggers and Detection Timing

The following events trigger failover in a Firepower high availability pair:

- More than 50% of the Snort instances on the active unit are down.
- Disk space on the active unit is more than 90% full.
- The **no failover active** command is run on the active unit or the **failover active** command is run on the standby unit.
- The active unit has more failed interfaces than the standby unit.
- Interface failure on the active device exceeds the threshold configured.

  By default, failure of a single interface causes failover. You can change the default value by configuring a threshold for the number of interfaces or a percentage of monitored interfaces that must fail for the failover to occur. If the threshold breaches on the active device, failover occurs. If the threshold breaches on the standby device, the unit moves to **Fail** state.

  To change the default failover criteria, enter the following command in global configuration mode:

Table 63:

Command	Purpose
failover interface-policy num [%] `hostname (config)# failover interface-policy 20%`	Changes the default failover criteria. When specifying a specific number of interfaces, the *num* argument can be from 1 to 250. When specifying a percentage of interfaces, the *num* argument can be from 1 to 100.

The following table shows the failover triggering events and associated failure detection timing. If failover occurs, you can view the reason for the failover in the Message Center, along with various operations pertaining to the high availability pair. You can configure these thresholds to a value within the specified minimum-maximum range.

Table 64: FTD Failover Times

Failover Triggering Event	Minimum	Default	Maximum
Active unit loses power, hardware goes down, or the software reloads or crashes. When any of these occur, the monitored interfaces or failover link do not receives any hello message.	800 milliseconds	15 seconds	45 seconds

Failover Triggering Event	Minimum	Default	Maximum
Active unit interface physical link down.	500 milliseconds	5 seconds	15 seconds
Active unit interface up, but connection problem causes interface testing.	5 seconds	25 seconds	75 seconds

About Active/Standby Failover

Active/Standby failover lets you use a standby FTD device to take over the functionality of a failed unit. When the active unit fails, the standby unit becomes the active unit.

Primary/Secondary Roles and Active/Standby Status

When setting up Active/Standby failover, you configure one unit to be primary and the other to be secondary. During configuration, the primary unit's policies are synchronized to the secondary unit. At this point, the two units act as a single device for device and policy configuration. However, for events, dashboards, reports and health monitoring, they continue to display as separate devices.

The main differences between the two units in a failover pair are related to which unit is active and which unit is standby, namely which IP addresses to use and which unit actively passes traffic.

However, a few differences exist between the units based on which unit is primary (as specified in the configuration) and which unit is secondary:

- The primary unit always becomes the active unit if both units start up at the same time (and are of equal operational health).
- The primary unit MAC addresses are always coupled with the active IP addresses. The exception to this rule occurs when the secondary unit becomes active and cannot obtain the primary unit MAC addresses over the failover link. In this case, the secondary unit MAC addresses are used.

Active Unit Determination at Startup

The active unit is determined by the following:

- If a unit boots and detects a peer already running as active, it becomes the standby unit.
- If a unit boots and does not detect a peer, it becomes the active unit.
- If both units boot simultaneously, then the primary unit becomes the active unit, and the secondary unit becomes the standby unit.

Failover Events

In Active/Standby failover, failover occurs on a unit basis.

The following table shows the failover action for each failure event. For each failure event, the table shows the failover policy (failover or no failover), the action taken by the active unit, the action taken by the standby unit, and any special notes about the failover condition and actions.

Table 65: Failover Events

Failure Event	Policy	Active Unit Action	Standby Unit Action	Notes
Active unit failed (power or hardware)	Failover	n/a	Become active Mark active as failed	No hello messages are received on any monitored interface or the failover link.
Formerly active unit recovers	No failover	Become standby	No action	None.
Standby unit failed (power or hardware)	No failover	Mark standby as failed	n/a	When the standby unit is marked as failed, then the active unit does not attempt to fail over, even if the interface failure threshold is surpassed.
Failover link failed during operation	No failover	Mark failover link as failed	Mark failover link as failed	You should restore the failover link as soon as possible because the unit cannot fail over to the standby unit while the failover link is down.
Failover link failed at startup	No failover	Become active Mark failover link as failed	Become active Mark failover link as failed	If the failover link is down at startup, both units become active.
State link failed	No failover	No action	No action	State information becomes out of date, and sessions are terminated if a failover occurs.
Interface failure on active unit above threshold	Failover	Mark active as failed	Become active	None.
Interface failure on standby unit above threshold	No failover	No action	Mark standby as failed	When the standby unit is marked as failed, then the active unit does not attempt to fail over even if the interface failure threshold is surpassed.

Requirements and Prerequisites for High Availability

Model Support

FTD

Supported Domains

Any

User Roles

Admin

Network Admin

Guidelines for High Availability

Model Support

- Firepower 9300—Intra-chassis High Availability is not supported.
- The FTDv on public cloud networks such as Microsoft Azure and Amazon Web Services are not supported with High Availability because Layer 2 connectivity is required.

Additional Guidelines

- When the active unit fails over to the standby unit, the connected switch port running Spanning Tree Protocol (STP) can go into a blocking state for 30 to 50 seconds when it senses the topology change. To avoid traffic loss while the port is in a blocking state, you can enable the STP PortFast feature on the switch:

  interface *interface_id* **spanning-tree portfast**

  This workaround applies to switches connected to both routed mode and bridge group interfaces. The PortFast feature immediately transitions the port into STP forwarding mode upon linkup. The port still participates in STP. So if the port is to be a part of the loop, the port eventually transitions into STP blocking mode.
- Configuring port security on the switches connected to the FTD device failover pair can cause communication problems when a failover event occurs. This problem occurs when a secure MAC address configured or learned on one secure port moves to another secure port, a violation is flagged by the switch port security feature.
- For Active/Standby High Availability and a VPN IPsec tunnel, you cannot monitor both the active and standby units using SNMP over the VPN tunnel. The standby unit does not have an active VPN tunnel, and will drop traffic destined for the NMS. You can instead use SNMPv3 with encryption so the IPsec tunnel is not required.
- Make sure each unit in the high-availability pair uses a unique hostname; the FMC cannot add the secondary unit if it has the same name as the primary unit.
- Immediately after failover, the source address of syslog messages will be the failover interface address for a few seconds.
- For better convergence (during a failover), you must shut down the interfaces on a HA pair that are not associated with any configuration or instance.
- When using SNMPv3 with failover, if you replace a failover unit, then SNMPv3 users are not replicated to the new unit. You must remove the users, re-add them, and then redeploy your configuration to force the users to replicate to the new unit.
- If you have a very large number of access control and NAT rules, the size of the configuration can prevent efficient configuration replication, resulting in the standby unit taking an excessively long time to reach

standby ready state. This can also impact your ability to connect to the standby unit during replication through the console or SSH session. To enhance configuration replication performance, enable transactional commit for both access rules and NAT, using the **asp rule-engine transactional-commit access-group** and **asp rule-engine transactional-commit nat** commands.

- A unit in a High Availability pair transitioning to the standby role synchronizes its clock with the active unit.

  Example:

```
firepower#show clock
01:00:52 UTC Mar 1 2022

...
01:01:18 UTC Mar 1 2022 <======= Incorrect (previous) clock
Cold Standby               Sync Config               Detected an Active mate

19:38:21 UTC Apr 9 2022 <======= Updated clock
Sync Config                Sync File System          Detected an Active mate
...
firepower/sec/stby#show clock
19:38:40 UTC Apr 9 2022
```

- The units in High Availability do not dynamically synchronize the clock. Here are some examples of events when synchronization takes place:
  - A new High Availability pair is created.
  - High Availability is broken and re-created.
  - Communication over the failover link was disrupted and reestablished.
  - Failover status was manually changed at the CLI using the **no failover/failover** or **configure high-availability suspend/resume** (FTD) commands.
- Enabling High Availability forces all routes to be deleted and are re-added after the High Availability progression changes to the Active state. You could experience connection loss during this phase.
- If you replace the primary unit, then when you re-create high-availability, you should set the replacement unit as the *secondary* unit so that the configurations are replicated from the former secondary unit to the replacement unit. If you set the replacement unit as primary, you will accidentally overwrite the configuration that is present on the operational unit.
- Deploying Firepower 1100 and 2100 devices in high availability with hundreds of interfaces configured on them can result in increased delay in the failover time (seconds).
- In the High Availability configuration, short-lived connections, usually using port 53, are closed quickly and never transferred or synchronized from Active to Standby, so there might be a difference in the number of connections on both High Availability devices. This is expected behavior for short-lived connections. You can try to compare the connections that are long-lived (for example, more than 30-60 seconds).

Add a Firepower Threat Defense High Availability Pair

When establishing an Active/Standby High Availability pair, you designate one of the devices as primary and the other as secondary. The system applies a merged configuration to the paired devices. If there is a conflict, the system applies the configuration from the device you designated as primary.

In a multidomain deployment, devices in a high availability pair must belong to the same domain.

Note The system uses the failover link to sync configuration, while the stateful failover link is used to sync application content between peers. The failover link and the stateful failover link are in a private IP space and are only used for communication between peers in a high availability pair.After high availability is established, selected interface links and encryption settings cannot be modified without breaking the high availability pair and reconfiguring it.

Caution Creating or breaking a Firepower Threat Defense high availability pair immediately restarts the Snort process on the primary and secondary devices, temporarily interrupting traffic inspection on both devices. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Before you begin

Confirm that both devices:

- Are the same model.
- Have the same number and type of interfaces.
- Are in the same domain and group.
- Have normal health status and are running the same software.
- Are either in routed or transparent mode.
- Have the same NTP configuration. See Configure NTP Time Synchronization for Threat Defense, on page 976.
- Are fully deployed with no uncommitted changes.
- Do not have DHCP or PPPoE configured in any of their interfaces.

Procedure

Step 1 Add both devices to the Firepower Management Center according to Add a Device to the FMC, on page 229.

Step 2 Choose **Devices** > **Device Management**.

Step 3 From the **Add** drop-down menu, choose **High Availability**.

Step 4 Enter a display **Name** for the high availability pair.

Step 5 Under **Device Type**, choose **Firepower Threat Defense**.

Step 6 Choose the **Primary Peer** device for the high availability pair.

Step 7 Choose the **Secondary Peer** device for the high availability pair.

Step 8 Click **Continue**.

Step 9 Under LAN Failover Link, choose an **Interface** with enough bandwidth to reserve for failover communications.

Note Only interfaces that do not have a logical name and do not belong to a security zone,will be listed in the **Interface** drop-down in the **Add High Availability Pair** dialog.

Step 10 Type any identifying **Logical Name**.

Step 11 Type a **Primary IP** address for the failover link on the active unit.

This address should be on an unused subnet.

Note 169.254.0.0/16 and fd00:0:0:*::/64 are internally used subnets and cannot be used for the failover or state links.

Step 12 Optionally, choose **Use IPv6 Address**.

Step 13 Type a **Secondary IP** address for the failover link on the standby unit. This IP address must be in the same subnet as the primary IP address.

Step 14 If IPv4 addresses are used, type a **Subnet Mask** that applies to both the primary and secondary IP addresses.

Step 15 Optionally, under Stateful Failover Link, choose the same **Interface**, or choose a different interface and enter the high availability configuration information.

Note 169.254.0.0/16 and fd00:0:0:*::/64 are internally used subnets and cannot be used for the failover or state links.

Step 16 Optionally, choose **Enabled** and choose the **Key Generation** method for IPsec Encryption between the failover links.

Step 17 Click **OK**. This process takes a few minutes as the process synchronizes system data.

Configure Optional High Availability Parameters

You can view the initial High Availability Configuration on the Firepower Management Center. You cannot edit these settings without breaking the high availability pair and then re-establishing it.

You can edit the Failover Trigger Criteria to improve failover results. Interface Monitoring allows you to determine which interfaces are better suited for failover.

Configure Standby IP Addresses and Interface Monitoring

For each interface, set a standby IP address. Although recommended, the standby address is not required. Without a standby IP address, the active unit cannot perform network tests to check the standby interface health; it can only track the link state.

By default, monitoring is enabled on all physical interfaces with logical names configured. You might want to exclude interfaces attached to less critical networks from affecting your failover policy.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device high-availability pair you want to edit, click the **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **High Availability** tab.

Step 4 In the **Monitored Interfaces** area, click the **Edit** ( ) next to the interface you want to edit.

Step 5 Check the **Monitor this interface for failures** check box.

Step 6 On the **IPv4** tab, enter the **Standby IP Address**.

This address must be a free address on the same network as the active IP address.

Step 7 If you configured the IPv6 address manually, on the **IPv6** tab, click the **Edit** ( ) next to the active IP address, enter the **Standby IP Address**, and click **OK**.

This address must be a free address on the same network as the active IP address. For autogenerated and **Enforce EUI 64** addresses, the standby address is automatically generated.

Step 8 Click **OK**.

Edit High Availability Failover Criteria

You can customize failover criteria based on your network deployment.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device high-availability pair you want to edit, click the **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Choose **High Availability**.

Step 4 Next to **Failover Trigger Criteria**, click the **Edit** ( ).

Step 5 Under **Interface Failure Threshold**, choose the number or percentage of interfaces that must fail before the device fails over.

Step 6 Under **Hello packet Intervals**, choose how often hello packets are sent over the failover link.

Note If you use remote access VPN on the Firepower 2100, use the default hello packet intervals. Otherwise, you might see high CPU usage that can cause a failover to occur.

Step 7 Click **OK**.

Configure Virtual MAC addresses

You can configure active and standby MAC addresses for failover in two places on the Firepower Management Center:

- The Advanced tab of the Edit Interface page during interface configuration; see Configure the MAC Address, on page 646.
- The Add Interface MAC Address page accessed from the High Availability page; see

If active and standby MAC addresses are configured in both locations, the addresses defined during interface configuration takes preference for failover.

You can minimize loss of traffic during failover by designating active and standby mac addresses to the physical interface. This feature offers redundancy against IP address mapping for failover.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device high-availability pair you want to edit, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Choose **High Availability**.

Step 4 Choose **Add** ( )next to Interface Mac Addresses.

Step 5 Choose a **Physical Interface**.

Step 6 Type an **Active Interface Mac Address**.

Step 7 Type a **Standby Interface Mac Address**.

Step 8 Click **OK**.

Manage High Availability

This section describes how to manage High Availability units after you enable High Availability, including how to change the High Availability setup and how to force failover from one unit to another.

Switch the Active Peer in a Firepower Threat Defense High Availability Pair

After you establish a Firepower Threat Defense high availability pair, you can manually switch the active and standby units, effectively forcing failover for reasons such as persistent fault or health events on the current active unit. Both units should be fully deployed before you complete this procedure.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the high availability pair where you want to change the active peer, click the **Switch Active Peer**.

Step 3 You can:

- Click **Yes** to immediately make the standby device the active device in the high availability pair.
- Click **No** to cancel and return to the Device Management page.

Suspend and Resume High Availability

You can suspend a unit in a high availability pair. This is useful when:

- Both units are in an active-active situation and fixing the communication on the failover link does not correct the problem.
- You want to troubleshoot an active or standby unit and do not want the units to fail over during that time.

When you suspend high availability, you stop the pair of devices from behaving as a failover unit. The currently active device remains active, handling all user connections. However, failover criteria are no longer monitored, and the system will never fail over to the now pseudo-standby device. The standby device will retain its configuration, but it will remain inactive.

The key difference between suspending HA and breaking HA is that on a suspended HA device, the high availability configuration is retained. When you break HA, the configuration is erased. Thus, you have the option to resume HA on a suspended system, which enables the existing configuration and makes the two devices function as a failover pair again.

To suspend HA, use the **configure high-availability suspend** command.

```
> configure high-availability suspend
Please ensure that no deployment operation is in progress before suspending
high-availability.
Please enter 'YES' to continue if there is no deployment operation in
progress and 'NO' if you wish to abort: YES
Successfully suspended high-availability.
```

If you suspend high availability from the active unit, the configuration is suspended on both the active and standby unit. If you suspend it from the standby unit, it is suspended on the standby unit only, but the active unit will not attempt to fail over to a suspended unit.

To resume failover, use the **configure high-availability resume** command.

```
> configure high-availability resume
Successfully resumed high-availablity.
```

You can resume a unit only if it is in Suspended state. The unit will negotiate active/standby status with the peer unit.



Note Suspending high availability is a temporary state. If you reload a unit, it resumes the high-availability configuration automatically and negotiates the active/standby state with the peer.

Replace a Unit in an FTD High Availability Pair

You must break high availability before you replace a failed unit in a Firepower Threat Defense high availability pair. Then, register the replacement device to the Firepower Management Center and reestablish high availability. The process varies depending on whether the device is primary or secondary:

- Replace a Primary FTD HA Unit, on page 704
- Replace a Secondary FTD HA Unit, on page 704

Replace a Primary FTD HA Unit

Follow the steps below to replace a failed primary unit in a Firepower Threat Defense high availability pair. Failing to follow these steps can overwrite the existing high availability configuration.



Caution Creating or breaking a Firepower Threat Defense high availability pair immediately restarts the Snort process on the primary and secondary devices, temporarily interrupting traffic inspection on both devices. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Procedure

Step 1 Choose **Force Break** to separate the high availability pair; see Separate Units in a High Availability Pair, on page 705.

Note The break operation removes all the configuration related to HA from Firepower Threat Defense and Firepower Management Center, and you need to recreate it manually later. To successfully configure the same HA pair, ensure that you save the IPs, MAC addresses, and monitoring configuration of all the interfaces/subinterfaces prior to executing the HA break operation.

Step 2 Unregister the failed primary Firepower Threat Defense device from the Firepower Management Center; see Delete a Device from the FMC, on page 231.

Step 3 Register the replacement Firepower Threat Defense to the Firepower Management Center; see Add a Device to the FMC, on page 229.

Step 4 Configure high availability, using the existing secondary/active unit as the primary device and the replacement device as the secondary/standby device during registration; see Add a Firepower Threat Defense High Availability Pair, on page 699.

Replace a Secondary FTD HA Unit

Follow the steps below to replace a failed secondary unit in a Firepower Threat Defense high availability pair.

Caution Creating or breaking a Firepower Threat Defense high availability pair immediately restarts the Snort process on the primary and secondary devices, temporarily interrupting traffic inspection on both devices. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Procedure

Step 1 Choose **Force Break** to separate the high availability pair; see Separate Units in a High Availability Pair, on page 705.

Note The break operation removes all the configuration related to HA from Firepower Threat Defense and Firepower Management Center, and you need to recreate it manually later. To successfully configure the same HA pair, ensure that you save the IPs, MAC addresses, and monitoring configuration of all the interfaces/subinterfaces prior to executing the HA break operation.

Step 2 Unregister the secondary Firepower Threat Defense device from the Firepower Management Center; see Delete a Device from the FMC, on page 231.

Step 3 Register the replacement Firepower Threat Defense to the Firepower Management Center; see Add a Device to the FMC, on page 229.

Step 4 Configure high availability, using the existing primary/active unit as the primary device and the replacement device as the secondary/standby device during registration; see Add a Firepower Threat Defense High Availability Pair, on page 699.

Separate Units in a High Availability Pair

When you break a high availability pair, the active device retains full deployed functionality. The standby device loses its failover and interface configurations, and becomes a standalone device.When you a break a high availability pair, policies that were yet to be deployed to the active device, prior to the break operation are automatically deployed to the active device when the break operation is completed.

Tip An exception to this is the FlexConfig policy. A FlexConfig policy deployed on the active device may show a deployment failure after the break HA operation. You must alter and re-deploy the FlexConfig policy on the active device.

Note If you cannot reach the high availability pair using the Firepower Management Center, use the CLI command **configure high-availability disable** to remove the failover configuration from both devices.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the high-availability pair you want to break, click the **Break HA**.

Step 3 Optionally, check the check box to force break, if the standby peer does not respond.

Step 4 Click **Yes**. The device high-availability pair is separated.

The Break operation removes the failover configuration from the active and standby devices.

What to do next

(Optional) If you are using a flex-config policy on the active device, alter and re-deploy the flex-config policy to eliminate deployment errors.

Unregister a High Availability Pair

You can delete the pair from the Firepower Management Center and disable High Availability on each unit using the CLI.

Before you begin

This procedure requires CLI access.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the high-availability pair you want to unregister, click **Delete** ().

Step 3 Click **Yes**. The device high availability pair is deleted.

Step 4 On each unit, access the Firepower Threat Defense CLI, and enter the following command:

configure high-availability disable

If you do not enter this command, you cannot re-register the units and form a new HA pair.

Note Enter this command *before* you change the firewall mode; if you change the mode, the unit will not later let you enter the **configure high-availability disable** command, and the Firepower Management Center cannot re-form the HA pair without this command.

Monitoring High Availability

This section lets you monitor the High Availability status.

View Failover History

You can view the failover history of both high availability devices in a single view. The history displays in chronological order and includes the reason for any failover.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device high-availability pair you want to edit, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Choose **Summary**.

Step 4 Under General, click **View** ( ).

View Stateful Failover Statistics

You can view the stateful failover link statistics of both the primary and secondary devices in the high availability pair.

Procedure

Step 1 Choose **Devices** > **Device Management**.

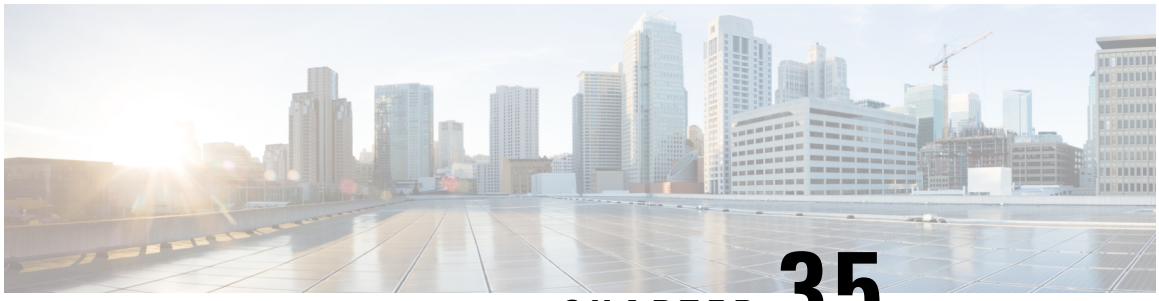
Step 2 Next to the device high-availability pair you want to edit, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Choose **High Availability**.

Step 4 Under Stateful Failover Link, click **View** ( ).

Step 5 Choose a device to view statistics.



CHAPTER 35

Clustering for the Firepower Threat Defense

Clustering lets you group multiple Firepower Threat Defense units together as a single logical device. Clustering is only supported for the Firepower Threat Defense device on the Firepower 9300 and the Firepower 4100 series. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices.

Note Some features are not supported when using clustering. See Unsupported Features with Clustering, on page 740.

- About Clustering on the Firepower 4100/9300 Chassis, on page 709
- Licenses for Clustering, on page 713
- Requirements and Prerequisites for Clustering, on page 714
- Clustering Guidelines and Limitations, on page 715
- Configure Clustering, on page 718
- FXOS: Remove a Cluster Node, on page 730
- FMC: Manage Cluster Members, on page 731
- FMC: Monitoring the Cluster, on page 734
- Examples for Clustering, on page 734
- Reference for Clustering, on page 739
- History for Clustering, on page 749

About Clustering on the Firepower 4100/9300 Chassis

When you deploy a cluster on the Firepower 4100/9300 chassis, it does the following:

- Creates a *cluster-control link* (by default, port-channel 48) for node-to-node communication.

  For a cluster isolated to security modules within one Firepower 9300 chassis, this link utilizes the Firepower 9300 backplane for cluster communications.

  For clustering with multiple chassis, you need to manually assign physical interface(s) to this EtherChannel for communications between chassis.

- Creates the cluster bootstrap configuration within the application.

When you deploy the cluster, the chassis supervisor pushes a minimal bootstrap configuration to each unit that includes the cluster name, cluster control link interface, and other cluster settings.

- Assigns data interfaces to the cluster as *Spanned* interfaces.

  For a cluster isolated to security modules within one Firepower 9300 chassis, spanned interfaces are not limited to EtherChannels, like it is for clustering with multiple chassis. The Firepower 9300 supervisor uses EtherChannel technology internally to load-balance traffic to multiple modules on a shared interface, so any data interface type works for Spanned mode. For clustering with multiple chassis, you must use Spanned EtherChannels for all data interfaces.



Note Individual interfaces are not supported, with the exception of a management interface.

- Assigns a management interface to all units in the cluster.

See the following sections for more information about clustering.

Bootstrap Configuration

When you deploy the cluster, the Firepower 9300 chassis supervisor pushes a minimal bootstrap configuration to each unit that includes the cluster name, cluster control link interface, and other cluster settings.

Cluster Members

Cluster members work together to accomplish the sharing of the security policy and traffic flows.

One member of the cluster is the **control** unit. The control unit is determined automatically. All other members are **data** units.

You must perform all configuration on the control unit only; the configuration is then replicated to the data units.

Some features do not scale in a cluster, and the control unit handles all traffic for those features. .

Cluster Control Link

The cluster control link is automatically created using the Port-channel 48 interface.

For a cluster isolated to security modules within one Firepower 9300 chassis, this interface has no member interfaces. This Cluster type EtherChannel utilizes the Firepower 9300 backplane for cluster communications. For clustering with multiple chassis, you must add one or more interfaces to the EtherChannel.

For a cluster with two chassis, do not directly connect the cluster control link from one chassis to the other chassis. If you directly connect the interfaces, then when one unit fails, the cluster control link fails, and thus the remaining healthy unit fails. If you connect the cluster control link through a switch, then the cluster control link remains up for the healthy unit.

Cluster control link traffic includes both control and data traffic.

Size the Cluster Control Link

If possible, you should size the cluster control link to match the expected throughput of each chassis so the cluster control link can handle the worst-case scenarios.

Cluster control link traffic is comprised mainly of state update and forwarded packets. The amount of traffic at any given time on the cluster control link varies. The amount of forwarded traffic depends on the load-balancing efficacy or whether there is a lot of traffic for centralized features. For example:

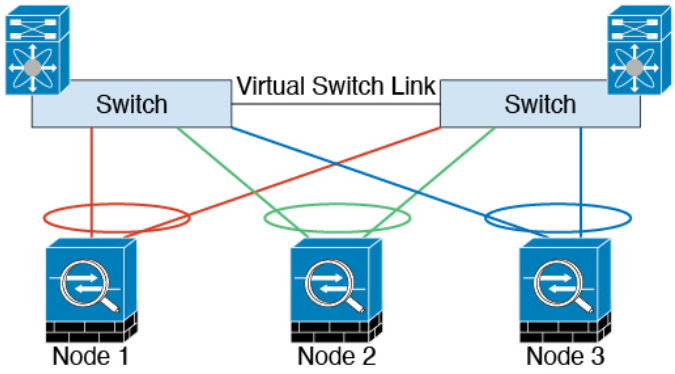
- NAT results in poor load balancing of connections, and the need to rebalance all returning traffic to the correct units.
- When membership changes, the cluster needs to rebalance a large number of connections, thus temporarily using a large amount of cluster control link bandwidth.

A higher-bandwidth cluster control link helps the cluster to converge faster when there are membership changes and prevents throughput bottlenecks.

Note If your cluster has large amounts of asymmetric (rebalanced) traffic, then you should increase the cluster control link size.

Cluster Control Link Redundancy

The following diagram shows how to use an EtherChannel as a cluster control link in a Virtual Switching System (VSS), Virtual Port Channel (vPC), StackWise, or StackWise Virtual environment. All links in the EtherChannel are active. When the switch is part of a redundant system, then you can connect firewall interfaces within the same EtherChannel to separate switches in the redundant system. The switch interfaces are members of the same EtherChannel port-channel interface, because the separate switches act like a single switch. Note that this EtherChannel is device-local, not a Spanned EtherChannel.



Cluster Control Link Reliability for Inter-Chassis Clustering

To ensure cluster control link functionality, be sure the round-trip time (RTT) between units is less than 20 ms. This maximum latency enhances compatibility with cluster members installed at different geographical sites. To check your latency, perform a ping on the cluster control link between units.

The cluster control link must be reliable, with no out-of-order or dropped packets; for example, for inter-site deployment, you should use a dedicated link.

Cluster Control Link Network

The Firepower 4100/9300 chassis auto-generates the cluster control link interface IP address for each unit based on the chassis ID and slot ID: 127.2.*chassis_id*.*slot_id*. You cannot set this IP address manually, either in FXOS or within the application. The cluster control link network cannot include any routers between units; only Layer 2 switching is allowed.

Management Network

We recommend connecting all units to a single management network. This network is separate from the cluster control link.

Management Interface

You must assign a Management type interface to the cluster. This interface is a special individual interface as opposed to a Spanned interface. The management interface lets you connect directly to each unit. This Management logical interface is separate from the other interfaces on the device. It is used to set up and register the device to the Firepower Management Center. It uses its own local authentication, IP address, and static routing. Each cluster member uses a separate IP address on the management network that you set as part of the bootstrap configuration.

The management interface is shared between the Management logical interface and the *Diagnostic* logical interface. The Diagnostic logical interface is optional and is not configured as part of the bootstrap configuration. The Diagnostic interface can be configured along with the rest of the data interfaces. If you choose to configure the Diagnostic interface, configure a Main cluster IP address as a fixed address for the cluster that always belongs to the current control unit. You also configure a range of addresses so that each unit, including the current control unit, can use a Local address from the range. The Main cluster IP address provides consistent diagnostic access to an address; when a control unit changes, the Main cluster IP address moves to the new control unit, so access to the cluster continues seamlessly. For outbound management traffic such as TFTP or syslog, each unit, including the control unit, uses the Local IP address to connect to the server.

Cluster Interfaces

For a cluster isolated to security modules within one Firepower 9300 chassis, you can assign both physical interfaces or EtherChannels (also known as port channels) to the cluster. Interfaces assigned to the cluster are Spanned interfaces that load-balance traffic across all members of the cluster.

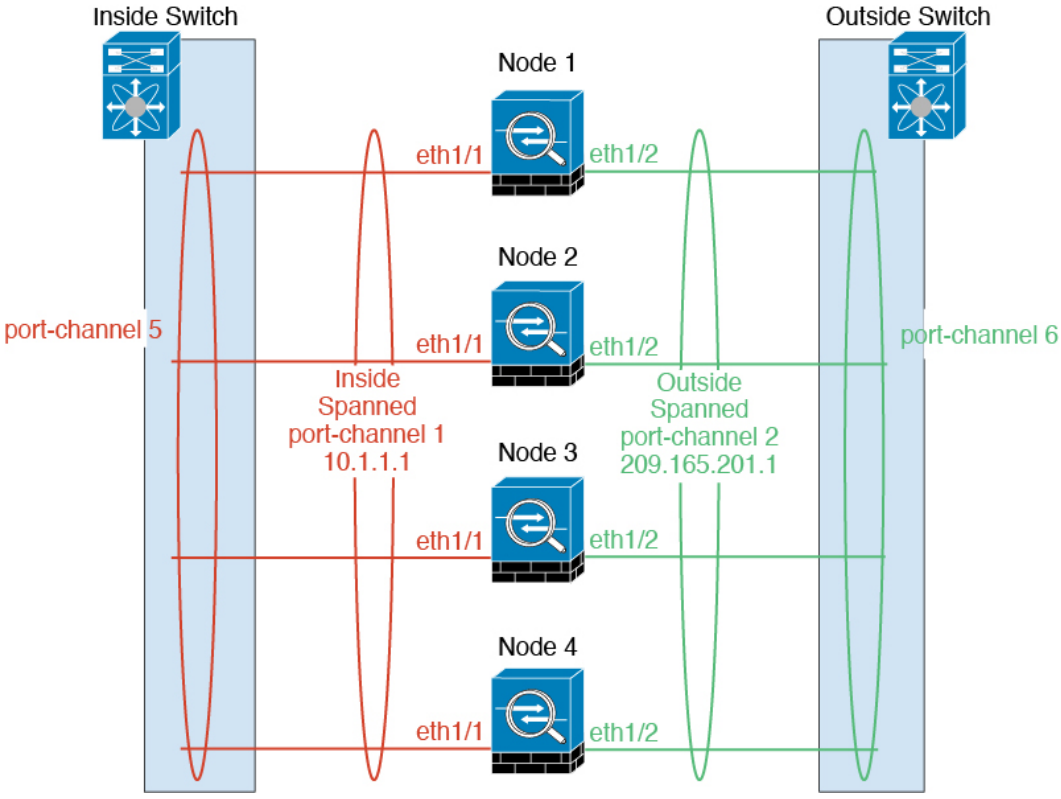
For clustering with multiple chassis, you can only assign data EtherChannels to the cluster. These Spanned EtherChannels include the same member interfaces on each chassis; on the upstream switch, all of these interfaces are included in a single EtherChannel, so the switch does not know that it is connected to multiple devices.

Individual interfaces are not supported, with the exception of a management interface.

Spanned EtherChannels

You can group one or more interfaces per chassis into an EtherChannel that spans all chassis in the cluster. The EtherChannel aggregates the traffic across all the available active interfaces in the channel. A Spanned EtherChannel can be configured in both routed and transparent firewall modes. In routed mode, the EtherChannel is configured as a routed interface with a single IP address. In transparent mode, the IP address

is assigned to the BVI, not to the bridge group member interface. The EtherChannel inherently provides load balancing as part of basic operation.



Connecting to a Redundant Switch System

We recommend connecting EtherChannels to a redundant switch system such as a VSS, vPC, StackWise, or StackWise Virtual system to provide redundancy for your interfaces.

Configuration Replication

All nodes in the cluster share a single configuration. You can only make configuration changes on the control node (with the exception of the bootstrap configuration), and changes are automatically synced to all other nodes in the cluster.

Licenses for Clustering

The Firepower Threat Defense uses Smart Licensing. You assign licenses to the cluster as a whole, not to individual units. However, each unit of the cluster consumes a separate license for each feature. The clustering feature itself does not require any licenses.

When you add a cluster member to the FMC, you can specify the feature licenses you want to use for the cluster. If you choose different feature licenses for each device before you form them into a cluster in the FMC, then the licenses you chose for the control unit are used for the cluster. You can modify licenses for the cluster in the **Devices** > **Device Management** > **Cluster** > **License** area.

Note If you add the cluster before the FMC is licensed (and running in Evaluation mode), then when you license the FMC, you can experience traffic disruption when you deploy policy changes to the cluster. Changing to licensed mode causes all data units to leave the cluster and then rejoin.

Requirements and Prerequisites for Clustering

Cluster Model Support

The FTD supports clustering on the following models:

- Firepower 9300—You can include up to 6 units in the cluster. For example, you can use 1 module in 6 chassis, or 2 modules in 3 chassis, or any combination that provides a maximum of 6 modules. Supports intra-chassis and inter-chassis clustering.
- Firepower 4100 series—Supported for up to 6 units using inter-chassis clustering.

User Roles

- Admin
- Access Admin
- Network Admin

Inter-Chassis Clustering Hardware and Software Requirements

All chassis in a cluster:

- For the Firepower 4100: All chassis must be the same model. For the Firepower 9300: All security modules must be the same type. For example, if you use clustering, all modules in the Firepower 9300 must be SM-40s. You can have different quantities of installed security modules in each chassis, although all modules present in the chassis must belong to the cluster including any empty slots.
- Must run the identical FXOS and application software except at the time of an image upgrade. Mismatched software versions can lead to poor performance, so be sure to upgrade all nodes in the same maintenance window.
- Must include the same interface configuration for interfaces you assign to the cluster, such as the same Management interface, EtherChannels, active interfaces, speed and duplex, and so on. You can use different network module types on the chassis as long as the capacity matches for the same interface IDs and interfaces can successfully bundle in the same spanned EtherChannel. Note that all data interfaces must be EtherChannels in clusters with multiple chassis. If you change the interfaces in FXOS after you enable clustering (by adding or removing interface modules, or configuring EtherChannels, for example), then perform the same changes on each chassis, starting with the data nodes, and ending with the control node.
- Must use the same NTP server. For Firepower Threat Defense, the FMC must also use the same NTP server. Do not set the time manually.

Switch Requirements for Inter-Chassis Clustering

- Be sure to complete the switch configuration and successfully connect all the EtherChannels from the chassis to the switch(es) before you configure clustering on the Firepower 4100/9300 chassis.
- For supported switch characteristics, see Cisco FXOS Compatibility.

Clustering Guidelines and Limitations

Switches for Clustering

- Make sure connected switches match the MTU for both cluster data interfaces and the cluster control link interface. You should configure the cluster control link interface MTU to be at least 100 bytes higher than the data interface MTU, so make sure to configure the cluster control link connecting switch appropriately. Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead.
- For Cisco IOS XR systems, if you want to set a non-default MTU, set the IOS XR interface MTU to be 14 bytes higher than the cluster device MTU. Otherwise, OSPF adjacency peering attempts may fail unless the **mtu-ignore** option is used. Note that the cluster device MTU should match the IOS XR *IPv4* MTU. This adjustment is not required for Cisco Catalyst and Cisco Nexus switches.
- On the switch(es) for the cluster control link interfaces, you can optionally enable Spanning Tree PortFast on the switch ports connected to the cluster unit to speed up the join process for new units.
- On the switch, we recommend that you use one of the following EtherChannel load-balancing algorithms: **source-dest-ip** or **source-dest-ip-port** (see the Cisco Nexus OS and Cisco IOS-XE **port-channel load-balance** command). Do not use a **vlan** keyword in the load-balance algorithm because it can cause unevenly distributed traffic to the devices in a cluster.
- If you change the load-balancing algorithm of the EtherChannel on the switch, the EtherChannel interface on the switch temporarily stops forwarding traffic, and the Spanning Tree Protocol restarts. There will be a delay before traffic starts flowing again.
- Switches on the cluster control link path should not verify the L4 checksum. Redirected traffic over the cluster control link does not have a correct L4 checksum. Switches that verify the L4 checksum could cause traffic to be dropped.
- Port-channel bundling downtime should not exceed the configured keepalive interval.
- On Supervisor 2T EtherChannels, the default hash distribution algorithm is adaptive. To avoid asymmetric traffic in a VSS design, change the hash algorithm on the port-channel connected to the cluster device to fixed:

  router(config)# **port-channel** *id* **hash-distribution fixed**

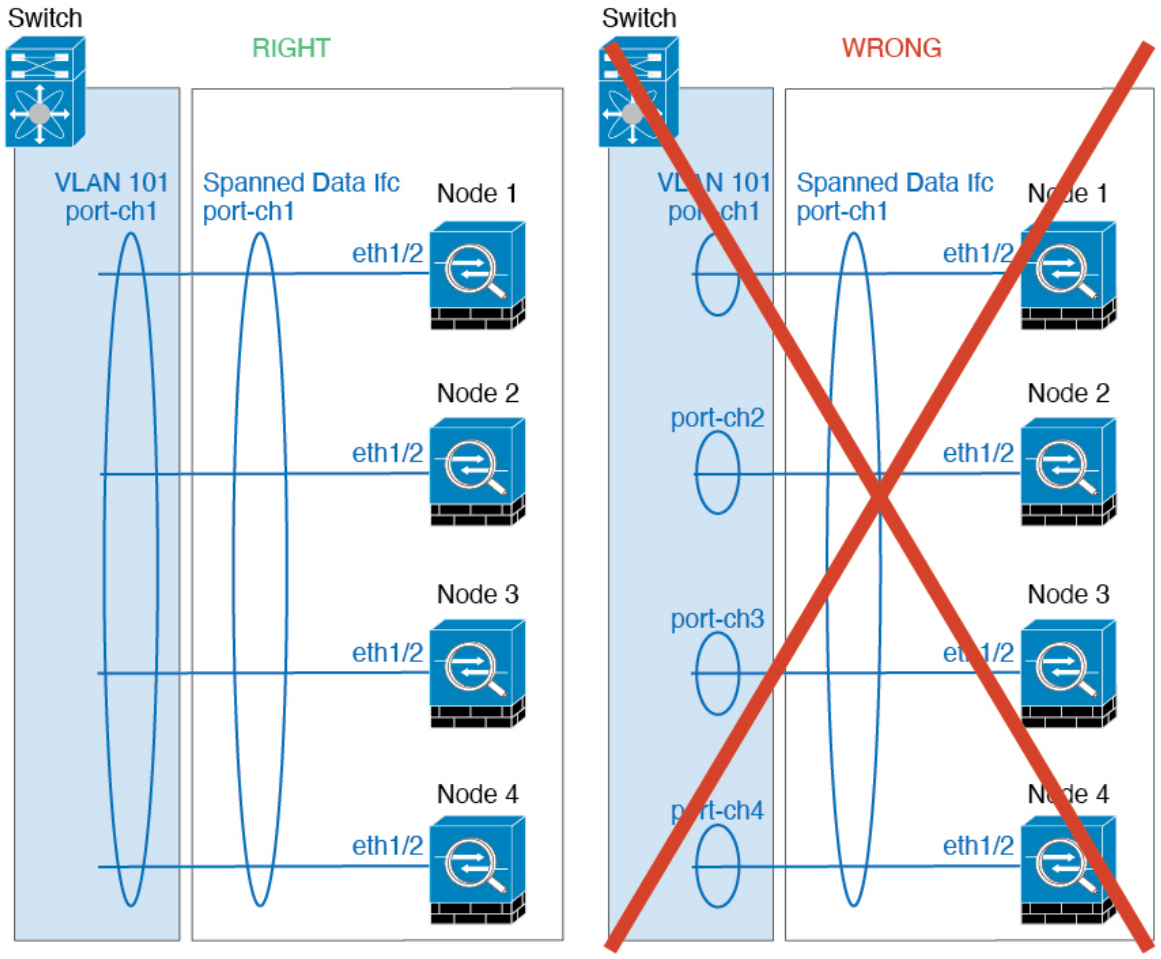
  Do not change the algorithm globally; you may want to take advantage of the adaptive algorithm for the VSS peer link.

- Firepower 4100/9300 clusters support LACP graceful convergence. So you can leave LACP graceful convergence enabled on connected Cisco Nexus switches.
- When you see slow bundling of a Spanned EtherChannel on the switch, you can enable LACP rate fast for an individual interface on the switch. FXOS EtherChannels have the LACP rate set to fast by default.

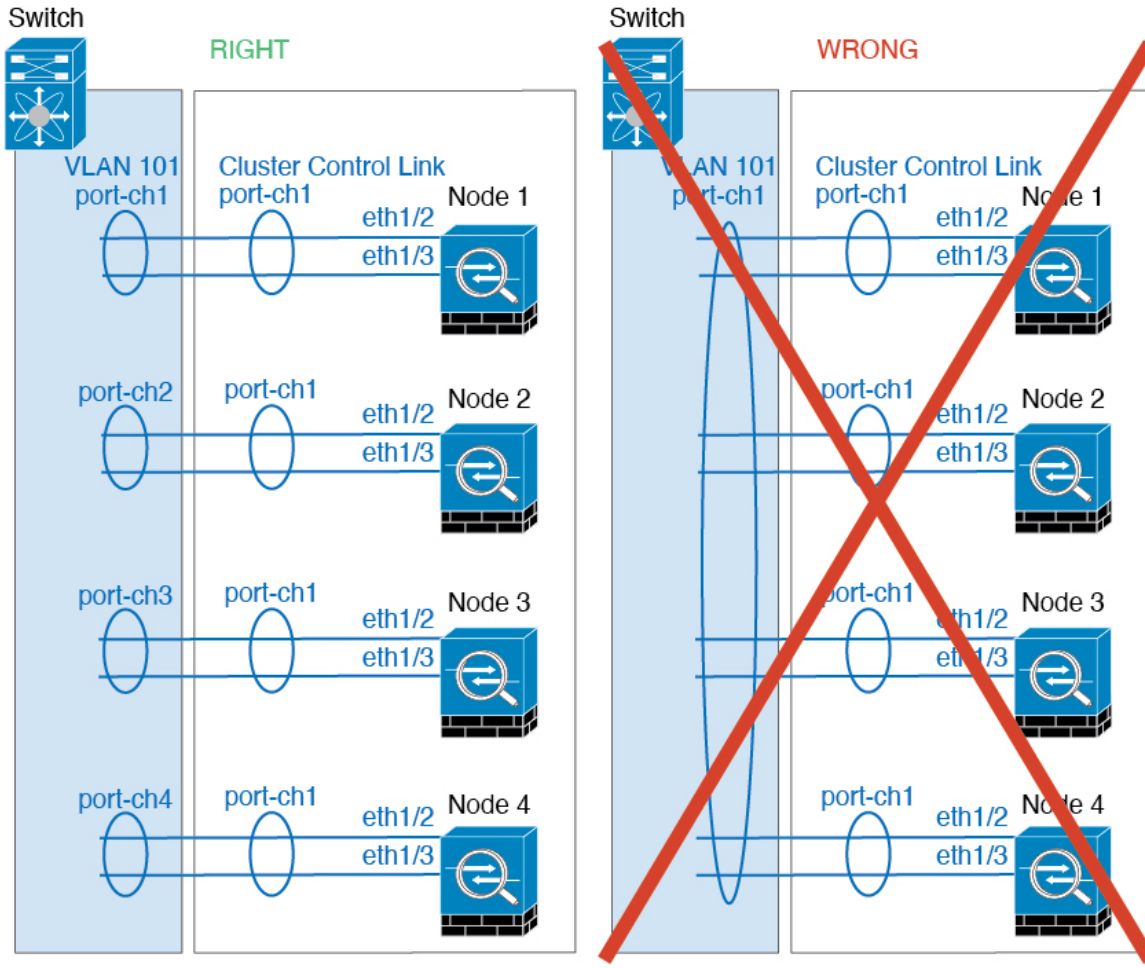
Note that some switches, such as the Nexus series, do not support LACP rate fast when performing in-service software upgrades (ISSUs), so we do not recommend using ISSUs with clustering.

EtherChannels for Clustering

- For connecting switches, set the EtherChannel mode to Active; On mode is not supported on the Firepower 4100/9300 chassis, even for the cluster control link.
- In Catalyst 3750-X Cisco IOS software versions earlier than 15.1(1)S2, the cluster unit did not support connecting an EtherChannel to a switch stack. With default switch settings, if the cluster unit EtherChannel is connected cross stack, and if the control unit switch is powered down, then the EtherChannel connected to the remaining switch will not come up. To improve compatibility, set the **stack-mac persistent timer** command to a large enough value to account for reload time; for example, 8 minutes or 0 for indefinite. Or, you can upgrade to more a more stable switch software version, such as 15.1(1)S2.
- Spanned vs. Device-Local EtherChannel Configuration—Be sure to configure the switch appropriately for Spanned EtherChannels vs. Device-local EtherChannels.
  - Spanned EtherChannels—For cluster unit *Spanned* EtherChannels, which span across all members of the cluster, the interfaces are combined into a single EtherChannel on the switch. Make sure each interface is in the same channel group on the switch.



- Device-local EtherChannels—For cluster unit *Device-local* EtherChannels including any EtherChannels configured for the cluster control link, be sure to configure discrete EtherChannels on the switch; do not combine multiple cluster unit EtherChannels into one EtherChannel on the switch.



Additional Guidelines

- When adding a unit to an existing cluster, or when reloading a unit, there will be a temporary, limited packet/connection drop; this is expected behavior. In some cases, the dropped packets can hang connections; for example, dropping a FIN/ACK packet for an FTP connection will make the FTP client hang. In this case, you need to reestablish the FTP connection.
- If you use a Windows 2003 server connected to a Spanned EtherChannel interface, when the syslog server port is down, and the server does not throttle ICMP error messages, then large numbers of ICMP messages are sent back to the cluster. These messages can result in some units of the cluster experiencing high CPU, which can affect performance. We recommend that you throttle ICMP error messages.
- We recommend connecting EtherChannels to a VSS, vPC, StackWise, or StackWise Virtual for redundancy.

- Within a chassis, you cannot cluster some security modules and run other security modules in standalone mode; you must include all security modules in the cluster.
- For decrypted TLS/SSL connections, the decryption states are not synchronized, and if the connection owner fails, then decrypted connections will be reset. New connections will need to be established to a new unit. Connections that are not decrypted (they match a do-not-decrypt rule) are not affected and are replicated correctly.

Defaults

- The cluster health check feature is enabled by default with the holdtime of 3 seconds. Interface health monitoring is enabled on all interfaces by default.
- The cluster auto-rejoin feature for a failed cluster control link is set to unlimited attempts every 5 minutes.
- The cluster auto-rejoin feature for a failed data interface is set to 3 attempts every 5 minutes, with the increasing interval set to 2.
- Connection replication delay of 5 seconds is enabled by default for HTTP traffic.

Configure Clustering

You can easily deploy the cluster from the Firepower 4100/9300 chassis supervisor. All initial configuration is automatically generated for each unit. You can then add the units to the FMC and group them into a cluster.

FXOS: Add a FTD Cluster

You can add a cluster to a single Firepower 9300 chassis that is isolated to security modules within the chassis, or you can use multiple chassis.

For clusters on multiple chassis, you must configure each chassis separately. Add the cluster on one chassis; you can then copy the bootstrap configuration from the first chassis to the next chassis for ease of deployment

Create a FTD Cluster

You can easily deploy the cluster from the Firepower 4100/9300 chassis supervisor. All initial configuration is automatically generated for each unit.

For clustering on multiple chassis, you must configure each chassis separately. Deploy the cluster on one chassis; you can then copy the bootstrap configuration from the first chassis to the next chassis for ease of deployment.

In a Firepower 9300 chassis, you must enable clustering for all 3 module slots, even if you do not have a module installed. If you do not configure all 3 modules, the cluster will not come up.

Before you begin

- Download the application image you want to use for the logical device from Cisco.com, and then upload that image to the Firepower 4100/9300 chassis.
- Gather the following information:
  - Management interface ID, IP addresses, and network mask

- Gateway IP address
- FMC IP address and/or NAT ID of your choosing
- DNS server IP address
- FTD hostname and domain name

Procedure

Step 1 Configure interfaces.

Step 2 Choose **Logical Devices**.

Step 3 Click **Add** > **ClusterAdd Device**, and set the following parameters:

a) Provide a **Device Name**.

This name is used internally by the chassis supervisor to configure management settings and to assign interfaces; it is not the device name used in the application configuration.

b) For the **Template**, choose **Cisco Firepower Threat Defense**.
c) Choose the **Image Version**.
d) For the **Instance Type**, only the **Native** type is supported.
e) For the **Usage**, click the **Cluster** radio button.
f) Click the **Create New Cluster** radio button.
g) Click **OK**.

You see the Provisioning - *device name* window.

Step 4 Choose the interfaces you want to assign to this cluster.

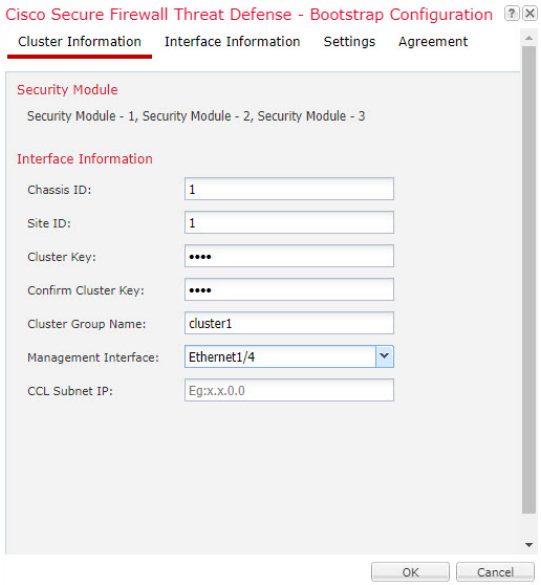
All valid interfaces are assigned by default.

Step 5 Click the device icon in the center of the screen.

A dialog box appears where you can configure initial bootstrap settings. These settings are meant for initial deployment only, or for disaster recovery. For normal operation, you can later change most values in the application CLI configuration.

Step 6 On the **Cluster Information** page, complete the following.

Figure 20:



a) For clustering on multiple chassis, in the **Chassis ID** field, enter a chassis ID. Each chassis in the cluster must use a unique ID.

   This field only appears if you added a member interface to cluster control link Port-Channel 48.

b) For inter-site clustering, in the **Site ID** field, enter the site ID for this chassis between 1 and 8. FlexConfig feature. Additional inter-site cluster customizations to enhance redundancy and stability, such as director localization, and cluster flow mobility, are only configurable using the FMC FlexConfig feature.

c) In the **Cluster Key** field, configure an authentication key for control traffic on the cluster control link.

   The shared secret is an ASCII string from 1 to 63 characters. The shared secret is used to generate the key. This option does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear.

d) Set the **Cluster Group Name**, which is the cluster group name in the logical device configuration.

   The name must be an ASCII string from 1 to 38 characters.

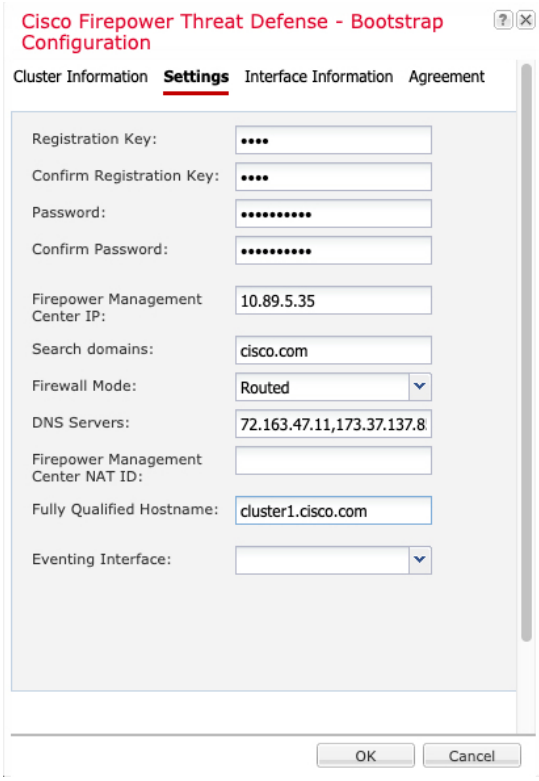
   Important From 2.4.1, spaces in cluster group name will be considered as special characters and may result in error while deploying the logical devices. To avoid this issue, you must rename the cluster group name without a space.

e) Choose the **Management Interface**.

   This interface is used to manage the logical device. This interface is separate from the chassis management port.

   If you assign a Hardware Bypass-capable interface as the Management interface, you see a warning message to make sure your assignment is intentional.

Step 7 On the **Settings** page, complete the following.



a) In the **Registration Key** field, enter the key to be shared between the FMC and the cluster members during registration.

   You can choose any text string for this key between 1 and 37 characters; you will enter the same key on the FMC when you add the Firepower Threat Defense.

b) Enter a **Password** for the Firepower Threat Defense admin user for CLI access.

c) In the **Firepower Management Center IP** field, enter the IP address of the managing FMC.

d) (Optional) In the **Search Domains** field, enter a comma-separated list of search domains for the management network.

e) (Optional) From the **Firewall Mode** drop-down list, choose **Transparent** or **Routed**.

   In routed mode, the Firepower Threat Defense is considered to be a router hop in the network. Each interface that you want to route between is on a different subnet. A transparent firewall, on the other hand, is a Layer 2 firewall that acts like a “bump in the wire,” or a “stealth firewall,” and is not seen as a router hop to connected devices.

   The firewall mode is only set at initial deployment. If you re-apply the bootstrap settings, this setting is not used.

f) (Optional) In the **DNS Servers** field, enter a comma-separated list of DNS servers.

   The Firepower Threat Defense uses DNS if you specify a hostname for the FMC, for example.

g) (Optional) In the **Fully Qualified Hostname** field, enter a fully qualified name for the Firepower Threat Defense device.

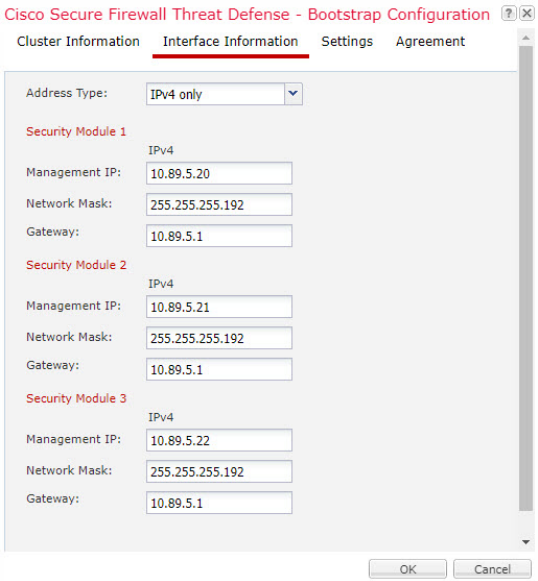
   Valid characters are the letters from a to z, the digits from 0 to 9, the dot (.), and the hyphen (-); maximum number of characters is 253.

h) (Optional) From the **Eventing Interface** drop-down list, choose the interface on which events should be sent. If not specified, the management interface will be used.

To specify a separate interface to use for events, you must configure an interface as a *firepower-eventing* interface. If you assign a Hardware Bypass-capable interface as the Eventing interface, you see a warning message to make sure your assignment is intentional.

Step 8 On the **Interface Information** page, configure a management IP address for each security module in the cluster. Select the type of address from the **Address Type** drop-down list and then complete the following for each security module.

Note You must set the IP address for all 3 module slots in a chassis, even if you do not have a module installed. If you do not configure all 3 modules, the cluster will not come up.



a) In the **Management IP** field, configure an IP address.

Specify a unique IP address on the same network for each module.

b) Enter a **Network Mask** or **Prefix Length**.

c) Enter a **Network Gateway** address.

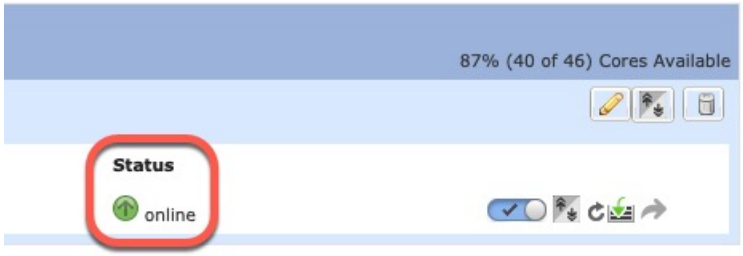
Step 9 On the **Agreement** tab, read and accept the end user license agreement (EULA).

Step 10 Click **OK** to close the configuration dialog box.

Step 11 Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for the status of the new logical device. When the logical device shows its **Status** as **online**, you can add the remaining cluster chassis, or for a cluster isolated to security modules within one Firepower 9300 chassis, start configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.





Step 12 For clustering on multiple chassis, add the next chassis to the cluster:

a) On the first chassis of the Firepower Chassis Manager, click the **Show Configuration icon** at the top right; copy the displayed cluster configuration.

b) Connect to the Firepower Chassis Manager on the next chassis, and add a logical device according to this procedure.

c) Choose **Join an Existing Cluster**.

d) lick **OK**.

e) In the **Copy Cluster Details** box, paste in the cluster configuration from the first chassis, and click **OK**.

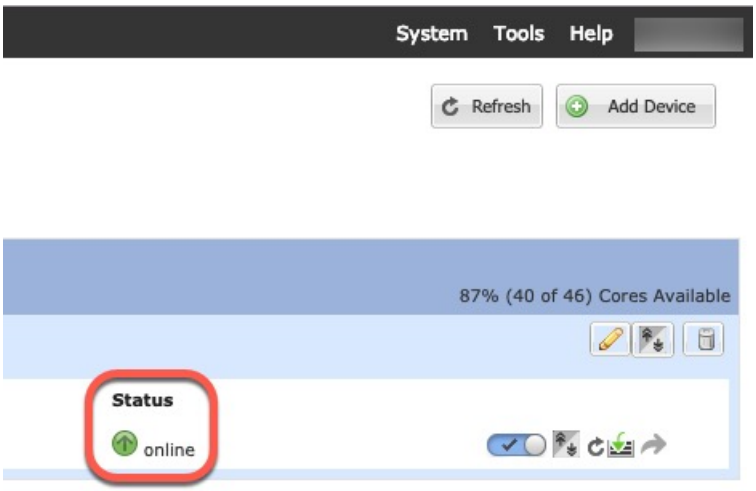
f) Click the device icon in the center of the screen. The cluster information is mostly pre-filled, but you must change the following settings:

- **Chassis ID**—Enter a unique chassis ID.
- **Site ID**—For inter-site clustering, enter the site ID for this chassis between 1 and 8. Additional inter-site cluster customizations to enhance redundancy and stability, such as director localization, and cluster flow mobility, are only configurable using the FMC FlexConfig feature.
- **Cluster Key**—(Not prefilled) Enter the same cluster key.
- **Management IP**—Change the management address for each module to be a unique IP address on the same network as the other cluster members.

Click **OK**.

g) Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for each cluster member for the status of the new logical device. When the logical device for each cluster member shows its **Status** as **online**, you can start configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.



Step 13 Add each unit separately to the FMC using the management IP addresses, and then group them into a cluster at the web interface.

All cluster units must be in a successfully-formed cluster on FXOS prior to adding them to FMC.

Add More Cluster Nodes

Add or replace the Firepower Threat Defense cluster node in an existing cluster.

Note The FXOS steps in this procedure only apply to adding a new *chassis*; if you are adding a new module to a Firepower 9300 where clustering is already enabled, the module will be added automatically.

Before you begin

- In the case of a replacement, you must delete the old cluster node from the FMC. When you replace it with a new node, it is considered to be a new device on the FMC.
- The interface configuration must be the same on the new chassis. You can export and import FXOS chassis configuration to make this process easier.

Procedure

Step 1 On an existing cluster chassis Firepower Chassis Manager, click **Logical Devices**.

Step 2 Click the **Show Configuration icon** at the top right; copy the displayed cluster configuration.

Step 3 Connect to the Firepower Chassis Manager on the new chassis, and click **Add** > **Cluster**.

Step 4 For the **Device Name**, provide a name for the logical device.

Step 5

Step 6

Step 7

Step 8

Step 9 lick **OK**.

Step 10 In the **Copy Cluster Details** box, paste in the cluster configuration from the first chassis, and click **OK**.

Step 11 Click the device icon in the center of the screen. The cluster information is partly pre-filled, but you must fill in the following settings:

Figure 21: Cluster Information

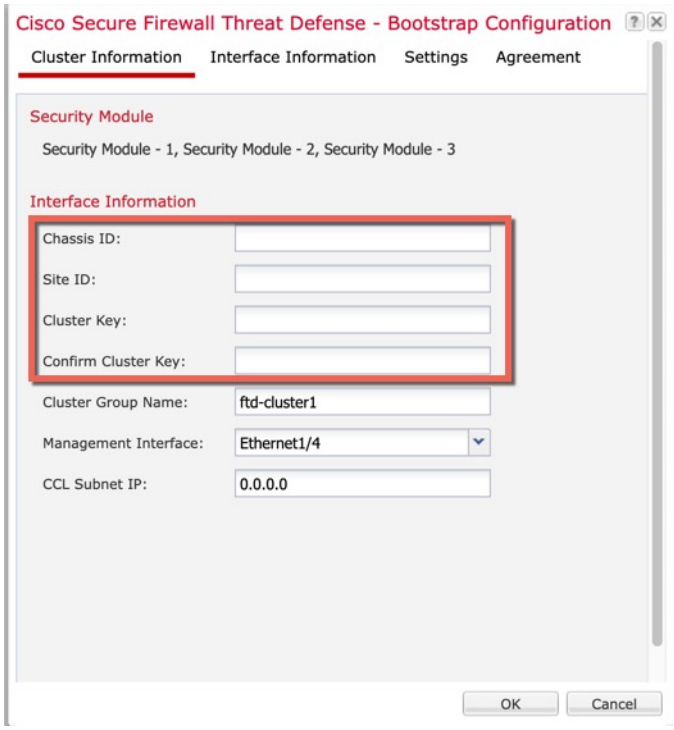


Figure 22: Interface Information

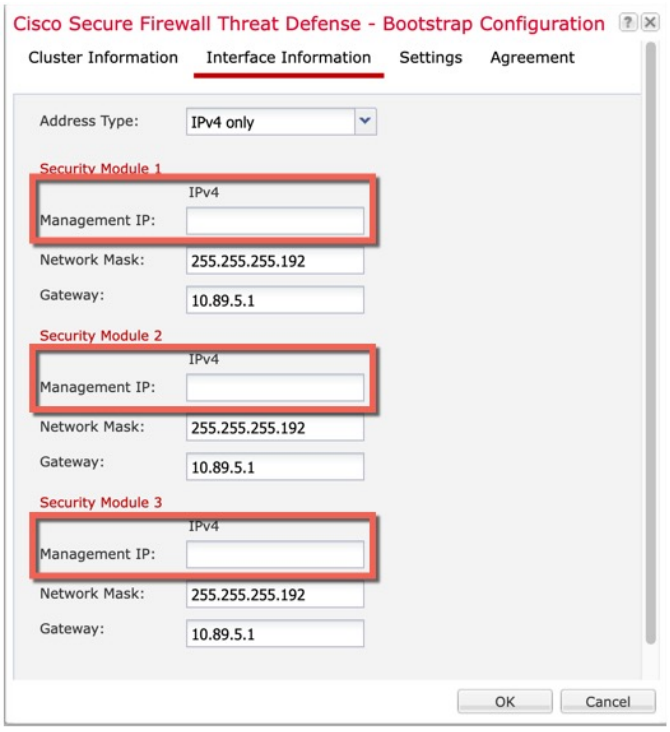
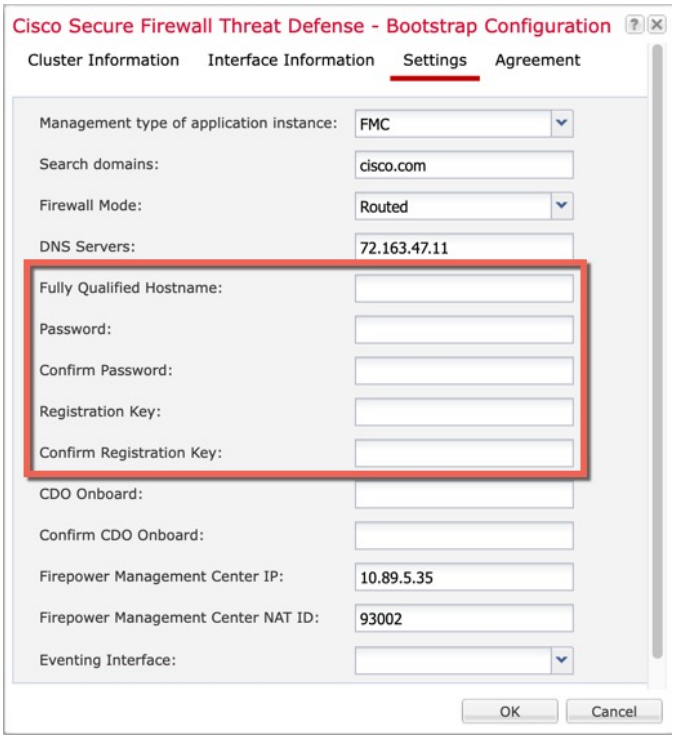


Figure 23: Settings



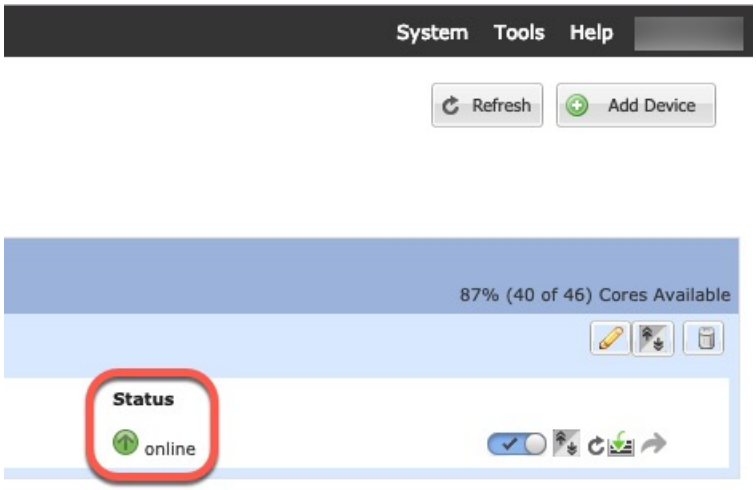
- **Chassis ID**—Enter a *unique* chassis ID.

- **Site ID**—For inter-site clustering, enter the site ID for this chassis between 1 and 8. This feature is only configurable using the FMC FlexConfig feature.
- **Cluster Key**—Enter the *same* cluster key.
- **Management IP**—Change the management address for each module to be a *unique* IP address on the same network as the other cluster members.
- **Fully Qualified Hostname**—Enter the *same* hostname.
- **Password**—Enter the *same* password.
- **Registration Key**—Enter the *same* registration key.

Click **OK**.

Step 12 Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for each cluster member for the status of the new logical device. When the logical device for each cluster member shows its **Status** as **online**, you can start configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.



FMC: Add a Cluster

Add the logical devices to the FMC, and then group them into a cluster.

Before you begin

- Refer to the Firepower Chassis Manager **Logical Devices** screen to see which unit is the control unit.
- All cluster units must be in a successfully formed cluster on FXOS prior to adding them to the FMC.

Procedure

Step 1 In the FMC, choose **Devices** > **Device Management**, and choose **Add** > **Add Device** to add each unit as a separate managed device using the management IP addresses you assigned when you deployed the cluster.

Note If you use FMC High Availability, make sure the standby FMC also successfully registers each unit before you continue and form the cluster on the active FMC: Log into the standby FMC to check the registration status of each unit.

Step 2 Choose **Add** > **Add Cluster** to group the units into a cluster.

a) Choose the control device from the drop-down list.

All other eligible members are added to the data devices box.

b) Specify a **Name** for the cluster.

c) Click **OK**.

The cluster object is added to the **Devices** screen, with the member units underneath. The current control unit is indicated by its role after the unit name.

Note If you add more units to the cluster later on the FXOS chassis, then you must add each unit to the FMC, and then add them as data nodes of the cluster as soon as possible.

Step 3 To configure device-specific settings, click **Edit** ( ) for the cluster; you can only configure the cluster as a whole, and not member units in the cluster.

Step 4 On the **Devices** > **Device Management** > **Cluster**, you can see **General**, **License**, and **Health** settings. This is useful for setting license entitlements.

Step 5 On the **Devices** > **Device Management** > **Devices**, you can choose each member in the cluster from the top right drop-down menu.

If you change the management IP address in the device configuration, you must match the new address in the FMC so that it can reach the device on the network; edit the **Host** address in the **Management** area.

FMC: Configure Cluster, Data, and Diagnostic Interfaces

This procedure configures basic parameters for each data interface that you assigned to the cluster when you deployed it in FXOS. For inter-chassis clustering, data interfaces are always Spanned EtherChannel interfaces. For the cluster control link interface for inter-chassis clustering, you must increase the MTU from the default. You can also configure the Diagnostic interface, which is the only interface that can run as an individual interface.

Note When using Spanned EtherChannels for inter-chassis clustering, the port-channel interface will not come up until clustering is fully enabled. This requirement prevents traffic from being forwarded to a unit that is not an active unit in the cluster.

Procedure

Step 1 Choose **Devices** > **Device Management**, and click **Edit** ( ) next to the cluster.

Step 2 Click **Interfaces**.

Step 3 Configure the cluster control link.

For inter-chassis clustering, set the cluster control link MTU to be at least 100 bytes higher than the highest MTU of the data interfaces. Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead. We suggest setting the MTU to the maximum of 9184; the minimum value is 1400 bytes. For example, because the maximum MTU is 9184, then the highest data interface MTU can be 9084, while the cluster control link can be set to 9184.

The cluster control link interface is Port-Channel48 by default.

a) Click **Edit** ( ) for the cluster control link interface.
b) On the **General** page, in the **MTU** field, enter a value between 1400 and 9184. We suggest using the maximum, 9184.
c) Click **OK**.

Step 4 Configure data interfaces.

a) (Optional) Configure VLAN subinterfaces on the data interface. The rest of this procedure applies to the subinterfaces.

b) Click **Edit** ( ) for the data interface.
c) Configure the name, IP address, and other parameters.

Note If the cluster control link interface MTU is not at least 100 bytes higher than the data interface MTU, you will see an error that you must reduce the MTU of the data interface. See step Step 3, on page 729 to increase the cluster control link MTU, after which you can continue configuring the data interfaces.

d) For inter-chassis clusters, set a manual global MAC address for the EtherChannel. Click **Advanced**, and in the **Active Mac Address** field, enter a MAC address in H.H.H format, where H is a 16-bit hexadecimal digit.

For example, the MAC address 00-0C-F1-42-4C-DE would be entered as 000C.F142.4CDE. The MAC address must not have the multicast bit set, that is, the second hexadecimal digit from the left cannot be an odd number.

Do not set the **Standby Mac Address**; it is ignored.

You must configure a MAC address for a Spanned EtherChannel to avoid potential network connectivity problems. With a manually-configured MAC address, the MAC address stays with the current control unit. If you do not configure a MAC address, then if the control unit changes, the new control unit uses a new MAC address for the interface, which can cause a temporary network outage.

e) Click **OK**. Repeat the above steps for other data interfaces.

Step 5 (Optional) Configure the Diagnostic interface.

The Diagnostic interface is the only interface that can run in Individual interface mode. You can use this interface for syslog messages or SNMP, for example.

a) Choose **Objects** > **Object Management** > **Address Pools** to add an IPv4 and/or IPv6 address pool.

Include at least as many addresses as there are units in the cluster. The Virtual IP address is not a part of this pool, but needs to be on the same network. You cannot determine the exact Local address assigned to each unit in advance.

b) On **Devices** > **Device Management** > **Interfaces**, click **Edit** ( ) for the Diagnostic interface.
c) On the **IPv4**, enter the **IP Address** and mask. This IP address is a fixed address for the cluster, and always belongs to the current control unit.
d) From the **IPv4 Address Pool** drop-down list, choose the address pool you created.
e) On **IPv6 > Basic**, from the **IPv6 Address Pool** drop-down list, choose the address pool you created.
f) Configure other interface settings as normal.

Step 6 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

FXOS: Remove a Cluster Node

The following sections describe how to remove nodes temporarily or permanently from the cluster.

Temporary Removal

A cluster node will be automatically removed from the cluster due to a hardware or network failure, for example. This removal is temporary until the conditions are rectified, and it can rejoin the cluster. You can also manually disable clustering.

To check whether a device is currently in the cluster, check the cluster status on the Firepower Chassis Manager **Logical Devices** page:



For FTD using the FMC, you should leave the device in the FMC device list so that it can resume full functionality after you reenable clustering.

- Disable clustering in the application—You can disable clustering using the application CLI. Enter the **cluster remove unit** *name* command to remove any node other than the one you are logged into. The bootstrap configuration remains intact, as well as the last configuration synced from the control node,

so you can later re-add the node without losing your configuration. If you enter this command on a data node to remove the control node, a new control node is elected.

When a device becomes inactive, all data interfaces are shut down; only the Management interface can send and receive traffic. To resume traffic flow, re-enable clustering. The Management interface remains up using the IP address the node received from the bootstrap configuration. However if you reload, and the node is still inactive in the cluster , the Management interface is disabled.

To reenable clustering, on the FTD enter **cluster enable**.

- Disable the application instance—In the Firepower Chassis Manager on the **Logical Devices** page, click the **Slider enabled** (). You can later reenable it using the **Slider disabled** ().
- Shut down the security module/engine—In the Firepower Chassis Manager on the **Security Module/Engine** page, click the **Power Off icon**.
- Shut down the chassis—In the Firepower Chassis Manager on the **Overview** page, click the **Shut Down icon**.

Permanent Removal

You can permanently remove a cluster node using the following methods.

For FTD using the FMC, be sure to remove the node from the FMC device list after you disable clustering on the chassis.

- Delete the logical device—In the Firepower Chassis Manager on the **Logical Devices** page, click the **Delete** (). You can then deploy a standalone logical device, a new cluster, or even add a new logical device to the same cluster.
- Remove the chassis or security module from service—If you remove a device from service, you can add replacement hardware as a new node of the cluster.

FMC: Manage Cluster Members

After you deploy the cluster, you can change the configuration and manage cluster members.

Add or Replace a Cluster Member

You can add a new cluster member to an existing cluster, for example, when you add an additional module to the Firepower 9300 device, or an additional chassis, or replace a unit. This procedure also applies to a unit that was reinitialized; in this case, although the hardware remains the same, it appears to be a new member.

Before you begin

- Add the units to the cluster on the FXOS chassis, and make sure they are in the FXOS cluster before you add them to the FMC. Make sure the interface configuration is the same as other chassis.

Procedure

Step 1 In the case of a replacement or a reinitialized unit, you must delete the old cluster member from the Firepower Management Center.

When you replace it with a new or reinitialized unit, it is considered to be a new device on the Firepower Management Center

a) In the FMC, choose **Devices** > **Device Management**, and click **Delete** ( ) next to the data unit.
b) Confirm that you want to delete the unit.

The unit is removed from the cluster and from the FMC devices list.

Step 2 Choose **Devices** > **Device Management**, and choose **Add** > **Add Device** to add the new logical device.

Step 3 Choose **Add** > **Add Cluster**.

Step 4 Choose the current control device from the drop-down list.

When you choose a control device that is already in a cluster, then the existing cluster name is auto-filled, and all eligible data devices are added to the data devices box, including the new unit you just added to the FMC.

Step 5 Click **Add**, and then **Deploy**.

The cluster is updated to include the new member(s).

Deactivate a Member

You may want to deactivate a member in preparation for deleting the unit, or temporarily for maintenance. This procedure is meant to temporarily deactivate a member; the unit will still appear in the FMC device list.

To deactivate a member other than the unit you are logged into, perform the following steps at the Firepower Threat Defense CLI.

Note When a unit becomes inactive, all data interfaces are shut down; only the Management interface can send and receive traffic. To resume traffic flow, reenable clustering. The Management interface remains up using the IP address the unit received from the bootstrap configuration. However if you reload, and the unit is still inactive in the cluster, the management interface is disabled. You must use the console for any further configuration.

Procedure

Step 1 Access the Firepower Threat Defense CLI.

Step 2 Remove the unit from the cluster:

cluster remove unit *unit_name*

The bootstrap configuration remains intact, as well as the last configuration synched from the control unit, so that you can later re-add the unit without losing your configuration. If you enter this command on a data unit to remove the control unit, a new control unit is elected.

To view member names, enter **cluster remove unit ?**, or enter the **show cluster info** command.

Example:

```
> cluster remove unit ?

Current active units in the cluster:
ftd1
ftd2
ftd3

> cluster remove unit ftd2
WARNING: Clustering will be disabled on unit ftd2. To bring it back
to the cluster please logon to that unit and re-enable clustering
```

Step 3 To reenable clustering, see Rejoin the Cluster, on page 733.

Rejoin the Cluster

If a unit was removed from the cluster, for example for a failed interface, you must manually rejoin the cluster by accessing the unit CLI. Make sure the failure is resolved before you try to rejoin the cluster. See Rejoining the Cluster, on page 745 for more information about why a unit can be removed from a cluster.

Procedure

Step 1 Access the CLI of the unit that needs to rejoin the cluster, either from the console port or using SSH to the Management interface. Log in with the username **admin** and the password you set during initial setup.

Step 2 Enable clustering:

cluster enable

Delete a Data Unit

If you need to permanently remove a cluster member (for example, if you remove a module on the Firepower 9300, or remove a chassis), then you should delete it from the FMC.

Do not delete the member if it is still a healthy part of the cluster, or if you only want to disable the member temporarily. To delete it permanently from the cluster in FXOS, see FXOS: Remove a Cluster Node, on page 730. If you remove it from the FMC, and it is still part of the cluster, it will continue to pass traffic, and could even become the control unit—a control unit that the FMC can no longer manage.

Before you begin

To manually deactivate the unit, see Deactivate a Member, on page 732. Before you delete a unit, the unit must be inactive, either manually or because of a health failure.

Procedure

Step 1 In the FMC for the data unit you want to delete, choose **Devices** > **Device Management**, and click **Delete** () .

Step 2 Confirm that you want to delete the unit.

The unit is removed from the cluster and from the FMC devices list.

FMC: Monitoring the Cluster

You can monitor the cluster in Firepower Management Center and at the Firepower Threat Defense CLI.

- **Devices** > **Device Management** > *cluster_name*.

  When you expand the cluster on the devices listing page, you can see all member units, including the control unit shown with its role next to the IP address.

- **show cluster** {**access-list** [*acl_name*] | **conn** [**count**] | **cpu** [**usage**] | **history** | **interface-mode** | **memory** | **resource usage** | **service-policy** | **traffic** | **xlate count**}

  To view aggregated data for the entire cluster or other information, use the **show cluster** command.

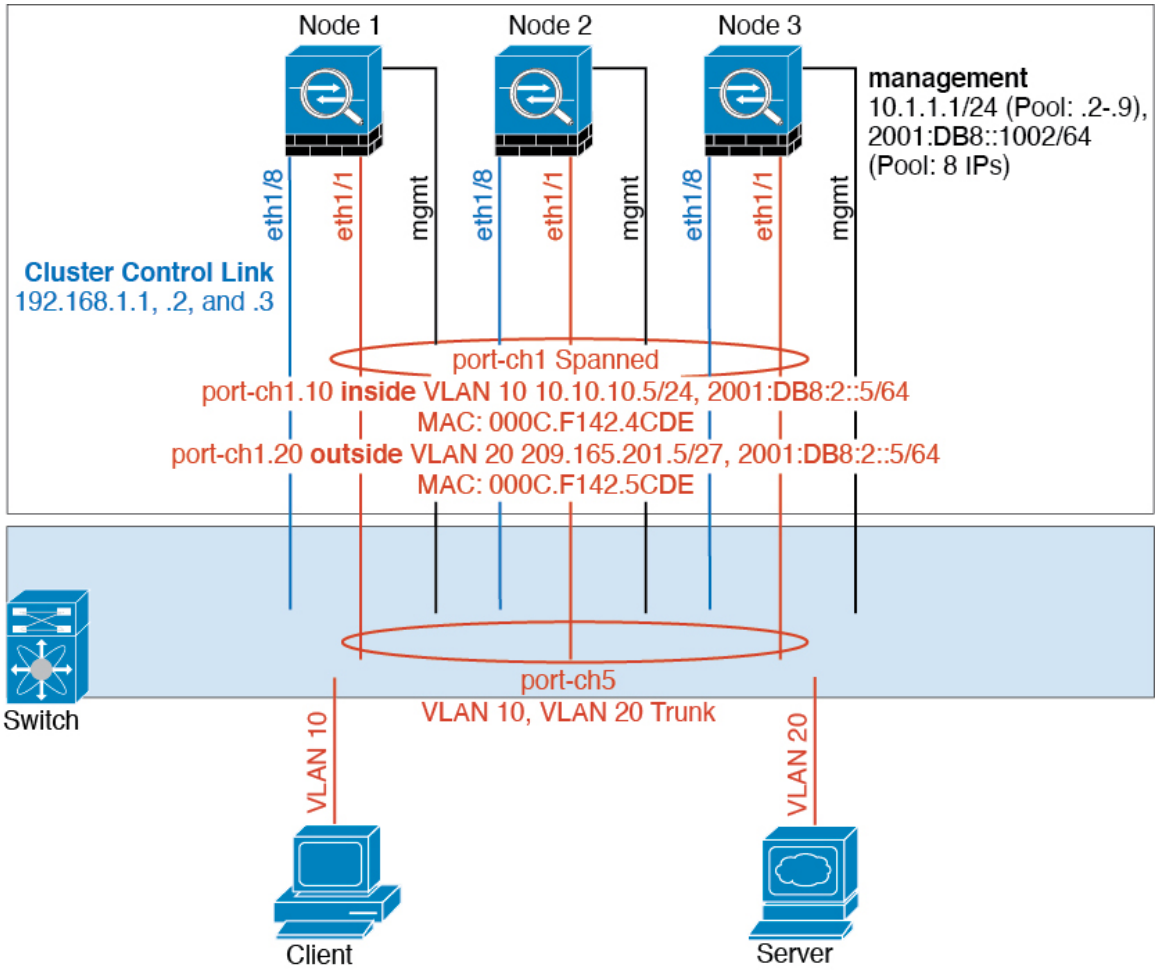
- **show cluster info** [**auto-join** | **clients** | **conn-distribution** | **flow-mobility counters** | **goid** [*options*] | **health** | **incompatible-config** | **loadbalance** | **old-members** | **packet-distribution** | **trace** [*options*] | **transport** { **asp** | **cp**}]

  To view cluster information, use the **show cluster info** command.

Examples for Clustering

These examples include typical deployments.

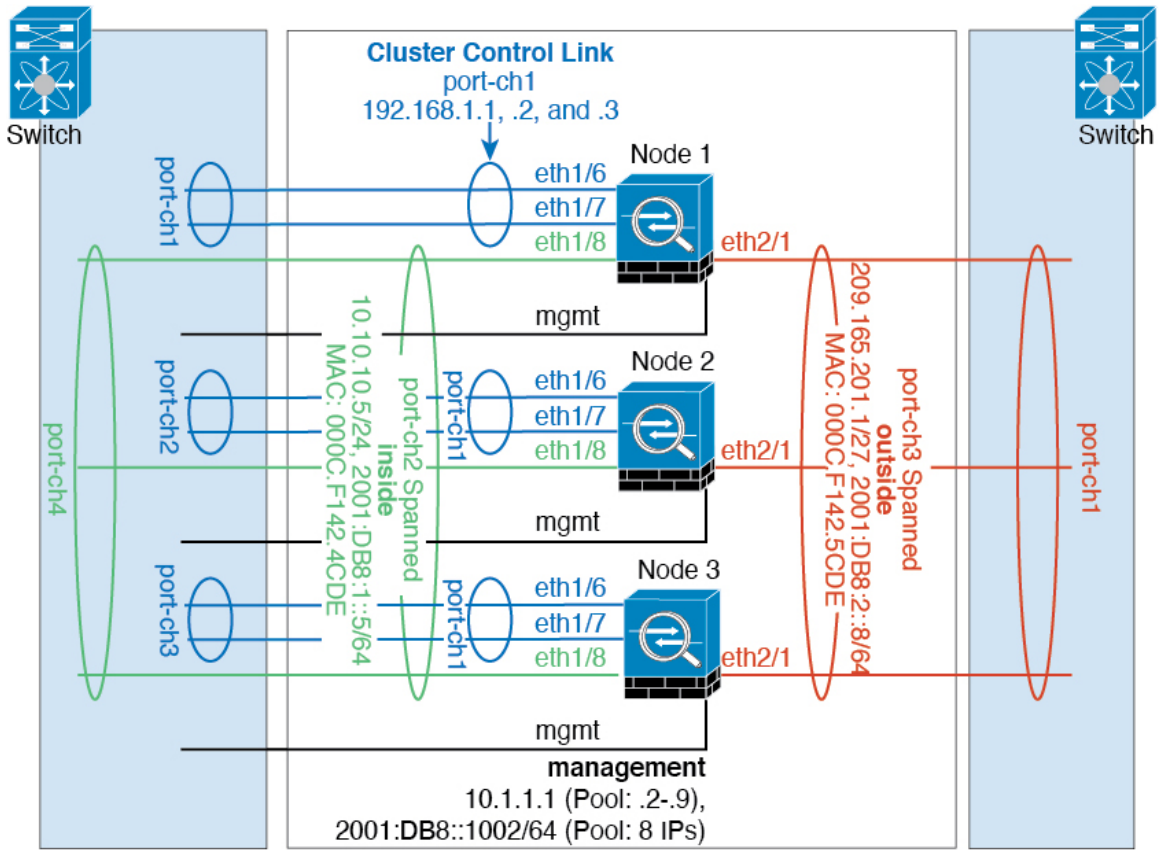
Firewall on a Stick



Data traffic from different security domains are associated with different VLANs, for example, VLAN 10 for the inside network and VLAN 20 for the outside network. Each has a single physical port connected to the external switch or router. Trunking is enabled so that all packets on the physical link are 802.1q encapsulated. The is the firewall between VLAN 10 and VLAN 20.

When using Spanned EtherChannels, all data links are grouped into one EtherChannel on the switch side. If the becomes unavailable, the switch will rebalance traffic between the remaining units.

Traffic Segregation

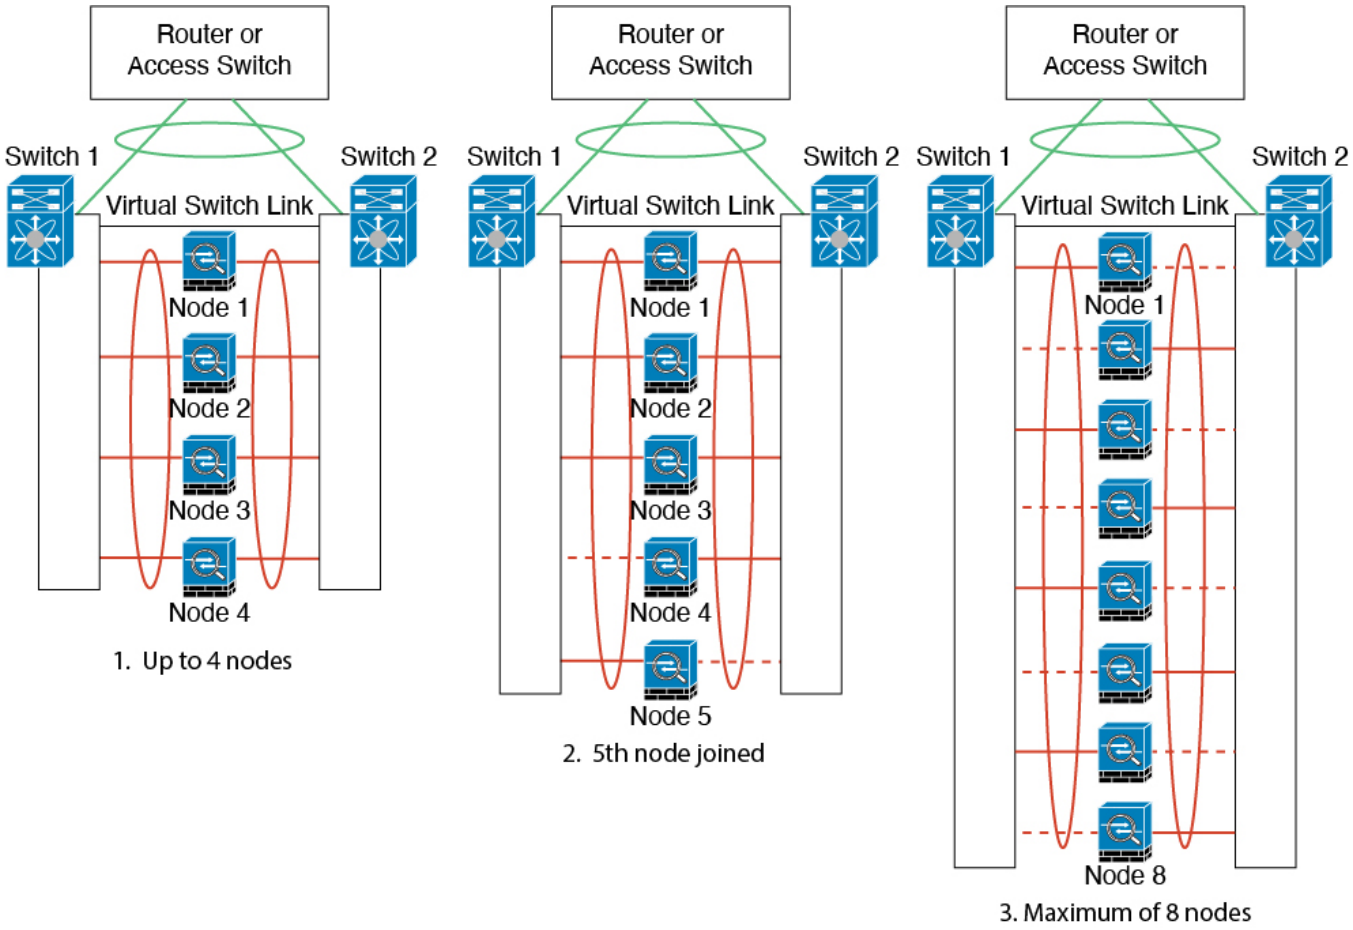


You may prefer physical separation of traffic between the inside and outside network.

As shown in the diagram above, there is one Spanned EtherChannel on the left side that connects to the inside switch, and the other on the right side to outside switch. You can also create VLAN subinterfaces on each EtherChannel if desired.

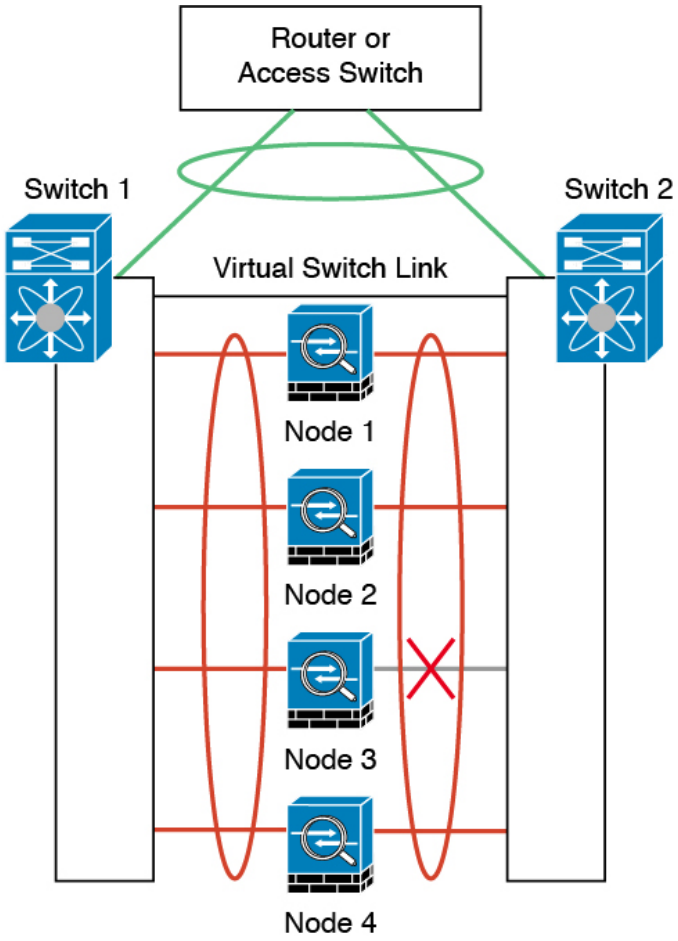
Spanned EtherChannel with Backup Links (Traditional 8 Active/8 Standby)

The maximum number of active ports in a traditional EtherChannel is limited to 8 from the switch side. If you have an 8-unit cluster, and you allocate 2 ports per unit to the EtherChannel, for a total of 16 ports total, then 8 of them have to be in standby mode. The FTD uses LACP to negotiate which links should be active or standby. If you enable multi-switch EtherChannel using VSS, vPC, StackWise, or StackWise Virtual, you can achieve inter-switch redundancy. On the FTD, all physical ports are ordered first by the slot number then by the port number. In the following figure, the lower ordered port is the “control” port (for example, Ethernet 1/1), and the other one is the “data” port (for example, Ethernet 1/2). You must guarantee symmetry in the hardware connection: all control links must terminate on one switch, and all data links must terminate on another switch if a redundant switch system is used. The following diagram shows what happens when the total number of links grows as more units join the cluster:

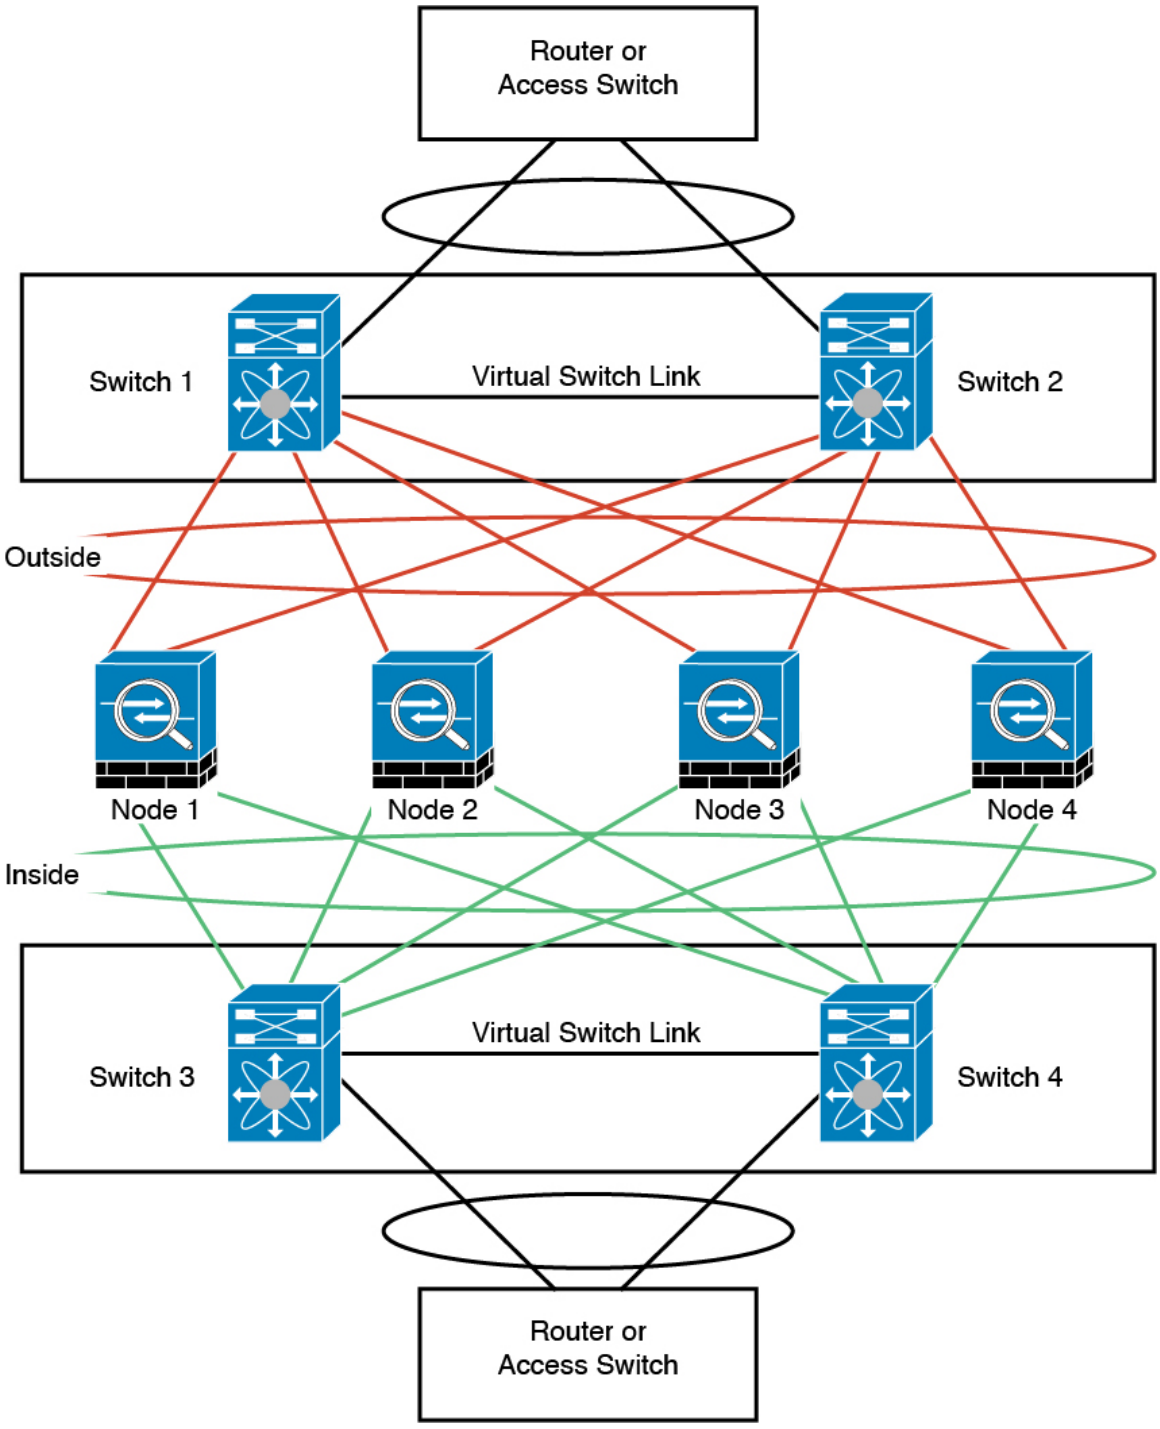


The principle is to first maximize the number of active ports in the channel, and secondly keep the number of active control ports and the number of active data ports in balance. Note that when a 5th unit joins the cluster, traffic is not balanced evenly between all units.

Link or device failure is handled with the same principle. You may end up with a less-than-perfect load balancing situation. The following figure shows a 4-unit cluster with a single link failure on one of the units.



There could be multiple EtherChannels configured in the network. The following diagram shows an EtherChannel on the inside and one on the outside. The FTD is removed from the cluster if both control and data links in one EtherChannel fail. This prevents the FTD from receiving traffic from the outside network when it has already lost connectivity to the inside network.



Reference for Clustering

This section includes more information about how clustering operates.

Firepower Threat Defense Features and Clustering

Some Firepower Threat Defense features are not supported with clustering, and some are only supported on the control unit. Other features might have caveats for proper usage.

Unsupported Features with Clustering

These features cannot be configured with clustering enabled, and the commands will be rejected.

Note To view FlexConfig features that are also not supported with clustering, for example WCCP inspection, see the ASA general operations configuration guide. FlexConfig lets you configure many ASA features that are not present in the FMC GUI.

- Site-to-site VPN
- DHCP client, server, and proxy. DHCP relay is supported.
- High Availability
- Integrated Routing and Bridging
- FMC UCAPL/CC mode

Centralized Features for Clustering

The following features are only supported on the control unit, and are not scaled for the cluster.

Note Traffic for centralized features is forwarded from member units to the control unit over the cluster control link.

If you use the rebalancing feature, traffic for centralized features may be rebalanced to non-control units before the traffic is classified as a centralized feature; if this occurs, the traffic is then sent back to the control unit.

For centralized features, if the control unit fails, all connections are dropped, and you have to re-establish the connections on the new control unit.

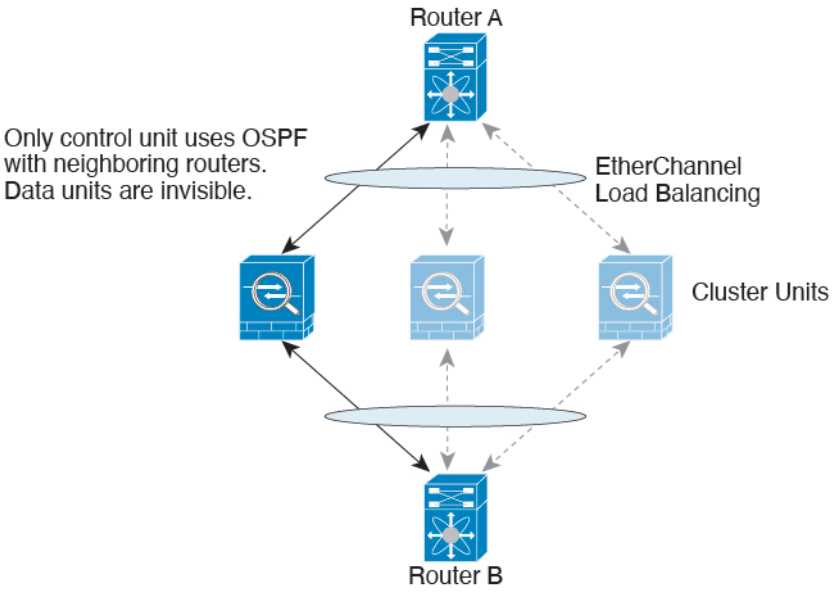
- The following application inspections:
  - DCERPC
  - NetBIOS
  - RSH
  - SUNRPC
  - TFTP
  - XDMCP
- Dynamic routing

- Static route monitoring

Dynamic Routing and Clustering

The routing process only runs on the control unit, and routes are learned through the control unit and replicated to secondaries. If a routing packet arrives at a data unit, it is redirected to the control unit.

Figure 24: Dynamic Routing



After the data units learn the routes from the control unit, each unit makes forwarding decisions independently.

The OSPF LSA database is not synchronized from the control unit to data units. If there is a control unit switchover, the neighboring router will detect a restart; the switchover is not transparent. The OSPF process picks an IP address as its router ID. Although not required, you can assign a static router ID to ensure a consistent router ID is used across the cluster. See the OSPF Non-Stop Forwarding feature to address the interruption.

Connection Settings

Connection limits are enforced cluster-wide. Each node has an estimate of the cluster-wide counter values based on broadcast messages. Due to efficiency considerations, the configured connection limit across the cluster might not be enforced exactly at the limit number. Each node may overestimate or underestimate the cluster-wide counter value at any given time. However, the information will get updated over time in a load-balanced cluster.

FTP and Clustering

- If FTP data channel and control channel flows are owned by different cluster members, then the data channel owner will periodically send idle timeout updates to the control channel owner and update the idle timeout value. However, if the control flow owner is reloaded, and the control flow is re-hosted, the parent/child flow relationship will not longer be maintained; the control flow idle timeout will not be updated.

NAT and Clustering

NAT can affect the overall throughput of the cluster. Inbound and outbound NAT packets can be sent to different FTDs in the cluster, because the load balancing algorithm relies on IP addresses and ports, and NAT causes inbound and outbound packets to have different IP addresses and/or ports. When a packet arrives at the FTD that is not the NAT owner, it is forwarded over the cluster control link to the owner, causing large amounts of traffic on the cluster control link. Note that the receiving node does not create a forwarding flow to the owner, because the NAT owner may not end up creating a connection for the packet depending on the results of security and policy checks.

If you still want to use NAT in clustering, then consider the following guidelines:

- NAT pool address distribution for dynamic PAT—The control node evenly pre-distributes addresses across the cluster. If a member receives a connection and they have no addresses assigned, then the connection is forwarded to the control node for PAT. If a cluster member leaves the cluster (due to failure), a backup member will get the PAT IP address, and if the backup exhausts its normal PAT IP address, it can make use of the new address. Make sure to include at least as many NAT addresses as there are nodes in the cluster, plus at least one extra address, to ensure that each node receives an address, and that a failed node can get a new address if its old address is in use by the member that took over the address. Use the **show nat pool cluster** command in the device CLI to see the address allocations.
- Reusing a PAT pool in multiple rules—To use the same PAT pool in multiple rules, you must be careful about the interface selection in the rules. You must either use specific interfaces in all rules, or "any" in all rules. You cannot mix specific interfaces and "any" across the rules, or the system might not be able to match return traffic to the right node in the cluster. Using unique PAT pools per rule is the most reliable option.
- No round-robin—Round-robin for a PAT pool is not supported with clustering.
- Dynamic NAT xlates managed by the control node—The control node maintains and replicates the xlate table to data nodes. When a data node receives a connection that requires dynamic NAT, and the xlate is not in the table, it requests the xlate from the control node. The data node owns the connection.
- Stale xlates—The xlate idle time on the connection owner does not get updated. Thus, the idle time might exceed the idle timeout. An idle timer value higher than the configured timeout with a refcnt of 0 is an indication of a stale xlate.
- No static PAT for the following inspections—
  - FTP
  - RSH
  - SQLNET
  - TFTP
  - XDMCP
  - SIP

SIP Inspection and Clustering

A control flow can be created on any node (due to load balancing); its child data flows must reside on the same node.

SNMP and Clustering

An SNMP agent polls each individual FTD by its interface Local IP address. You cannot poll consolidated data for the cluster.

You should always use the Local address, and not the Main cluster IP address for SNMP polling. If the SNMP agent polls the Main cluster IP address, if a new control node is elected, the poll to the new control node will fail.

When using SNMPv3 with clustering, if you add a new cluster node after the initial cluster formation, then SNMPv3 users are not replicated to the new node. You must remove the users, and re-add them, and then redeploy your configuration to force the users to replicate to the new node.

Syslog and Clustering

- Each node in the cluster generates its own syslog messages. You can configure logging so that each node uses either the same or a different device ID in the syslog message header field. For example, the hostname configuration is replicated and shared by all nodes in the cluster. If you configure logging to use the hostname as the device ID, syslog messages generated by all nodes look as if they come from a single node. If you configure logging to use the local-node name that is assigned in the cluster bootstrap configuration as the device ID, syslog messages look as if they come from different nodes.

TLS/SSL Connections and Clustering

The decryption states of TLS/SSL connections are not synchronized, and if the connection owner fails, then the decrypted connections will be reset. New connections will need to be established to a new unit. Connections that are not decrypted (they match a do-not-decrypt rule) are not affected and are replicated correctly.

Cisco TrustSec and Clustering

Only the control node learns security group tag (SGT) information. The control node then populates the SGT to data nodes, and data nodes can make a match decision for SGT based on the security policy.

Performance Scaling Factor

When you combine multiple units into a cluster, you can expect the total cluster performance to be approximately 80% of the maximum combined throughput.

For example, for TCP throughput, the Firepower 9300 with 3 SM-40 modules can handle approximately 135 Gbps of real world firewall traffic when running alone. For 2 chassis, the maximum combined throughput will be approximately 80% of 270 Gbps (2 chassis x 135 Gbps): 216 Gbps.

Control Unit Election

Members of the cluster communicate over the cluster control link to elect a control unit as follows:

1. When you deploy the cluster, each unit broadcasts an election request every 3 seconds.
2. Any other units with a higher priority respond to the election request; the priority is set when you deploy the cluster and is not configurable.
3. If after 45 seconds, a unit does not receive a response from another unit with a higher priority, then it becomes the control unit.

Note If multiple units tie for the highest priority, the cluster unit name and then the serial number is used to determine the control unit.

4. If a unit later joins the cluster with a higher priority, it does not automatically become the control unit; the existing control unit always remains as the control unit unless it stops responding, at which point a new control unit is elected.
5. In a "split brain" scenario when there are temporarily multiple control units, then the unit with highest priority retains the role while the other units return to data unit roles.

Note You can manually force a unit to become the control unit. For centralized features, if you force a control unit change, then all connections are dropped, and you have to re-establish the connections on the new control unit.

High Availability Within the Cluster

Clustering provides high availability by monitoring chassis, unit, and interface health and by replicating connection states between units.

Chassis-Application Monitoring

Chassis-application health monitoring is always enabled. The Firepower 4100/9300 chassis supervisor checks the FTD application periodically (every second). If the FTD device is up and cannot communicate with the Firepower 4100/9300 chassis supervisor for 3 seconds, the FTD device generates a syslog message and leaves the cluster.

If the Firepower 4100/9300 chassis supervisor cannot communicate with the application after 45 seconds, it reloads the FTD device. If the FTD device cannot communicate with the supervisor, it removes itself from the cluster.

Unit Health Monitoring

Each unit periodically sends a broadcast keepalivekeepalive packet over the cluster control link. If the control node does not receive any keepalivekeepalive packets or other packets from a data node within the timeout period, then the control node removes the data node from the cluster. If the data nodes do not receive packets from the control node, then a new control node is elected from the remaining node.

If nodes cannot reach each other over the cluster control link because of a network failure and not because a node has actually failed, then the cluster may go into a "split brain" scenario where isolated data nodes will elect their own control nodes. For example, if a router fails between two cluster locations, then the original control node at location 1 will remove the location 2 data nodes from the cluster. Meanwhile, the nodes at location 2 will elect their own control node and form their own cluster. Note that asymmetric traffic may fail in this scenario. After the cluster control link is restored, then the control node that has the higher priority will keep the control node's role. See Control Unit Election, on page 743 for more information.

Interface Monitoring

Each node monitors the link status of all hardware interfaces in use, and reports status changes to the control node. For clustering on multiple chassis, Spanned EtherChannels use the cluster Link Aggregation Control Protocol (cLACP). Each chassis monitors the link status and the cLACP protocol messages to determine if the port is still active in the EtherChannel, and informs the FTD application if the interface is down. All physical interfaces are monitored by default (including the main EtherChannel for EtherChannel interfaces). Only named interfaces that are in an Up state can be monitored. For example, all member ports of an EtherChannel must fail before a *named* EtherChannel is removed from the cluster.

If a monitored interface fails on a particular node, but it is active on other nodes, then the node is removed from the cluster. The amount of time before the FTD device removes a node from the cluster depends on whether the node is an established member or is joining the cluster. The FTD device does not monitor interfaces for the first 90 seconds that a node joins the cluster. Interface status changes during this time will not cause the FTD device to be removed from the cluster. For an established member, the node is removed after 500 ms.

For clustering on multiple chassis, if you add or delete an EtherChannel from the cluster, interface health-monitoring is suspended for 95 seconds to ensure that you have time to make the changes on each chassis.

Decorator Application Monitoring

When you install a decorator application on an interface, such as the Radware DefensePro application, then both the FTD device and the decorator application must be operational to remain in the cluster. The unit does not join the cluster until both applications are operational. Once in the cluster, the unit monitors the decorator application health every 3 seconds. If the decorator application is down, the unit is removed from the cluster.

Status After Failure

When a node in the cluster fails, the connections hosted by that node are seamlessly transferred to other nodes; state information for traffic flows is shared over the control node's cluster control link.

If the control node fails, then another member of the cluster with the highest priority (lowest number) becomes the control node.

The FTD automatically tries to rejoin the cluster, depending on the failure event.



Note When the FTD becomes inactive and fails to automatically rejoin the cluster, all data interfaces are shut down; only the Management/Diagnostic interface can send and receive traffic.

Rejoining the Cluster

After a cluster member is removed from the cluster, how it can rejoin the cluster depends on why it was removed:

- Failed cluster control link when initially joining—After you resolve the problem with the cluster control link, you must manually rejoin the cluster by re-enabling clustering.
- Failed cluster control link after joining the cluster—The FTD automatically tries to rejoin every 5 minutes, indefinitely.

- Failed data interface—The Firepower Threat Defense automatically tries to rejoin at 5 minutes, then at 10 minutes, and finally at 20 minutes. If the join is not successful after 20 minutes, then the Firepower Threat Defense application disables clustering. After you resolve the problem with the data interface, you have to manually enable clustering.
- Failed unit—If the unit was removed from the cluster because of a unit health check failure, then rejoining the cluster depends on the source of the failure. For example, a temporary power failure means the unit will rejoin the cluster when it starts up again as long as the cluster control link is up. The Firepower Threat Defense application attempts to rejoin the cluster every 5 seconds.
- Failed Chassis-Application Communication—When the Firepower Threat Defense application detects that the chassis-application health has recovered, it tries to rejoin the cluster automatically.
- Internal error—Internal failures include: application sync timeout; inconsistent application statuses; and so on. After you resolve the problem, you must manually rejoin the cluster by re-enabling clustering.
- Failed configuration deployment—If you deploy a new configuration from FMC, and the deployment fails on some cluster members but succeeds on others, then the units that failed are removed from the cluster. You must manually rejoin the cluster by re-enabling clustering. If the deployment fails on the control unit, then the deployment is rolled back, and no members are removed.

Data Path Connection State Replication

Every connection has one owner and at least one backup owner in the cluster. The backup owner does not take over the connection in the event of a failure; instead, it stores TCP/UDP state information, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner is usually also the director.

Some traffic requires state information above the TCP or UDP layer. See the following table for clustering support or lack of support for this kind of traffic.

Table 66: Features Replicated Across the Cluster

Traffic	State Support	Notes
Up time	Yes	Keeps track of the system up time.
ARP Table	Yes	—
MAC address table	Yes	—
User Identity	Yes	—
IPv6 Neighbor database	Yes	—
Dynamic routing	Yes	—
SNMP Engine ID	**No**	—

How the Cluster Manages Connections

Connections can be load-balanced to multiple nodes of the cluster. Connection roles determine how connections are handled in both normal operation and in a high availability situation.

Connection Roles

See the following roles defined for each connection:

- Owner—Usually, the node that initially receives the connection. The owner maintains the TCP state and processes packets. A connection has only one owner. If the original owner fails, then when new nodes receive packets from the connection, the director chooses a new owner from those nodes.
- Backup owner—The node that stores TCP/UDP state information received from the owner, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner does not take over the connection in the event of a failure. If the owner becomes unavailable, then the first node to receive packets from the connection (based on load balancing) contacts the backup owner for the relevant state information so it can become the new owner.

  As long as the director (see below) is not the same node as the owner, then the director is also the backup owner. If the owner chooses itself as the director, then a separate backup owner is chosen.

  For clustering on the Firepower 9300, which can include up to 3 cluster nodes in one chassis, if the backup owner is on the same chassis as the owner, then an additional backup owner will be chosen from another chassis to protect flows from a chassis failure.
- Director—The node that handles owner lookup requests from forwarders. When the owner receives a new connection, it chooses a director based on a hash of the source/destination IP address and ports (see below for ICMP hash details), and sends a message to the director to register the new connection. If packets arrive at any node other than the owner, the node queries the director about which node is the owner so it can forward the packets. A connection has only one director. If a director fails, the owner chooses a new director.

  As long as the director is not the same node as the owner, then the director is also the backup owner (see above). If the owner chooses itself as the director, then a separate backup owner is chosen.

  ICMP/ICMPv6 hash details:

  - For Echo packets, the source port is the ICMP identifier, and the destination port is 0.
  - For Reply packets, the source port is 0, and the destination port is the ICMP identifier.
  - For other packets, both source and destination ports are 0.
- Forwarder—A node that forwards packets to the owner. If a forwarder receives a packet for a connection it does not own, it queries the director for the owner, and then establishes a flow to the owner for any other packets it receives for this connection. The director can also be a forwarder. Note that if a forwarder receives the SYN-ACK packet, it can derive the owner directly from a SYN cookie in the packet, so it does not need to query the director. (If you disable TCP sequence randomization, the SYN cookie is not used; a query to the director is required.) For short-lived flows such as DNS and ICMP, instead of querying, the forwarder immediately sends the packet to the director, which then sends them to the owner. A connection can have multiple forwarders; the most efficient throughput is achieved by a good load-balancing method where there are no forwarders and all packets of a connection are received by the owner.



Note We do not recommend disabling TCP sequence randomization when using clustering. There is a small chance that some TCP sessions won't be established, because the SYN/ACK packet might be dropped.

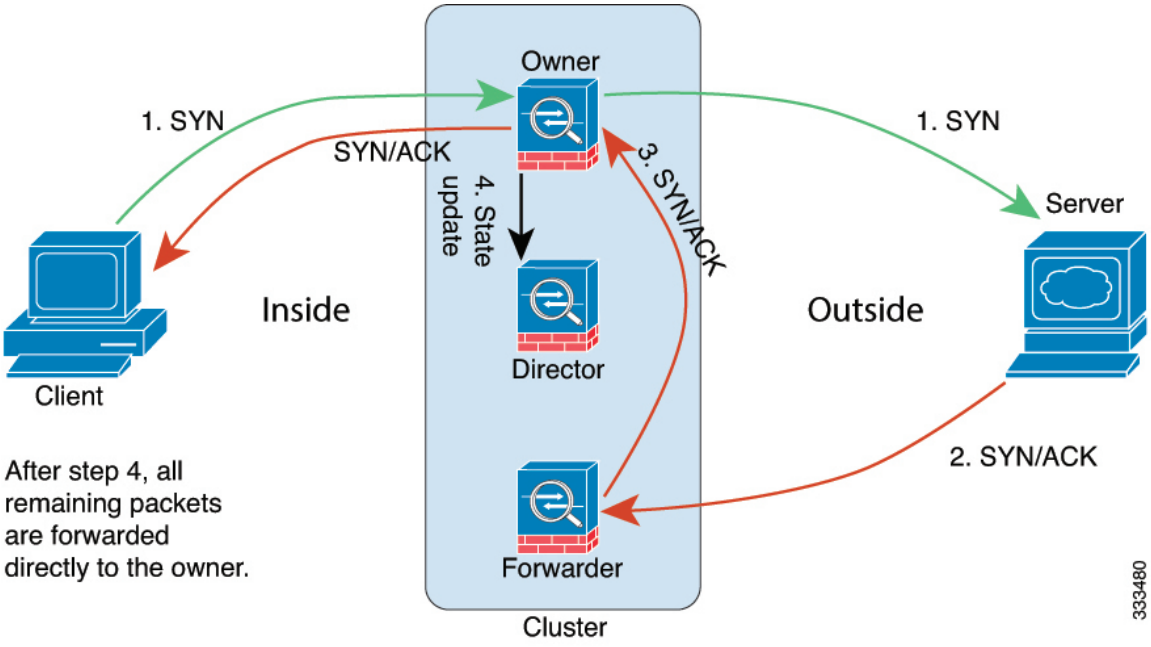
- Fragment Owner—For fragmented packets, cluster nodes that receive a fragment determine a fragment owner using a hash of the fragment source IP address, destination IP address, and the packet ID. All fragments are then forwarded to the fragment owner over the cluster control link. Fragments may be load-balanced to different cluster nodes, because only the first fragment includes the 5-tuple used in the switch load balance hash. Other fragments do not contain the source and destination ports and may be load-balanced to other cluster nodes. The fragment owner temporarily reassembles the packet so it can determine the director based on a hash of the source/destination IP address and ports. If it is a new connection, the fragment owner will register to be the connection owner. If it is an existing connection, the fragment owner forwards all fragments to the provided connection owner over the cluster control link. The connection owner will then reassemble all fragments.

New Connection Ownership

When a new connection is directed to a node of the cluster via load balancing, that node owns both directions of the connection. If any connection packets arrive at a different node, they are forwarded to the owner node over the cluster control link. If a reverse flow arrives at a different node, it is redirected back to the original node.

Sample Data Flow for TCP

The following example shows the establishment of a new connection.



1. The SYN packet originates from the client and is delivered to one FTD (based on the load balancing method), which becomes the owner. The owner creates a flow, encodes owner information into a SYN cookie, and forwards the packet to the server.
2. The SYN-ACK packet originates from the server and is delivered to a different FTD (based on the load balancing method). This FTD is the forwarder.
3. Because the forwarder does not own the connection, it decodes owner information from the SYN cookie, creates a forwarding flow to the owner, and forwards the SYN-ACK to the owner.

4. The owner sends a state update to the director, and forwards the SYN-ACK to the client.
5. The director receives the state update from the owner, creates a flow to the owner, and records the TCP state information as well as the owner. The director acts as the backup owner for the connection.
6. Any subsequent packets delivered to the forwarder will be forwarded to the owner.
7. If packets are delivered to any additional nodes, it will query the director for the owner and establish a flow.
8. Any state change for the flow results in a state update from the owner to the director.

History for Clustering

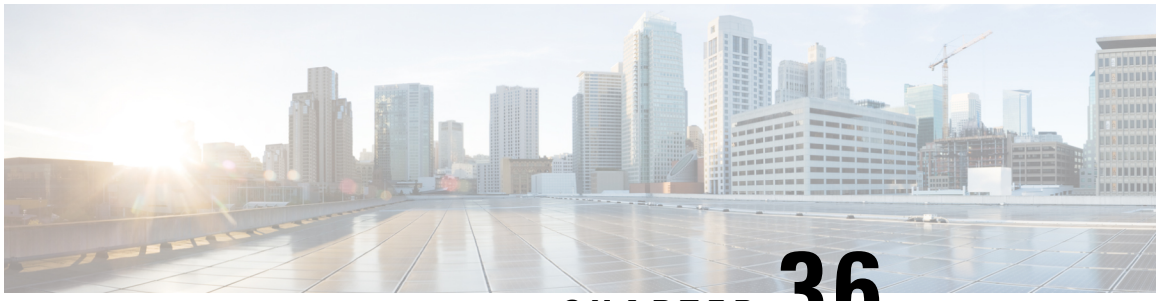
Feature	Version	Details
Inter-chassis clustering for 6 modules; Firepower 4100 support	6.2	With FXOS 2.1.1, you can now enable inter-chassis clustering on the Firepower 9300 and 4100. For the Firepower 9300, you can include up to 6 modules. For example, you can use 1 module in 6 chassis, or 2 modules in 3 chassis, or any combination that provides a maximum of 6 modules. For the Firepower 4100, you can include up to 6 chassis. **Note** Inter-site clustering is also supported. However, customizations to enhance redundancy and stability, such as site-specific MAC and IP addresses, director localization, and cluster flow mobility, are only configurable using the FlexConfig feature. No modified screens. Supported platforms: Firepower Threat Defense on the Firepower 4100/9300
Intra-chassis Clustering for the Firepower 9300	6.0.1	You can cluster up to 3 security modules within the Firepower 9300 chassis. All modules in the chassis must belong to the cluster. New/Modified screens: **Devices** > **Device Management** > **Add** > **Add Cluster** **Devices** > **Device Management** > **Cluster** Supported platforms: Firepower Threat Defense on the Firepower 9300



PART X

Firepower Threat Defense Routing

- Routing Overview for Firepower Threat Defense, on page 753
- Static and Default Routes for Firepower Threat Defense, on page 763
- OSPF for Firepower Threat Defense, on page 769
- BGP for Firepower Threat Defense, on page 797
- RIP for Firepower Threat Defense, on page 813
- Multicast Routing for Firepower Threat Defense, on page 819



CHAPTER 36

Routing Overview for Firepower Threat Defense

This chapter describes underlying concepts of how routing behaves within the Firepower Threat Defense.

- Path Determination, on page 753
- Supported Route Types, on page 753
- Supported Internet Protocols for Routing, on page 755
- Routing Table, on page 755
- Routing Table for Management Traffic, on page 759
- Equal-Cost Multi-Path (ECMP) Routing, on page 759
- About Route Maps, on page 760

Path Determination

Routing protocols use metrics to evaluate what path will be the best for a packet to travel. A metric is a standard of measurement, such as path bandwidth, that is used by routing algorithms to determine the optimal path to a destination. To aid the process of path determination, routing algorithms initialize and maintain routing tables, which include route information. Route information varies depending on the routing algorithm used.

Routing algorithms fill routing tables with a variety of information. Destination or next hop associations tell a router that a particular destination can be reached optimally by sending the packet to a particular router representing the next hop on the way to the final destination. When a router receives an incoming packet, it checks the destination address and attempts to associate this address with a next hop.

Routing tables also can include other information, such as data about the desirability of a path. Routers compare metrics to determine optimal routes, and these metrics differ depending on the design of the routing algorithm used.

Routers communicate with one another and maintain their routing tables through the transmission of a variety of messages. The routing update message is one such message that generally consists of all or a portion of a routing table. By analyzing routing updates from all other routers, a router can build a detailed picture of network topology. A link-state advertisement, another example of a message sent between routers, informs other routers of the state of the sender links. Link information also can be used to build a complete picture of network topology to enable routers to determine optimal routes to network destinations.

Supported Route Types

There are several route types that a router can use. The FTD device uses the following route types:

- Static Versus Dynamic
- Single-Path Versus Multipath
- Flat Versus Hierarchical
- Link-State Versus Distance Vector

Static Versus Dynamic

Static routing algorithms are actually table mappings established by the network administrator. These mappings do not change unless the network administrator alters them. Algorithms that use static routes are simple to design and work well in environments where network traffic is relatively predictable and where network design is relatively simple.

Because static routing systems cannot react to network changes, they generally are considered unsuitable for large, constantly changing networks. Most of the dominant routing algorithms are dynamic routing algorithms, which adjust to changing network circumstances by analyzing incoming routing update messages. If the message indicates that a network change has occurred, the routing software recalculates routes and sends out new routing update messages. These messages permeate the network, stimulating routers to rerun their algorithms and change their routing tables accordingly.

Dynamic routing algorithms can be supplemented with static routes where appropriate. A router of last resort (a default route for a router to which all unrouteable packets are sent), for example, can be designated to act as a repository for all unrouteable packets, ensuring that all messages are at least handled in some way.

Single-Path Versus Multipath

Some sophisticated routing protocols support multiple paths to the same destination. Unlike single-path algorithms, these multipath algorithms permit traffic multiplexing over multiple lines. The advantages of multipath algorithms are substantially better throughput and reliability, which is generally called load sharing.

Flat Versus Hierarchical

Some routing algorithms operate in a flat space, while others use routing hierarchies. In a flat routing system, the routers are peers of all others. In a hierarchical routing system, some routers form what amounts to a routing backbone. Packets from non-backbone routers travel to the backbone routers, where they are sent through the backbone until they reach the general area of the destination. At this point, they travel from the last backbone router through one or more non-backbone routers to the final destination.

Routing systems often designate logical groups of nodes, called domains, autonomous systems, or areas. In hierarchical systems, some routers in a domain can communicate with routers in other domains, while others can communicate only with routers within their domain. In very large networks, additional hierarchical levels may exist, with routers at the highest hierarchical level forming the routing backbone.

The primary advantage of hierarchical routing is that it mimics the organization of most companies and therefore supports their traffic patterns well. Most network communication occurs within small company groups (domains). Because intradomain routers need to know only about other routers within their domain, their routing algorithms can be simplified, and, depending on the routing algorithm being used, routing update traffic can be reduced accordingly.

Link-State Versus Distance Vector

Link-state algorithms (also known as shortest path first algorithms) flood routing information to all nodes in the internetwork. Each router, however, sends only the portion of the routing table that describes the state of its own links. In link-state algorithms, each router builds a picture of the entire network in its routing tables. Distance vector algorithms (also known as Bellman-Ford algorithms) call for each router to send all or some portion of its routing table, but only to its neighbors. In essence, link-state algorithms send small updates everywhere, while distance vector algorithms send larger updates only to neighboring routers. Distance vector algorithms know only about their neighbors. Typically, link-state algorithms are used in conjunction with OSPF routing protocols.

Supported Internet Protocols for Routing

The FTD device supports several Internet protocols for routing. Each protocol is briefly described in this section.

- Enhanced Interior Gateway Routing Protocol (EIGRP)

  EIGRP is a Cisco proprietary protocol that provides compatibility and seamless interoperation with IGRP routers. An automatic-redistribution mechanism allows IGRP routes to be imported into Enhanced IGRP, and vice versa, so it is possible to add Enhanced IGRP gradually into an existing IGRP network.

- Open Shortest Path First (OSPF)

  OSPF is a routing protocol developed for Internet Protocol (IP) networks by the interior gateway protocol (IGP) working group of the Internet Engineering Task Force (IETF). OSPF uses a link-state algorithm to build and calculate the shortest path to all known destinations. Each router in an OSPF area includes an identical link-state database, which is a list of each of the router usable interfaces and reachable neighbors.

- Routing Information Protocol (RIP)

  RIP is a distance-vector protocol that uses hop count as its metric. RIP is widely used for routing traffic in the global Internet and is an interior gateway protocol (IGP), which means that it performs routing within a single autonomous system.

- Border Gateway Protocol (BGP)

  BGP is an interautonomous system routing protocol. BGP is used to exchange routing information for the Internet and is the protocol used between Internet service providers (ISP). Customers connect to ISPs, and ISPs use BGP to exchange customer and ISP routes. When BGP is used between autonomous systems (AS), the protocol is referred to as External BGP (EBGP). If a service provider is using BGP to exchange routes within an AS, then the protocol is referred to as Interior BGP (IBGP).

Routing Table

The FTD uses separate routing tables for data traffic (through-the-device) and for management traffic (from-the-device). This section describes how the routing tables work. For information about the management routing table, see also Routing Table for Management Traffic, on page 759.

How the Routing Table Is Populated

The FTD routing table can be populated by statically defined routes, directly connected routes, and routes discovered by the dynamic routing protocols. Because the FTD device can run multiple routing protocols in addition to having static and connected routes in the routing table, it is possible that the same route is discovered or entered in more than one manner. When two routes to the same destination are put into the routing table, the one that remains in the routing table is determined as follows:

- If the two routes have different network prefix lengths (network masks), then both routes are considered unique and are entered into the routing table. The packet forwarding logic then determines which of the two to use.

  For example, if the RIP and OSPF processes discovered the following routes:

  - RIP: 192.168.32.0/24
  - OSPF: 192.168.32.0/19

  Even though OSPF routes have the better administrative distance, both routes are installed in the routing table because each of these routes has a different prefix length (subnet mask). They are considered different destinations and the packet forwarding logic determines which route to use.

- If the FTD device learns about multiple paths to the same destination from a single routing protocol, such as RIP, the route with the better metric (as determined by the routing protocol) is entered into the routing table.

  Metrics are values associated with specific routes, ranking them from most preferred to least preferred. The parameters used to determine the metrics differ for different routing protocols. The path with the lowest metric is selected as the optimal path and installed in the routing table. If there are multiple paths to the same destination with equal metrics, load balancing is done on these equal cost paths.

- If the FTD device learns about a destination from more than one routing protocol, the administrative distances of the routes are compared, and the routes with lower administrative distance are entered into the routing table.

Administrative Distances for Routes

You can change the administrative distances for routes discovered by or redistributed into a routing protocol. If two routes from two different routing protocols have the same administrative distance, then the route with the lower *default* administrative distance is entered into the routing table. In the case of EIGRP and OSPF routes, if the EIGRP route and the OSPF route have the same administrative distance, then the EIGRP route is chosen by default.

Administrative distance is a route parameter that the FTD device uses to select the best path when there are two or more different routes to the same destination from two different routing protocols. Because the routing protocols have metrics based on algorithms that are different from the other protocols, it is not always possible to determine the best path for two routes to the same destination that were generated by different routing protocols.

Each routing protocol is prioritized using an administrative distance value. The following table shows the default administrative distance values for the routing protocols supported by the FTD device.

Table 67: Default Administrative Distance for Supported Routing Protocols

Route Source	Default Administrative Distance
Connected interface	0
VPN route	1
Static route	1
EIGRP Summary Route	5
External BGP	20
Internal EIGRP	90
OSPF	110
IS-IS	115
RIP	120
EIGRP external route	170
Internal and local BGP	200
Unknown	255

The smaller the administrative distance value, the more preference is given to the protocol. For example, if the FTD device receives a route to a certain network from both an OSPF routing process (default administrative distance - 110) and a RIP routing process (default administrative distance - 120), the FTD device chooses the OSPF route because OSPF has a higher preference. In this case, the router adds the OSPF version of the route to the routing table.

A VPN advertised route (V-Route/RRI)) is equivalent to a static route with the default administrative distance 1. But it has a higher preference as with the network mask 255.255.255.255.

In this example, if the source of the OSPF-derived route was lost (for example, due to a power shutdown), the FTD device would then use the RIP-derived route until the OSPF-derived route reappears.

The administrative distance is a local setting. For example, if you change the administrative distance of routes obtained through OSPF, that change would only affect the routing table for the FTD device on which the command was entered. The administrative distance is not advertised in routing updates.

Administrative distance does not affect the routing process. The routing processes only advertise the routes that have been discovered by the routing process or redistributed into the routing process. For example, the RIP routing process advertises RIP routes, even if routes discovered by the OSPF routing process are used in the routing table.

Backup Dynamic and Floating Static Routes

A backup route is registered when the initial attempt to install the route in the routing table fails because another route was installed instead. If the route that was installed in the routing table fails, the routing table maintenance process calls each routing protocol process that has registered a backup route and requests them to reinstall the route in the routing table. If there are multiple protocols with registered backup routes for the failed route, the preferred route is chosen based on administrative distance.

Because of this process, you can create floating static routes that are installed in the routing table when the route discovered by a dynamic routing protocol fails. A floating static route is simply a static route configured with a greater administrative distance than the dynamic routing protocols running on the FTD device. When the corresponding route discovered by a dynamic routing process fails, the static route is installed in the routing table.

How Forwarding Decisions Are Made

Forwarding decisions are made as follows:

- If the destination does not match an entry in the routing table, the packet is forwarded through the interface specified for the default route. If a default route has not been configured, the packet is discarded.
- If the destination matches a single entry in the routing table, the packet is forwarded through the interface associated with that route.
- If the destination matches more than one entry in the routing table, then the packet is forwarded out of the interface associated with the route that has the longer network prefix length.

For example, a packet destined for 192.168.32.1 arrives on an interface with the following routes in the routing table:

- 192.168.32.0/24 gateway 10.1.1.2
- 192.168.32.0/19 gateway 10.1.1.3

In this case, a packet destined to 192.168.32.1 is directed toward 10.1.1.2, because 192.168.32.1 falls within the 192.168.32.0/24 network. It also falls within the other route in the routing table, but 192.168.32.0/24 has the longest prefix within the routing table (24 bits verses 19 bits). Longer prefixes are always preferred over shorter ones when forwarding a packet.



Note Existing connections continue to use their established interfaces even if a new similar connection would result in different behavior due to a change in routes.

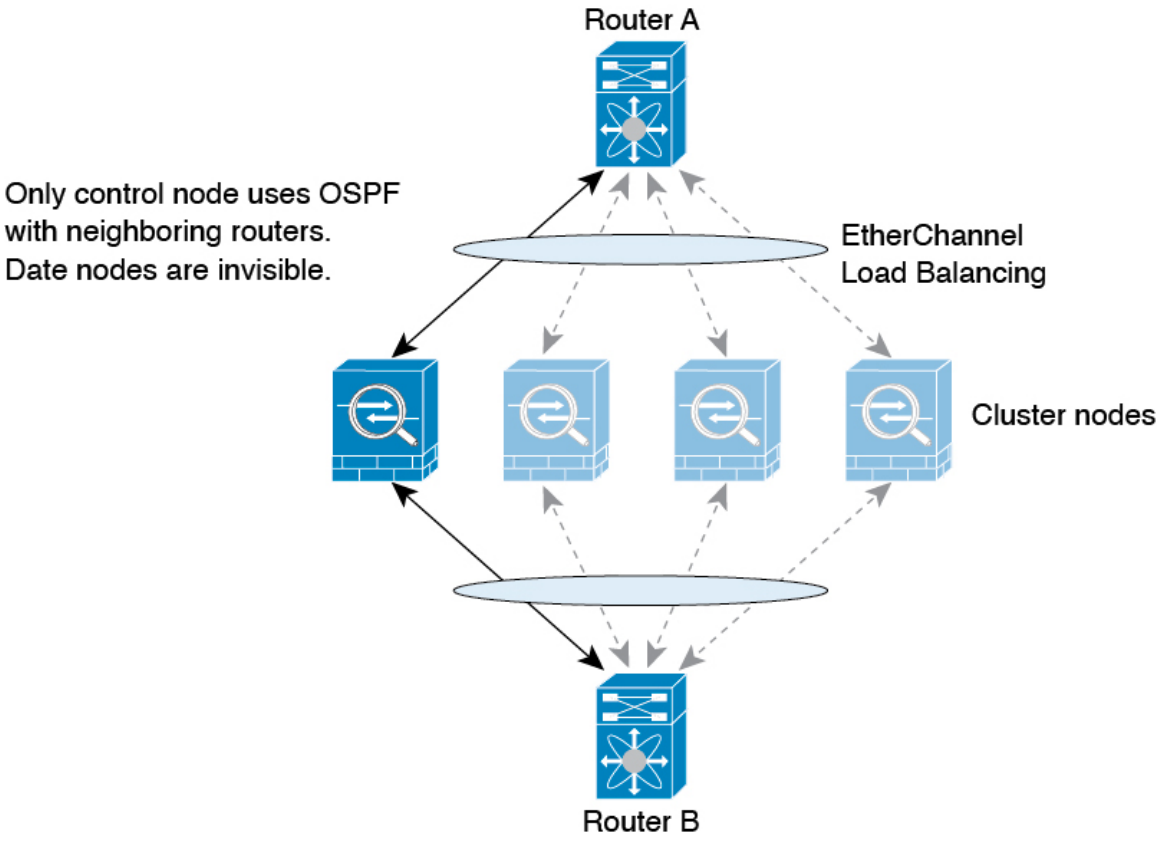
Dynamic Routing and High Availability

Dynamic routes are synchronized on the standby unit when the routing table changes on the active unit. This means that all additions, deletions, or changes on the active unit are immediately propagated to the standby unit. If the standby unit becomes active in an active/standby ready High Availability pair, it will already have an identical routing table as that of the former active unit because routes are synchronized as a part of the High Availability bulk synchronization and continuous replication processes.

Dynamic Routing in Clustering

The routing process only runs on the control node, and routes are learned through the control node and replicated to data nodes. If a routing packet arrives at a data node, it is redirected to the control node.

Figure 25: Dynamic Routing in Spanned EtherChannel Mode



After the data node learn the routes from the control node, each node makes forwarding decisions independently.

The OSPF LSA database is not synchronized from the control node to data nodes. If there is a control node switchover, the neighboring router will detect a restart; the switchover is not transparent. The OSPF process picks an IP address as its router ID. Although not required, you can assign a static router ID to ensure a consistent router ID is used across the cluster. See the OSPF Non-Stop Forwarding feature to address the interruption.

Routing Table for Management Traffic

Equal-Cost Multi-Path (ECMP) Routing

The FTD device supports Equal-Cost Multi-Path (ECMP) routing.

You can have up to 8 equal cost static or dynamic routes per interface. For example, you can configure multiple default routes on the outside interface that specify different gateways.

```
route for 0.0.0.0 0.0.0.0 through outside to 10.1.1.2
route for 0.0.0.0 0.0.0.0 through outside to 10.1.1.3
route for 0.0.0.0 0.0.0.0 through outside to 10.1.1.4
```

In this case, traffic is load-balanced on the outside interface between 10.1.1.2, 10.1.1.3, and 10.1.1.4. Traffic is distributed among the specified gateways based on an algorithm that hashes the source and destination IP addresses, incoming interface, protocol, source and destination ports.

About Route Maps

Route maps are used when redistributing routes into an OSPF, RIP, EIGRP or BGP routing process. They are also used when generating a default route into an OSPF routing process. A route map defines which of the routes from the specified routing protocol are allowed to be redistributed into the target routing process.

Route maps have many features in common with widely known ACLs. These are some of the traits common to both:

- They are an ordered sequence of individual statements, and each has a permit or deny result. Evaluation of an ACL or a route map consists of a list scan, in a predetermined order, and an evaluation of the criteria of each statement that matches. A list scan is aborted once the first statement match is found and an action associated with the statement match is performed.
- They are generic mechanisms. Criteria matches and match interpretation are dictated by the way that they are applied and the feature that uses them. The same route map applied to different features might be interpreted differently.

These are some of the differences between route maps and ACLs:

- Route maps are more flexible than ACLs and can verify routes based on criteria which ACLs can not verify. For example, a route map can verify if the type of route is internal.
- Each ACL ends with an implicit deny statement, by design convention. If the end of a route map is reached during matching attempts, the result depends on the specific application of the route map. Route maps that are applied to *redistribution* behave the same way as ACLs: if the route does not match any clause in a route map then the route redistribution is denied, as if the route map contained a deny statement at the end.

Permit and Deny Clauses

Route maps can have permit and deny clauses. The deny clause rejects route matches from redistribution. You can use an ACL as the matching criterion in the route map. Because ACLs also have permit and deny clauses, the following rules apply when a packet matches the ACL:

- ACL permit + route map permit: routes are redistributed.
- ACL permit + route map deny: routes are not redistributed.
- ACL deny + route map permit or deny: the route map clause is not matched, and the next route-map clause is evaluated.

Match and Set Clause Values

Each route map clause has two types of values:

- A match value selects routes to which this clause should be applied.

- A set value modifies information that will be redistributed into the target protocol.

For each route that is being redistributed, the router first evaluates the match criteria of a clause in the route map. If the match criteria succeeds, then the route is redistributed or rejected as dictated by the permit or deny clause, and some of its attributes might be modified by the values set from the set commands. If the match criteria fail, then this clause is not applicable to the route, and the software proceeds to evaluate the route against the next clause in the route map. Scanning of the route map continues until a clause is found that matches the route or until the end of the route map is reached.

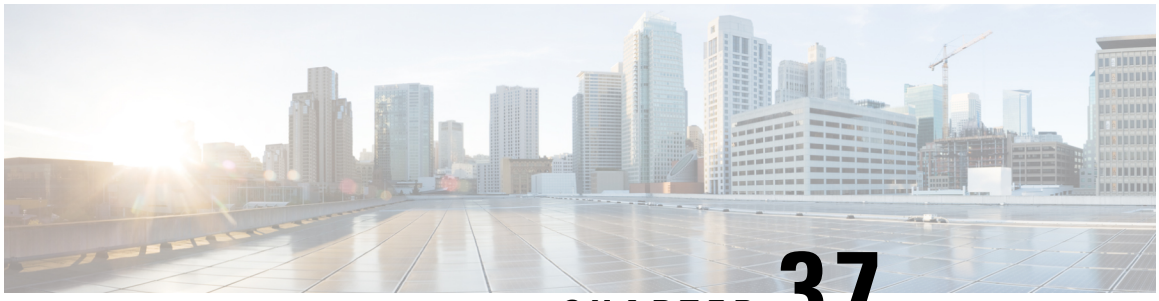
A match or set value in each clause can be missed or repeated several times, if one of these conditions exists:

- If several match entries are present in a clause, all must succeed for a given route in order for that route to match the clause (in other words, the logical AND algorithm is applied for multiple match commands).
- If a match entry refers to several objects in one entry, either of them should match (the logical OR algorithm is applied).
- If a match entry is not present, all routes match the clause.
- If a set entry is not present in a route map permit clause, then the route is redistributed without modification of its current attributes.



Note Do not configure a set entry in a route map deny clause because the deny clause prohibits route redistribution—there is no information to modify.

A route map clause without a match or set entry does perform an action. An empty permit clause allows a redistribution of the remaining routes without modification. An empty deny clause does not allow a redistribution of other routes (this is the default action if a route map is completely scanned, but no explicit match is found).



CHAPTER 37

Static and Default Routes for Firepower Threat Defense

This chapter describes how to configure static and default routes on the Firepower Threat Defense.

- About Static and Default Routes, on page 763
- Requirements and Prerequisites for Static Routes, on page 765
- Guidelines for Static and Default Routes, on page 766
- Add a Static Route, on page 766

About Static and Default Routes

To route traffic to a non-connected host or network, you must define a route to the host or network, either using static or dynamic routing. Generally, you must configure at least one static route: a default route for all traffic that is not routed by other means to a default network gateway, typically the next hop router.

Default Route

The simplest option is to configure a default static route to send all traffic to an upstream router, relying on the router to route the traffic for you. A default route identifies the gateway IP address to which the FTD device sends all IP packets for which it does not have a learned or static route. A default static route is simply a static route with 0.0.0.0/0 (IPv4) or ::/0 (IPv6) as the destination IP address.

You should always define a default route.

Because FTD uses separate routing tables for data traffic and for management traffic, you can optionally configure a default route for data traffic and another default route for management traffic. Note that from-the-device traffic uses either the management-only or data routing table by default depending on the type (see Routing Table for Management Traffic, on page 759), but will fall back to the other routing table if a route is not found. Default routes will always match traffic, and will prevent a fall back to the other routing table. In this case, you must specify the interface you want to use for egress traffic if that interface is not in the default routing table. The Diagnostic interface is included in the management-only table. The special Management interface uses a separate Linux routing table, and has its own default route. See the **configure network** commands.

Static Routes

You might want to use static routes in the following cases:

- Your networks use an unsupported router discovery protocol.
- Your network is small and you can easily manage static routes.
- You do not want the traffic or CPU overhead associated with routing protocols.
- In some cases, a default route is not enough. The default gateway might not be able to reach the destination network, so you must also configure more specific static routes. For example, if the default gateway is outside, then the default route cannot direct traffic to any inside networks that are not directly connected to the FTD device.
- You are using a feature that does not support dynamic routing protocols.

Route to null0 Interface to Drop Unwanted Traffic

Access rules let you filter packets based on the information contained in their headers. A static route to the null0 interface is a complementary solution to access rules. You can use a null0 route to forward unwanted or undesirable traffic so the traffic is dropped.

Static null0 routes have a favorable performance profile. You can also use static null0 routes to prevent routing loops. BGP can leverage the static null0 route for Remotely Triggered Black Hole routing.

Route Priorities

- Routes that identify a specific destination take precedence over the default route.
- When multiple routes exist to the same destination (either static or dynamic), then the administrative distance for the route determines priority. Static routes are set to 1, so they typically are the highest priority routes.
- When you have multiple static routes to the same destination with the same administrative distance, see Equal-Cost Multi-Path (ECMP) Routing, on page 759.
- For traffic emerging from a tunnel with the Tunneled option, this route overrides any other configured or learned default routes.

Transparent Firewall Mode and Bridge Group Routes

For traffic that originates on the FTD device and is destined through a bridge group member interface for a non-directly connected network, you need to configure either a default route or static routes so the FTD device knows out of which bridge group member interface to send traffic. Traffic that originates on the FTD device might include communications to a syslog server or SNMP server. If you have servers that cannot all be reached through a single default route, then you must configure static routes. For transparent mode, you cannot specify the BVI as the gateway interface; only member interfaces can be used. For bridge groups in routed mode, you must specify the BVI in a static route; you cannot specify a member interface. See #unique_1130 for more information.

Static Route Tracking

One of the problems with static routes is that there is no inherent mechanism for determining if the route is up or down. They remain in the routing table even if the next hop gateway becomes unavailable. Static routes are only removed from the routing table if the associated interface on the FTD device goes down.

The static route tracking feature provides a method for tracking the availability of a static route and installing a backup route if the primary route should fail. For example, you can define a default route to an ISP gateway and a backup default route to a secondary ISP in case the primary ISP becomes unavailable.

The FTD device implements static route tracking by associating a static route with a monitoring target host on the destination network that the FTD device monitors using ICMP echo requests. If an echo reply is not received within a specified time period, the host is considered down, and the associated route is removed from the routing table. An untracked backup route with a higher metric is used in place of the removed route.

When selecting a monitoring target, you need to make sure that it can respond to ICMP echo requests. The target can be any network object that you choose, but you should consider using the following:

- The ISP gateway (for dual ISP support) address
- The next hop gateway address (if you are concerned about the availability of the gateway)
- A server on the target network, such as a syslog server, that the FTD device needs to communicate with
- A persistent network object on the destination network



Note A PC that may be shut down at night is not a good choice.

You can configure static route tracking for statically defined routes or default routes obtained through DHCP or PPPoE. You can only enable PPPoE clients on multiple interfaces with route tracking configured.

Requirements and Prerequisites for Static Routes

Model Support

FTD

Supported Domains

Any

User Roles

Admin

Network Admin

Guidelines for Static and Default Routes

Firewall Mode and Bridge Groups

- In transparent mode, static routes must use the bridge group member interface as the gateway; you cannot specify the BVI.
- In routed mode, you must specify the BVI as the gateway; you cannot specify the member interface.
- Static route tracking is not supported for bridge group member interfaces or on the BVI.

Supported Network Address

- Static route tracking is not supported for IPv6.
- ASA does not support CLASS E routing. Hence, CLASS E network is not routable as static routes.

Clustering and Multiple Context Mode

- In clustering, static route tracking is only supported on the primary unit.
- Static route tracking is not supported in multiple context mode.

ASP and RIB Route Entries

All routes and its distance installed on the device are captured in the ASP routing table. This is common for all static and dynamic routing protocols. Only the best distance route is captured in the RIB table.

Add a Static Route

A static route defines where to send traffic for specific destination networks. You should at a minimum define a default route. A default route is simply a static route with 0.0.0.0/0 as the destination IP address.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Click **Routing**.

Step 3 Select **Static Route**.

Step 4 Click **Add Routes**.

Step 5 Click **IPv4** or **IPv6** depending on the type of static route that you are adding.

Step 6 Choose the **Interface** to which this static route applies.

For transparent mode, choose a bridge group member interface name. For routed mode with bridge groups, you can choose either the bridge group member interface for the BVI name. To “black hole” unwanted traffic, choose the **Null0** interface.

Step 7 In the **Available Network** list, choose the destination network.

To define a default route, create an object with the address 0.0.0.0/0 and select it here.

Step 8 In the **Gateway** or **IPv6 Gateway** field, enter or choose the gateway router which is the next hop for this route. You can provide an IP address or a Networks/Hosts object.

Step 9 In the **Metric** field, enter the number of hops to the destination network. Valid values range from 1 to 255; the default value is 1. The metric is a measurement of the "expense" of a route, based on the number of hops (hop count) to the network on which a specific host resides. Hop count is the number of networks that a network packet must traverse, including the destination network, before it reaches its final destination. The metric is used to compare routes among different routing protocols. The default administrative distance for static routes is 1, giving it precedence over routes discovered by dynamic routing protocols but not directly connected routes. The default administrative distance for routes discovered by OSPF is 110. If a static route has the same administrative distance as a dynamic route, the static route takes precedence. Connected routes always take precedence over static or dynamically discovered routes.

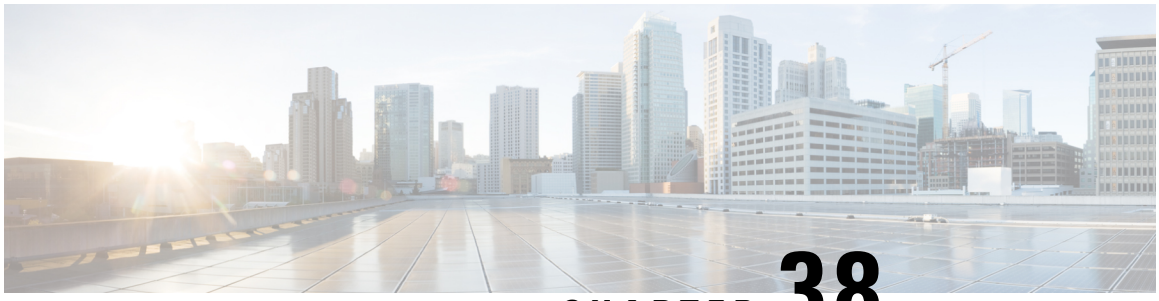
Step 10 (Optional) For a default route, click the **Tunneled** checkbox to define a separate default route for VPN traffic.

You can define a separate default route for VPN traffic if you want your VPN traffic to use a different default route than your non VPN traffic. For example, traffic incoming from VPN connections can be easily directed towards internal networks, while traffic from internal networks can be directed towards the outside. When you create a default route with the tunneled option, all traffic from a tunnel terminating on the device that cannot be routed using learned or static routes, is sent to this route. You can configure only one default tunneled gateway per device. ECMP for tunneled traffic is not supported.

Step 11 (IPv4 static route only) To monitor route availability, enter or choose the name of an SLA (service level agreement) Monitor object that defines the monitoring policy, in the **Route Tracking** field.

See SLA Monitor Objects, on page 461.

Step 12 Click **Ok**.



CHAPTER 38

OSPF for Firepower Threat Defense

This chapter describes how to configure the Firepower Threat Defense to route data, perform authentication, and redistribute routing information using the Open Shortest Path First (OSPF) routing protocol.

- OSPF for Firepower Threat Defense, on page 769
- Requirements and Prerequisites for OSPF, on page 772
- Guidelines for OSPF, on page 772
- Configure OSPFv2, on page 774
- Configure OSPFv3, on page 786

OSPF for Firepower Threat Defense

This chapter describes how to configure the Firepower Threat Defense to route data, perform authentication, and redistribute routing information using the Open Shortest Path First (OSPF) routing protocol.

About OSPF

OSPF is an interior gateway routing protocol that uses link states rather than distance vectors for path selection. OSPF propagates link-state advertisements rather than routing table updates. Because only LSAs are exchanged instead of the entire routing tables, OSPF networks converge more quickly than RIP networks.

OSPF uses a link-state algorithm to build and calculate the shortest path to all known destinations. Each router in an OSPF area contains an identical link-state database, which is a list of each of the router usable interfaces and reachable neighbors.

The advantages of OSPF over RIP include the following:

- OSPF link-state database updates are sent less frequently than RIP updates, and the link-state database is updated instantly, rather than gradually, as stale information is timed out.
- Routing decisions are based on cost, which is an indication of the overhead required to send packets across a certain interface. The FTD device calculates the cost of an interface based on link bandwidth rather than the number of hops to the destination. The cost can be configured to specify preferred paths.

The disadvantage of shortest path first algorithms is that they require a lot of CPU cycles and memory.

The FTD device can run two processes of OSPF protocol simultaneously on different sets of interfaces. You might want to run two processes if you have interfaces that use the same IP addresses (NAT allows these interfaces to coexist, but OSPF does not allow overlapping addresses). Or you might want to run one process

on the inside and another on the outside, and redistribute a subset of routes between the two processes. Similarly, you might need to segregate private addresses from public addresses.

You can redistribute routes into an OSPF routing process from another OSPF routing process, a RIP routing process, or from static and connected routes configured on OSPF-enabled interfaces.

The FTD device supports the following OSPF features:

- Intra-area, inter-area, and external (Type I and Type II) routes.
- Virtual links.
- LSA flooding.
- Authentication to OSPF packets (both password and MD5 authentication).
- Configuring the FTD device as a designated router or a designated backup router. The FTD device also can be set up as an ABR.
- Stub areas and not-so-stubby areas.
- Area boundary router Type 3 LSA filtering.

OSPF supports MD5 and clear text neighbor authentication. Authentication should be used with all routing protocols when possible because route redistribution between OSPF and other protocols (such as RIP) can potentially be used by attackers to subvert routing information.

If NAT is used, if OSPF is operating on public and private areas, and if address filtering is required, then you need to run two OSPF processes—one process for the public areas and one for the private areas.

A router that has interfaces in multiple areas is called an Area Border Router (ABR). A router that acts as a gateway to redistribute traffic between routers using OSPF and routers using other routing protocols is called an Autonomous System Boundary Router (ASBR).

An ABR uses LSAs to send information about available routes to other OSPF routers. Using ABR Type 3 LSA filtering, you can have separate private and public areas with the ASA acting as an ABR. Type 3 LSAs (inter-area routes) can be filtered from one area to other, which allows you to use NAT and OSPF together without advertising private networks.



Note Only Type 3 LSAs can be filtered. If you configure the FTD device as an ASBR in a private network, it will send Type 5 LSAs describing private networks, which will get flooded to the entire AS, including public areas.

If NAT is employed but OSPF is only running in public areas, then routes to public networks can be redistributed inside the private network, either as default or Type 5 AS external LSAs. However, you need to configure static routes for the private networks protected by the FTD device. Also, you should not mix public and private networks on the same FTD device interface.

You can have two OSPF routing processes, one RIP routing process, and one EIGRP routing process running on the FTD device at the same time.

OSPF Support for Fast Hello Packets

The OSPF Support for Fast Hello Packets feature provides a way to configure the sending of hello packets in intervals less than one second. Such a configuration would result in faster convergence in an Open Shortest Path First (OSPF) network.

Prerequisites for OSPF Support for Fast Hello Packets

OSPF must be configured in the network already or configured at the same time as the OSPF Support for Fast Hello Packets feature.

OSPF Hello Interval and Dead Interval

OSPF hello packets are packets that an OSPF process sends to its OSPF neighbors to maintain connectivity with those neighbors. The hello packets are sent at a configurable interval (in seconds). The defaults are 10 seconds for an Ethernet link and 30 seconds for a non broadcast link. Hello packets include a list of all neighbors for which a hello packet has been received within the dead interval. The dead interval is also a configurable interval (in seconds), and defaults to four times the value of the hello interval. The value of all hello intervals must be the same within a network. Likewise, the value of all dead intervals must be the same within a network.

These two intervals work together to maintain connectivity by indicating that the link is operational. If a router does not receive a hello packet from a neighbor within the dead interval, it will declare that neighbor to be down.

OSPF Fast Hello Packets

OSPF fast hello packets refer to hello packets being sent at intervals of less than 1 second. To understand fast hello packets, you should already understand the relationship between OSPF hello packets and the dead interval. See OSPF Hello Interval and Dead Interval, on page 771.

OSPF fast hello packets are achieved by using the ospf dead-interval command. The dead interval is set to 1 second, and the hello-multiplier value is set to the number of hello packets you want sent during that 1 second, thus providing subsecond or "fast" hello packets.

When fast hello packets are configured on the interface, the hello interval advertised in the hello packets that are sent out this interface is set to 0. The hello interval in the hello packets received over this interface is ignored.

The dead interval must be consistent on a segment, whether it is set to 1 second (for fast hello packets) or set to any other value. The hello multiplier need not be the same for the entire segment as long as at least one hello packet is sent within the dead interval.

Benefits of OSPF Fast Hello Packets

The benefit of the OSPF Fast Hello Packets feature is that your OSPF network will experience faster convergence time than it would without fast hello packets. This feature allows you to detect lost neighbors within 1 second. It is especially useful in LAN segments, where neighbor loss might not be detected by the Open System Interconnection (OSI) physical layer and data-link layer.

Implementation Differences Between OSPFv2 and OSPFv3

OSPFv3 is not backward compatible with OSPFv2. To use OSPF to route both IPv4 and IPv6 traffic, you must run both OSPFv2 and OSPFv3 at the same time. They coexist with each other, but do not interact with each other.

The additional features that OSPFv3 provides include the following:

- Protocol processing per link.
- Removal of addressing semantics.
- Addition of flooding scope.
- Support for multiple instances per link.
- Use of the IPv6 link-local address for neighbor discovery and other features.
- LSAs expressed as prefix and prefix length.
- Addition of two LSA types.
- Handling of unknown LSA types.
- Authentication support using the IPsec ESP standard for OSPFv3 routing protocol traffic, as specified by RFC-4552.

Requirements and Prerequisites for OSPF

Model Support

FTD

FTDv

Supported Domains

Any

User Roles

Admin

Network Admin

Guidelines for OSPF

Firewall Mode Guidelines

OSPF supports routed firewall mode only. OSPF does not support transparent firewall mode.

High Availability Guidelines

OSPFv2 and OSPFv3 support Stateful High Availability.

IPv6 Guidelines

- OSPFv2 does not support IPv6.
- OSPFv3 supports IPv6.
- OSPFv3 uses IPv6 for authentication.
- The FTD device installs OSPFv3 routes into the IPv6 RIB, provided it is the best route.

OSPFv3 Hello Packets and GRE

Typically, OSPF traffic does not pass through GRE tunnel. When OSPFv3 on IPv6 is encapsulated inside GRE, the IPv6 header validation for security check such as Multicast Destination fails. The packet is dropped due to the implicit security check validation, as this packet has destination IPv6 multicast.

You may define a pre-filter rule to bypass GRE traffic. However, with pre-filter rule, inner packets would not be interrogated by the inspection engine.

Clustering Guidelines

- OSPFv3 encryption is not supported. An error message appears if you try to configure OSPFv3 encryption in a clustering environment.
- In Spanned interface mode, dynamic routing is not supported on management-only interfaces.
- When a control role change occurs in the cluster, the following behavior occurs:
  - In spanned interface mode, the router process is active only on the control unit and is in a suspended state on the data units. Each cluster unit has the same router ID because the configuration has been synchronized from the control unit. As a result, a neighboring router does not notice any change in the router ID of the cluster during a role change.

Multiprotocol Label Switching (MPLS) and OSPF Guidelines

When a MPLS-configured router sends Link State (LS) update packets containing opaque Type-10 link-state advertisements (LSAs) that include an MPLS header, authentication fails and the appliance silently drops the update packets, rather than acknowledging them. Eventually the peer router will terminate the neighbor relationship because it has not received any acknowledgments.

Make sure that non-stop forwarding (NSF) is disabled on the appliance to ensure that the neighbor relationship remains stable:

- Navigate to the **Non Stop Forwarding** page in FMC(**Devices** > **Device Management (select the desired device)** > **Routing** > **OSPF** > **Advanced** > **Non Stop Forwarding**).

  Ensure the **Non Stop Forwarding Capability** boxes are not checked.

Note The Firepower 4100/9300 models may have high latency when using MPLS because they lack load balancing across multiple receiving queues.

Route Redistribution Guidelines

- Redistribution of route maps with IPv4 or IPv6 prefix list on OSPFv2 or OSPFv3 is not supported. Use an access list in the route map on OSPF for redistribution.
- When OSPF is configured on a device that is a part of EIGRP network or vice versa, ensure that OSPF-router is configured to tag the route (EIGRP does not support route tag yet).

  When redistributing OSPF into EIGRP and EIGRP into OSPF, a routing loop occurs when there is an outage on one of the links, interfaces, or even when the route originator is down. To prevent the redistribution of routes from one domain back into the same domain, a router can tag a route that belongs to a domain while it is redistributing, and those routes can be filtered on the remote router based on the same tag. Because the routes will not be installed into the routing table, they will not be redistributed back into the same domain.

Additional Guidelines

- OSPFv2 and OSPFv3 support multiple instances on an interface.
- OSPFv3 supports encryption through ESP headers in a non-clustered environment.
- OSPFv3 supports Non-Payload Encryption.
- OSPFv2 supports Cisco NSF Graceful Restart and IETF NSF Graceful Restart mechanisms as defined in RFCs 4811, 4812 & 3623 respectively.
- OSPFv3 supports Graceful Restart mechanism as defined in RFC 5187.
- There is a limit to the number of intra area (type 1) routes that can be distributed. For these routes, a single type-1 LSA contains all prefixes. Because the system has a limit of 35 KB for packet size, 3000 routes result in a packet that exceeds the limit. Consider 2900 type 1 routes to be the maximum number supported.
- To avoid adjacency flaps due to route updates being dropped if the route update is larger than the minimum MTU on the link, ensure that you configure the same MTU on the interfaces on both sides of the link.

Configure OSPFv2

This section describes the tasks involved in configuring an OSPFv2 routing process.

Configure OSPF Areas, Ranges, and Virtual Links

You can configure several OSPF area parameters, which include setting authentication, defining stub areas, and assigning specific costs to the default summary route. You can enable up to two OSPF process instances. Each OSPF process has its own associated areas and networks. Authentication provides password-based protection against unauthorized access to an area.

Stub areas are areas into which information on external routes is not sent. Instead, there is a default external route generated by the ABR into the stub area for destinations outside the autonomous system. To take advantage of the OSPF stub area support, default routing must be used in the stub area.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Click **Routing**.

Step 3 Click **OSPF**.

Step 4 Select **Process 1**. You can enable up to two OSPF process instances for each context. You must choose an OSPF process to be able to configure the Area parameters.

Step 5 Choose the OSPF role from the drop-down list, and enter a description for it in the next field. The options are Internal, ABR, ASBR, and ABR and ASBR. See About OSPF, on page 769 for a description of the OSPF roles.

Step 6 Select **Area** > **Add**.

You can click **Edit** (), or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 7 Configure the following area options for each OSPF process:

- **OSPF Process**—Choose 1 or 2.
- **Area ID**—Designation of the area for which routes are to be summarized.
- **Area Type**—Choose one of the following:
  - **Normal**—(Default) Standard OSPF area.
  - **Stub**—A stub area does not have any routers or areas beyond it. Stub areas prevent Autonomous System (AS) External LSAs (Type 5 LSAs) from being flooded into the stub area. When you create a stub area, you can prevent summary LSAs (Types 3 and 4) from being flooded into the area by NOT checking the **Summary Stub** check box.
  - **NSSA**—Makes the area a not-so-stubby area (NSSA). NSSAs accept Type 7 LSAs. You can disable route redistribution by NOT checking the **Redistribute check box** and checking the **Default Information Originate** check box. You can prevent summary LSAs from being flooded into the area by NOT checking the **Summary NSSA** check box.
- **Metric Value**—The metric used for generating the default route. The default value is 10. Valid metric values range from 0 through 16777214.
- **Metric Type**—The metric type is the external link type that is associated with the default route that is advertised into the OSPF routing domain. The available options are 1 for a Type 1 external route or 2 for a Type 2 external route.
- **Available Network**—Choose one of the available networks and click **Add**, or click **Add** () to add a new network object. See Network Objects, on page 402 for the procedure for adding networks.
- **Authentication**—Choose the OSPF authentication:
  - **None**—(Default) Disables OSPF area authentication.

- **Password**—Provides a clear text password for area authentication, which is not recommended where security is a concern.
- **MD5**—Allows MD5 authentication.

- **Default Cost**—The default cost for the OSPF area, which is used to determine the shortest paths to the destination. Valid values range from 0 through 65535. The default value is 1.

Step 8 Click **OK** to save the area configuration.

Step 9 Select **Range** > **Add**.

- Choose one of the available networks and whether to advertise, or,
- Click **Add** ( ) to add a new network object. See Network Objects, on page 402 for the procedure for adding networks.

Step 10 Click **OK** to save the range configuration.

Step 11 Select **Virtual Link**, click **Add**, and configure the following options for each OSPF process:

- **Peer Router**—Choose the IP address of the peer router. To add a new peer router, click **Add** ( ). See Network Objects, on page 402 for the procedure for adding networks.
- **Hello Interval**—The time in seconds between the hello packets sent on an interface. The hello interval is an unsigned integer that is to be advertised in the hello packets. The value must be the same for all routers and access servers on a specific network. Valid values range from 1 through 65535. The default is 10.

  The smaller the hello interval, the faster topological changes are detected, but the more traffic is sent on the interface.
- **Transmit Delay**—The estimated time in seconds that is required to send an LSA packet on the interface. The integer value must be greater than zero. Valid values range from 1 through 8192. The default is 1.

  LSAs in the update packet have their own ages incremented by this amount before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links.
- **Retransmit Interval**—The time in seconds between LSA retransmissions for adjacencies that belong to the interface. The retransmit interval is the expected round-trip delay between any two routers on the attached network. The value must be greater than the expected round-trip delay, and can range from 1 through 65535. The default is 5.

  When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links.
- **Dead Interval**—The time in seconds that hello packets are not seen before a neighbor indicates that the router is down. The dead interval is an unsigned integer. The default is four times the hello interval, or 40 seconds. The value must be the same for all routers and access servers that are attached to a common network. Valid values range from 1 through 65535.
- **Authentication**—Choose the OSPF virtual link authentication from the following:

- **None**—(Default) Disables virtual link area authentication.
- **Area Authentication**—Enables area authentication using MD5. Click **Add**, and enter the key ID, key, confirm the key, and then click **OK**.
- **Password**—Provides a clear text password for virtual link authentication, which is not recommended where security is a concern.
- **MD5**—Allows MD5 authentication. Click **Add**, and enter the key ID, key, confirm the key, and then click **OK**.

  Note Ensure to enter only numbers as the MD5 key ID.

Step 12 Click **OK** to save the virtual link configuration.

Step 13 Click **Save** on the Routing page to save your changes.

What to do next

Continue with Configure OSPF Redistribution.

Configure OSPF Redistribution

The Firepower Threat Defense device can control the redistribution of routes between the OSPF routing processes. The rules for redistributing routes from one routing process into an OSPF routing process are displayed. You can redistribute routes discovered by RIP and BGP into the OSPF routing process. You can also redistribute static and connected routes into the OSPF routing process.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Click **Routing**.

Step 3 Click **OSPF**.

Step 4 Select **Redistribution** > **Add**.

You can click **Edit** ( ), or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 5 Configure the following redistribution options for each OSPF process:

-
- **Route Type**—Choose one of the following types:
  - **Static**—Redistributes static routes to the OSPF routing process.
  - **Connected**—Redistributes connected routes (routes established automatically by virtue of having the IP address enabled on the interface) to the OSPF routing process. Connected routes are redistributed as external to the device. You can select whether to use subnets under the Optional list.

- **OSPF**—Redistributes routes from another OSPF routing process, for example, internal, external 1 and 2, NSSA external 1 and 2, or whether to use subnets. You can select these options under the Optional list.
- **BGP**—Redistribute routes from the BGP routing process. Add the AS number and whether to use subnets.
- **RIP**—Redistributes routes from the RIP routing process. You can select whether to use subnets under the Optional list.

- **Metric Value**—Metric value for the routes being distributed. The default value is 10. Valid values range from 0 to 16777214.

  When redistributing from one OSPF process to another OSPF process on the same device, the metric will be carried through from one process to the other if no metric value is specified. When redistributing other processes to an OSPF process, the default metric is 20 when no metric value is specified.
- **Metric Type**—The metric type is the external link type that is associated with the default route that is advertised into the OSPF routing domain. The available options are 1 for a Type 1 external route or 2 for a Type 2 external route.
- **Tag Value**—Tag specifies the 32-bit decimal value attached to each external route that is not used by OSPF itself, but which may be used to communicate information between ASBRs. If none is specified, then the remote autonomous system number is used for routes from BGP and EGP. For other protocols, zero is used. Valid values are from 0 to 4294967295.
- **RouteMap**—Checks for filtering the importing of routes from the source routing protocol to the current routing protocol. If this parameter is not specified, all routes are redistributed. If this parameter is specified, but no route map tags are listed, no routes are imported. Or you can add a new route map by clicking **Add** (). See Route Maps to add a new route map.

Step 6 Click **OK** to save the redistribution configuration.

Step 7 Click **Save** on the Routing page to save your changes.

What to do next

Continue with Configure OSPF Inter-Area Filtering, on page 778.

Configure OSPF Inter-Area Filtering

ABR type 3 LSA filtering extends the capability of an ABR that is running OSPF to filter type 3 LSAs between different OSPF areas. Once a prefix list is configured, only the specified prefixes are sent from one OSPF area to another OSPF area. All other prefixes are restricted to their OSPF area. You can apply this type of area filtering to traffic going into or coming out of an OSPF area, or to both the incoming and outgoing traffic for that area.

When multiple entries of a prefix list match a given prefix, the entry with the lowest sequence number is used. For efficiency, you may want to put the most common matches or denials near the top of the list by manually assigning them a lower sequence number. By default, sequence numbers are automatically generated in increments of 5, beginning with 5.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Click **Routing**.

Step 3 Click **OSPF**.

Step 4 Select **InterArea** > **Add**.

You can click **Edit** ( ), or use the right-click menu to cut, copy, past, insert, and delete inter-areas.

Step 5 Configure the following inter-area filtering options for each OSPF process:

- **OSPF Process**—Choose 1 or 2.
- **Area ID**—The area for which routes are to be summarized.
- **PrefixList**—The name of the prefix. To add a new prefix list object, see Step 5.
- **Traffic Direction**—Inbound or outbound. Choose Inbound to filter LSAs coming into an OSPF area, or Outbound to filter LSAs coming out of an OSPF area. If you are editing an existing filter entry, you cannot modify this setting.

Step 6 Click **Add** ( ), and enter a name for the new prefix list, and whether to allow overrides.

You must configure a prefix list before you can configure a prefix rule.

Step 7 Click **Add** to configure prefix rules, and configure the following parameters:

- **Action**—Select **Block** or **Allow** for the redistribution access.
- **Sequence No**—The routing sequence number. By default, sequence numbers are automatically generated in increments of 5, beginning with 5.
- **IP Address**—Specify the prefix number in the format of IP address/mask length.
- **Min Prefix Length**—(Optional) The minimum prefix length.
- **Max Prefix Length**—(Optional) The maximum prefix length.

Step 8 Click **OK** to save the inter-area filtering configuration.

Step 9 Click **Save** on the Routing page to save your changes.

What to do next

Continue with Configure OSPF Filter Rules, on page 779.

Configure OSPF Filter Rules

You can configure ABR Type 3 LSA filters for each OSPF process. ABR Type 3 LSA filters allow only specified prefixes to be sent from one area to another area and restrict all other prefixes. You can apply this type of area filtering out of a specific OSPF area, into a specific OSPF area, or into and out of the same OSPF

area at the same time. OSPF ABR Type 3 LSA filtering improves your control of route distribution between OSPF areas.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Click **Routing**.

Step 3 Click **OSPF**.

Step 4 Select **Filter Rule** > **Add**.

You can click **Edit** (), or use the right-click menu to cut, copy, past, insert, and delete filter rules.

Step 5 Configure the following filter rule options for each OSPF process:

- **OSPF Process**—Choose 1 or 2.
- **Access List**—The access list for this OSPF process. To add a new standard access list object, click **Add** () and see Configure Standard ACL Objects, on page 469.
- **Traffic Direction**—Choose In or Out for the traffic direction being filtered. Choose In to filter LSAs coming into an OSPF area, or Out to filter LSAs coming out of an OSPF area. If you are editing an existing filter entry, you cannot modify this setting.
- **Interface**—The interface for this filter rule.

Step 6 Click **OK** to save the filter rule configuration.

Step 7 Click **Save** on the Routing page to save your changes.

What to do next

Continue with Configure OSPF Summary Addresses, on page 780.

Configure OSPF Summary Addresses

When routes from other protocols are redistributed into OSPF, each route is advertised individually in an external LSA. However, you can configure the Firepower Threat Defense device to advertise a single route for all the redistributed routes that are included for a specified network address and mask. This configuration decreases the size of the OSPF link-state database. Routes that match the specified IP address mask pair can be suppressed. The tag value can be used as a match value for controlling redistribution through route maps.

Routes learned from other routing protocols can be summarized. The metric used to advertise the summary is the smallest metric of all the more specific routes. Summary routes help reduce the size of the routing table.

Using summary routes for OSPF causes an OSPF ASBR to advertise one external route as an aggregate for all redistributed routes that are covered by the address. Only routes from other routing protocols that are being redistributed into OSPF can be summarized.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Click **Routing**.

Step 3 Click **OSPF**.

Step 4 Select **Summary Address** > **Add**.

You can click **Edit** ( ) to edit, or use the right-click menu to cut, copy, past, insert, and delete summary addresses.

Step 5 Configure the following summary address options for each OSPF process:

- **OSPF Process**—Choose 1 or 2.
- **Available Network**—The IP address of the summary address. Select one from the Available networks list and click **Add**, or to add a new network, click **Add** ( ). See Network Objects, on page 402 for the procedure for adding networks.
- **Tag**—A 32-bit decimal value that is attached to each external route. This value is not used by OSPF itself, but may be used to communicate information between ASBRs.
- **Advertise—Advertises** the summary route. Uncheck this check box to suppress routes that fall under the summary address. By default, this check box is checked.

Step 6 Click **OK** to save the summary address configuration.

Step 7 Click **Save** on the Routing page to save your changes.

What to do next

Continue with Configure OSPF Interfaces and Neighbors, on page 781.

Configure OSPF Interfaces and Neighbors

You can change some interface-specific OSPFv2 parameters, if necessary. You are not required to change any of these parameters, but the following interface parameters must be consistent across all routers in an attached network: the hello interval, the dead interval, and the authentication key. If you configure any of these parameters, be sure that the configurations for all routers on your network have compatible values.

You need to define static OSPFv2 neighbors to advertise OSPFv2 routes over a point-to-point, non-broadcast network. This feature lets you broadcast OSPFv2 advertisements across an existing VPN connection without having to encapsulate the advertisements in a GRE tunnel.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Click **Routing**.

Step 3 Click **OSPF**.

Step 4 Select **Interface** > **Add**.

You can click **Edit** (), or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 5 Configure the following Interface options for each OSPF process:

- **Interface**—The interface you are configuring.
- **Default Cost**—The cost of sending a packet through the interface. The default value is 10.
- **Priority—Determines** the designated router for a network. Valid values range from 0 to 255. The default value is 1. Entering 0 for this setting makes the router ineligible to become the designated router or backup designated router.

  When two routers connect to a network, both attempt to become the designated router. The device with the higher router priority becomes the designated router. If there is a tie, the router with the higher router ID becomes the designated router. This setting does not apply to interfaces that are configured as point-to-point interfaces.
- **MTU Ignore—OSPF** checks whether neighbors are using the same MTU on a common interface. This check is performed when neighbors exchange DBD packets. If the receiving MTU in the DBD packet is higher than the IP MTU configured on the incoming interface, OSPF adjacency is not established.
- **Database Filter**—Use this setting to filter the outgoing LSA interface during synchronization and flooding. By default, OSPF floods new LSAs over all interfaces in the same area, except the interface on which the LSA arrives. In a fully meshed topology, this flooding can waste bandwidth and lead to excessive link and CPU usage. Checking this check box prevents OSPF flooding of the LSA on the selected interface.
- **Hello Interval**—Specifies the interval, in seconds, between hello packets sent on an interface. Valid values range 1–8192 seconds. The default value is 10 seconds.

  The smaller the hello interval, the faster topological changes are detected, but more traffic is sent on the interface. This value must be the same for all routers and access servers on a specific interface.
- **Transmit Delay**—Estimated time in seconds to send an LSA packet on the interface. Valid values range 1–65535 seconds. The default is 1 second.

  LSAs in the update packet have their ages increased by the amount specified by this field before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links.
- **Retransmit Interval**—Time in seconds between LSA retransmissions for adjacencies that belong to the interface. The time must be greater than the expected round-trip delay between any two routers on the attached network. Valid values range from 1 to 65535 seconds. The default is 5 seconds.

  When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links.
- **Dead Interval**—Time period in seconds for which hello packets must not be seen before neighbors indicate that the router is down. The value must be the same for all nodes on the network and can range 1–65535.
- **Hello Multiplier**—Specifies the number of Hello packets to be sent per second. Valid values are 3–20.

- **Point-to-Point**—Lets you transmit OSPF routes over VPN tunnels.
- **Authentication**—Choose the OSPF interface authentication from the following:
  - **None**—(Default) Disables interface authentication.
  - **Area Authentication**—Enables interface authentication using MD5. Click **Add**, and enter the key ID, key, confirm the key, and then click **OK**.
  - **Password**—Provides a clear text password for virtual link authentication, which is not recommended where security is a concern.
  - **MD5**—Allows MD5 authentication. Click **Add**, and enter the key ID, key, confirm the key, and then click **OK**.

    Note Ensure to enter only numbers as the MD5 key ID.
- **Enter Password**—The password you configure if you choose Password as the type of authentication.
- **Confirm Password**—Confirm the password that you chose.

Step 6 Select **Neighbor** > **Add**.

You can click **Edit** (), or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 7 Configure the following parameters for each OSPF process:

- **OSPF Process**—Choose 1 or 2.
- **Neighbor**—Choose one of the neighbors in the drop-down list, or click **Add** () to add a new neighbor; enter the name, description, network, whether to allow overrides, and then click **Save**.
- **Interface**—Choose the interface associated with the neighbor.

Step 8 Click **OK** to save the neighbor configuration.

Step 9 Click **Save** on the Routing page to save your changes.

Configure OSPF Advanced Properties

The Advanced Properties allows you to configure options, such as syslog message generation, administrative route distances, an LSA timer, and graceful restarts.

Graceful Restarts

The Firepower Threat Defense device may experience some known failure situations that should not affect packet forwarding across the switching platform. The Non-Stop Forwarding (NSF) capability allows data forwarding to continue along known routes, while the routing protocol information is being restored. This capability is useful when there is a scheduled hitless software upgrade. You can configure graceful restart on OSPFv2 by using either using NSF Cisco (RFC 4811 and RFC 4812) or NSF IETF (RFC 3623).

Note NSF capability is also useful in HA mode and clustering.

Configuring the NSF graceful-restart feature involves two steps; configuring capabilities and configuring a device as NSF-capable or NSF-aware. A NSF-capable device can indicate its own restart activities to neighbors and a NSF-aware device can help a restarting neighbor.

A device can be configured as NSF-capable or NSF-aware, depending on some conditions:

- A device can be configured as NSF-aware irrespective of the mode in which it is.
- A device has to be in either Failover or Spanned Etherchannel (L2) cluster mode to be configured as NSF-capable.
- For a device to be either NSF-aware or NSF-capable, it should be configured with the capability of handling opaque Link State Advertisements (LSAs)/ Link Local Signaling (LLS) block as required.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Click **Routing**.

Step 3 Click **OSPF** > **Advanced Settings**.

Step 4 Select **General**, and configure the following:

- **Router ID**—Choose **Automatic** or **IP address** for the router ID. If you choose IP address, enter the IP address in the IP Address field.
- **Ignore LSA MOSPF**—Suppresses syslog messages when the route receives unsupported LSA Type 6 multicast OSPF (MOSPF) packets.
- **RFC 1583 Compatible**—Configures RFC 1583 compatibility as the method used to calculate summary route costs. Routing loops can occur with RFC 1583 compatibility enabled. Disable it to prevent routing loops. All OSPF routers in an OSPF routing domain should have RFC compatibility set identically.
- **Adjacency Changes**—Defines the adjacency changes that cause syslog messages to be sent.

  By default, a syslog message is generated when an OSPF neighbor goes up or down. You can configure the router to send a syslog message when an OSPF neighbor goes down and also a syslog for each state.

  - **Log Adjacency Changes**—Causes the Firepower Threat Defense device to send a syslog message whenever an OSPF neighbor goes up or down. This setting is checked by default.
  - **Log Adjacency Change Details**—Causes the Firepower Threat Defense device to send a syslog message whenever any state change occurs, not just when a neighbor goes up or down. This setting is unchecked by default.

- **Administrative Route Distance**—Allows you to modify the settings that were used to configure administrative route distances for inter-area, intra-area, and external IPv6 routes. The administrative route distance is an integer from 1 to 254. The default is 110.
- **LSA Group Pacing**—Specifies the interval in seconds at which LSAs are collected into a group and refreshed, check summed, or aged. Valid values range from 10 to 1800. The default value is 240.

- **Enable Default Information Originate**—Check the **Enable** check box to generate a default external route into an OSPF routing domain and configure the following options:
  - **Always advertise the default route**—Ensures that the default route is always advertised.
  - **Metric Value**—Metric used for generating the default route. Valid metric values range from 0 to 16777214. The default value is 10.
  - **Metric Type**—The external link type that is associated with the default route that is advertised into the OSPFv3 routing domain. Valid values are 1 (Type 1 external route) and 2 (Type 2 external route). The default is Type 2 external route.
  - **RouteMap**—Choose the routing process that generates the default route if the route map is satisfied or click **Add** () to add a new one. See Route Maps to add a new route map.

Step 5 Click **OK** to save the general configuration.

Step 6 Select **Non Stop Forwarding**, and configure Cisco NSF graceful restart for OSPFv2, for an NSF-capable or NSF-aware device:

Note There are two graceful restart mechanisms for OSPFv2, Cisco NSF and IETF NSF. Only one of these graceful restart mechanisms can be configured at a time for an OSPF instance. An NSF-aware device can be configured as both Cisco NSF helper and IETF NSF helper but a NSF-capable device can be configured in either Cisco NSF or IETF NSF mode at a time for an OSPF instance.

a) Check the **Enable Cisco Non Stop Forwarding Capability** check box.
b) (Optional) Check the **Cancel NSF restart when non-NSF-aware neighboring networking devices are detected** check box if required.
c) (Optional) Make sure the **Enable Cisco Non Stop Forwarding Helper** mode check box is unchecked to disable the helper mode on an NSF-aware device.

Step 7 Configure IETF NSF Graceful Restart for OSPFv2, for an NSF-capable or NSF-aware device:

a) Check the **Enable IETF Non Stop Forwarding Capability** check box.
b) In the **Length of graceful restart interval (seconds)** field, enter the restart interval in seconds. The default value is 120 seconds. For a restart interval below 30 seconds, graceful restart will be terminated.
c) (Optional) Make sure the **Enable IETF nonstop forwarding (NSF) for helper mode** check box is unchecked to disable the IETF NSF helper mode on an NSF-aware device.
d) **Enable Strict Link State advertisement checking**—When enabled, it indicates that the helper router will terminate the process of restarting the router if it detects that there is a change to a LSA that would be flooded to the restarting router, or if there is a changed LSA on the retransmission list of the restarting router when the graceful restart process is initiated.
e) **Enable IETF Non Stop Forwarding**—Enables non stop forwarding, which allows for the forwarding of data packets to continue along known routes while the routing protocol information is being restored following a switchover. OSPF uses extensions to the OSPF protocol to recover its state from neighboring OSPF devices. For the recovery to work, the neighbors must support the NSF protocol extensions and be willing to act as "helpers" to the device that is restarting. The neighbors must also continue forwarding data traffic to the device that is restarting while protocol state recovery takes place.

Configure OSPFv3

This section describes the tasks involved in configuring an OSPFv3 routing process.

Configure OSPFv3 Areas, Route Summaries, and Virtual Links

To enable OSPFv3, you need to create an OSPFv3 routing process, create an area for OSPFv3, enable an interface for OSPFv3, and then redistribute the route into the targeted OSPFv3 routing process.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Select **Routing** > **OSPFv3**.

Step 3 By default **Enable Process 1** is selected. You can enable up to two OSPF process instances.

Step 4 Chose the OSPFv3 role from the drop-down list, and enter a description for it. The options are Internal, ABR, ASBR, and ABR and ASBR. See About OSPF, on page 769 for descriptions of the OSPFv3 roles.

Step 5 Select **Area** > **Add**.

You can click **Edit** (), or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 6 Select **General**, and configure the following options for each OSPF process:

- **Area ID**—The area for which routes are to be summarized.
- **Cost**—The metric or cost for the summary route, which is used during OSPF SPF calculations to determine the shortest paths to the destination. Valid values range from 0 to 16777215.
- **Type**—Specifies Normal, NSSA, or Stub. If you select Normal, there are no other parameters to configure. If you select Stub, you can choose to send summary LSAs in the area. If you select NSSA, you can configure the next three options:
  - **Allow Sending summary LSA into this area**—Allows the sending of summary LSAs into the area.
  - **Redistribute imports routes to normal and NSSA area**—Allows redistribution to import routes to normal and not to stubby areas.
  - **Defaults information originate**—Generates a default external route into an OSPFv3 routing domain.
- **Metric**—Metric used for generating the default route. The default value is 10. Valid metric values range from 0 to 16777214.
- **Metric Type**—The metric type is the external link type that is associated with the default route that is advertised into the OSPFv3 routing domain. The available options are 1 for a Type 1 external route or 2 for a Type 2 external route.

Step 7 Click **OK** to save the general configuration.

Step 8 Select **Route Summary** > **Add Route Summary**.

You can click **Edit** (), or use the right-click menu to cut, copy, past, insert, and delete route summaries.

Step 9 Configure the following route summary options for each OSPF process:

- **IPv6 Prefix/Length**—The IPv6 prefix. To add a new network object, click **Add** (). See Network Objects, on page 402 for the procedure for adding networks.
- **Cost**—The metric or cost for the summary route, which is used during OSPF SPF calculations to determine the shortest paths to the destination. Valid values range from 0 to 16777215.
- **Advertise**—Advertises the summary route. Uncheck this check box to suppress routes that fall under the summary address. By default, this check box is checked.

Step 10 Click **OK** to save the route summary configuration.

Step 11 Select **Virtual Link**, click **Add Virtual Link**, and configure the following options for each OSPF process:

- **Peer RouterID**—Choose the IP address of the peer router. To add a new network object, click **Add** (). See Network Objects, on page 402 for the procedure for adding networks.
- **TTL Security**—Enables TTL security check. The value for the hop-count is a number from 1 to 254. The default is 1.

  OSPF sends outgoing packets with an IP header Time to Live (TTL) value of 255 and discards incoming packets that have TTL values less than a configurable threshold. Because each device that forwards an IP packet decrements the TTL, packets received via a direct (one-hop) connection have a value of 255. Packets that cross two hops have a value of 254, and so on. The receive threshold is configured in terms of the maximum number of hops that a packet may have traveled.
- **Dead Interval**—The time in seconds that hello packets are not seen before a neighbor indicates that the router is down. The default is four times the hello interval, or 40 seconds. Valid values range from 1 to 65535.

  The dead interval is an unsigned integer. The value must be the same for all routers and access servers that are attached to a common network.
- **Hello Interval**—The time in seconds between the hello packets sent on an interface. Valid values range from 1 to 65535. The default is 10.

  The hello interval is an unsigned integer that is to be advertised in the hello packets. The value must be the same for all routers and access servers on a specific network. The smaller the hello interval, the faster topological changes are detected, but the more traffic is sent on the interface.
- **Retransmit Interval**—The time in seconds between LSA retransmissions for adjacencies that belong to the interface. The retransmit interval is the expected round-trip delay between any two routers on the attached network. The value must be greater than the expected round-trip delay, and can range from 1 to 65535. The default is 5.

  When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links.
- **Transmit Delay**—The estimated time in seconds that is required to send an LSA packet on the interface. The integer value must be greater than zero. Valid values range from 1 to 8192. The default is 1.

LSAs in the update packet have their own ages incremented by this amount before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links.

Step 12 Click **OK** to save the virtual link configuration.

Step 13 Click **Save** on the Router page to save your changes.

What to do next

Continue with Configure OSPFv3 Redistribution.

Configure OSPFv3 Redistribution

The Firepower Threat Defense device can control the redistribution of routes between the OSPF routing processes. The rules for redistributing routes from one routing process into an OSPF routing process are displayed. You can redistribute routes discovered by RIP and BGP into the OSPF routing process. You can also redistribute static and connected routes into the OSPF routing process.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Select **Routing** > **OSPF**.

Step 3 Select **Redistribution**, and click **Add.**

You can click **Edit** (), or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 4 Configure the following redistribution options for each OSPF process:

- **Source Protocol**—The source protocol from which routes are being redistributed. The supported protocols are connected, OSPF, static, and BGP. If you choose OSPF, you must enter the Process ID in the **Process ID** field. If you choose BCP, you must add the AS number in the **AS Number** field.
- **Metric** —Metric value for the routes being distributed. The default value is 10. Valid values range from 0 to 16777214.

  When redistributing from one OSPF process to another OSPF process on the same device, the metric will be carried through from one process to the other if no metric value is specified. When redistributing other processes to an OSPF process, the default metric is 20 when no metric value is specified.
- **Metric Type**—The metric type is the external link type that is associated with the default route that is advertised into the OSPF routing domain. The available options are 1 for a Type 1 external route or 2 for a Type 2 external route.
- **Tag** —Tag specifies the 32-bit decimal value attached to each external route that is not used by OSPF itself, but which may be used to communicate information between ASBRs. If none is specified, then the remote autonomous system number is used for routes from BGP and EGP. For other protocols, zero is used. Valid values are from 0 to 4294967295.

- **Route Map**—Checks for filtering the importing of routes from the source routing protocol to the current routing protocol. If this parameter is not specified, all routes are redistributed. If this parameter is specified, but no route map tags are listed, no routes are imported. Or you can add a new route map by clicking **Add** ( ). See Route Maps, on page 464 for the procedure to add a new route map.
- **Process ID**—The OSPF process ID, either 1 or 2.

  Note The Process ID is enabled the OSPFv3 process is redistributing a route learned by another OSPFv3 process.

- **Match**—Enables OSPF routes to be redistributed into other routing domains:
  - **Internal** for routes that are internal to a specific autonomous system.
  - **External 1** for routes that are external to the autonomous system, but are imported into OSPFv3 as Type 1 external routes.
  - **External 2** for routes that are external to the autonomous system, but are imported into OSPFv3 as Type 2 external routes.
  - **NSSA External 1** for routes that are external to the autonomous system, but are imported into OSPFv3 in an NSSA for IPv6 as Type 1 external routes.
  - **NSSA External 2** for routes that are external to the autonomous system, but are imported into OSPFv3 in an NSSA for IPv6 as Type 2 external routes.

Step 5 Click **OK** to save the redistribution configuration.

Step 6 Click **Save** on the Routing page to save your changes.

What to do next

Continue with Configure OSPFv3 Summary Prefixes, on page 789.

Configure OSPFv3 Summary Prefixes

You can configure the Firepower Threat Defense device to advertise routes that match a specified IPv6 prefix and mask pair.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Select **Routing** > **OSPFv3**.

Step 3 Select **Summary Prefix** > **Add**.

You can click **Edit** ( ), or use the right-click menu to cut, copy, past, insert, and delete summary prefixes.

Step 4 Configure the following summary prefix options for each OSPF process:

- **IPv6 Prefix/Length**—The IPv6 prefix and prefix length label. Select one from the list or click **Add** ( ) to add a new network object. See Network Objects, on page 402 for the procedure for adding networks.
- **Advertise**— Advertises routes that match the specified prefix and mask pair. Uncheck this check box to suppress routes that match the specified prefix and mask pair.
- (Optional) **Tag**—A value that you can use as a match value for controlling redistribution through route maps.

Step 5 Click **OK** to save the summary prefix configuration.

Step 6 Click **Save** on the Routing page to save your changes.

What to do next

Continue with Configure OSPFv3 Interfaces, Authentication, and Neighbors, on page 790.

Configure OSPFv3 Interfaces, Authentication, and Neighbors

You can change certain interface-specific OSPFv3 parameters, if necessary. You are not required to change any of these parameters, but the following interface parameters must be consistent across all routers in an attached network: the hello interval and the dead interval. If you configure any of these parameters, be sure that the configurations for all routers on your network have compatible values.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Select **Routing** > **OSPFv3**.

Step 3 Select **Interface** > **Add**.

You can click **Edit** to edit, or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 4 Configure the following interface options for each OSPFv3 process:

- **Interface**—The interface you are configuring.
- **Enable OSPFv3**—Enables OSPFv3.
- **OSPF Process**—Choose 1 or 2.
- **Area**—The area ID for this process.
- **Instance** —Specifies the area instance ID to be assigned to the interface. An interface can have only one OSPFv3 area. You can use the same area on multiple interfaces, and each interface can use a different area instance ID.

Step 5 Select **Properties**, and configuring the following options for each OSPFv3 process:

- **Filter Outgoing Link Status Advertisements**—Filters outgoing LSAs to an OSPFv3 interface. All outgoing LSAs are flooded to the interface by default.

- **Disable MTU mismatch detection**—Disables the OSPF MTU mismatch detection when DBD packets are received. OSPF MTU mismatch detection is enabled by default.
- **Flood Reduction**—Changes normal LSAs into Do Not Age LSAs, so that they don't get flooded every 3600 seconds across areas.

  OSPF LSAs are refreshed every 3600 seconds. In large OSPF networks, this can lead to large amounts of unnecessary LSA flooding from area to area.
- **Point-to-Point Network**—Lets you transmit OSPF routes over VPN tunnels. When an interface is configured as point-to-point, non-broadcast, the following restrictions apply:
  - You can define only one neighbor for the interface.
  - You need to manually configure the neighbor.
  - You need to define a static route pointing to the crypto endpoint.
  - If OSPF over a tunnel is running on the interface, regular OSPF with an upstream router cannot be run on the same interface.
  - You should bind the crypto map to the interface before specifying the OSPF neighbor to ensure that the OSPF updates are passed through the VPN tunnel. If you bind the crypto map to the interface after specifying the OSPF neighbor, use the **clear local-host all** command to clear OSPF connections so that the OSPF adjacencies can be established over the VPN tunnel.
- **Broadcast**— Specifies that the interface is a broadcast interface. By default, this check box is checked for Ethernet interfaces. Uncheck this check box to designate the interface as a point-to-point, nonbroadcast interface. Specifying an interface as point-to-point, nonbroadcast lets you transmit OSPF routes over VPN tunnels.
- **Cost**—Specifies the cost of sending a packet on the interface. Valid values for this setting range from 0 to 255. The default value is 1. Entering 0 for this setting makes the router ineligible to become the designated router or backup designated router. This setting does not apply to interfaces that are configured as point-to-point, nonbroadcast interfaces.

  When two routers connect to a network, both attempt to become the designated router. The device with the higher router priority becomes the designated router. If there is a tie, the router with the higher router ID becomes the designated router.
- **Priority**—Determines the designated router for a network. Valid values range from 0 to 255.
- **Dead Interval**—Time period in seconds for which hello packets must not be seen before neighbors indicate that the router is down. The value must be the same for all nodes on the network and can range from 1 to 65535.
- **Poll Interval**— Time period in seconds between OSPF packets that the router will send before adjacency is established with a neighbor. Once the routing device detects an active neighbor, the hello packet interval changes from the time specified in the poll interval to the time specified in the hello interval. Valid values range from 1 to 65535 seconds.
- **Retransmit Interval**—Time in seconds between LSA retransmissions for adjacencies that belong to the interface. The time must be greater than the expected round-trip delay between any two routers on the attached network. Valid values range from 1 to 65535 seconds. The default is 5 seconds.
- **Transmit Delay**—Estimated time in seconds to send a link-state update packet on the interface. Valid values range from 1 to 65535 seconds. The default is 1 second.

Step 6 Click **OK** to save the properties configuration.

Step 7 Select **Authentication**, and configure the following options for each OSPFv3 process:

- **Type**—Type of authentication. The available options are Area, Interface, and None. The None option indicates that no authentication is used.
- **Security Parameters Index**— A number from 256 to 4294967295. Configure this if you chose Interface as the type.
- **Authentication**—Type of authentication algorithm. Supported values are SHA-1 and MD5. Configure this if you chose Interface as the type.
- **Authentication Key**— When MD5 authentication is used, the key must be 32 hexadecimal digits (16 bytes) long. When SHA-1 authentication is used, the key must be 40 hexadecimal digits (20 bytes) long.
- **Encrypt Authentication Key**—Enables encryption of the authentication key.
- **Include Encryption**— Enables encryption.
- **Encryption Algorithm**—Type of encryption algorithm. Supported value is DES. The NULL entry indicates no encryption. Configure this if you chose **Include Encryption**.
- **Encryption Key**—Enter the encryption key. Configure this if you chose **Include Encryption**.
- **Encrypt Key**—Enables the key to be encrypted.

Step 8 Click **OK** to save the authentication configuration.

Step 9 Select **Neighbor**, click **Add**, and configure the following options for each OSPFv3 process:

- **Link Local Address**—The IPv6 address of the static neighbor.
- **Cost**—Enables cost. Enter the cost in the **Cost** field, and check the **Filter Outgoing Link State Advertisements** if you want to advertise.
- (Optional) **Poll Interval**—Enables the poll interval. Enter the **Priority** level and the **Poll Interval** in seconds.

Step 10 Click **Add** to add the neighbor.

Step 11 Click **OK** to save the Interface configuration.

Configure OSPFv3 Advanced Properties

The Advanced Properties allows you to configure options, such as syslog message generation, administrative route distances, passive OSPFv3 routing, LSA timers, and graceful restarts.

Graceful Restarts

The Firepower Threat Defense device may experience some known failure situations that should not affect packet forwarding across the switching platform. The Non-Stop Forwarding (NSF) capability allows data forwarding to continue along known routes, while the routing protocol information is being restored. This capability is useful when there is a scheduled hitless software upgrade. You can configure graceful restart on OSPFv3 using graceful-restart (RFC 5187).

Note NSF capability is also useful in HA mode and clustering.

Configuring the NSF graceful-restart feature involves two steps; configuring capabilities and configuring a device as NSF-capable or NSF-aware. A NSF-capable device can indicate its own restart activities to neighbors and a NSF-aware device can help a restarting neighbor.

A device can be configured as NSF-capable or NSF-aware, depending on some conditions:

- A device can be configured as NSF-aware irrespective of the mode in which it is.
- A device has to be in either Failover or Spanned Etherchannel (L2) cluster mode to be configured as NSF-capable.
- For a device to be either NSF-aware or NSF-capable, it should be configured with the capability of handling opaque Link State Advertisements (LSAs)/ Link Local Signaling (LLS) block as required.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Select **Routing** > **OSPFv3** > **Advanced**.

Step 3 For **Router ID**, choose **Automatic** or **IP address**. If you choose IP address, enter the IP address in the IP Address field.

Step 4 Check the **Ignore LSA MOSPF** check box if you want to suppress syslog messages when the route receives unsupported LSA Type 6 multicast OSPF (MOSPF) packets.

Step 5 Select **General**, and configure the following:

- **Adjacency Changes**—Defines the adjacency changes that cause syslog messages to be sent.

  By default, a syslog message is generated when an OSPF neighbor goes up or down. You can configure the router to send a syslog message when an OSPF neighbor goes down and also a syslog for each state.

  - **Adjacency Changes**—Causes the Firepower Threat Defense device to send a syslog message whenever an OSPF neighbor goes up or down. This setting is checked by default.
  - **Include Details**—Causes the Firepower Threat Defense device to send a syslog message whenever any state change occurs, not just when a neighbor goes up or down. This setting is unchecked by default.

- **Administrative Route Distances**—Allows you to modify the settings that were used to configure administrative route distances for inter-area, intra-area, and external IPv6 routes. The administrative route distance is an integer from 1 to 254. The default is 110.
- **Default Information Originate**—Check the **Enable** check box to generate a default external route into an OSPFv3 routing domain and configure the following options:
  - **Always Advertise**—Will always advertise the default route whether or not one exists.
  - **Metric**—Metric used for generating the default route. Valid metric values range from 0 to 16777214. The default value is 10.

- **Metric Type**—The external link type that is associated with the default route that is advertised into the OSPFv3 routing domain. Valid values are 1 (Type 1 external route) and 2 (Type 2 external route). The default is Type 2 external route.
- **Route Map**—Choose the routing process that generates the default route if the route map is satisfied or click **Add** ( ) to add a new one. See Route Maps, on page 464 to add a new route map.

Step 6 Click **OK** to save the general configuration.

Step 7 Select **Passive Interface**, select the interfaces on which you want to enable passive OSPFv3 routing from the Available Interfaces list, and click **Add** to move them to the Selected Interfaces list.

Passive routing assists in controlling the advertisement of OSPFv3 routing information and disables the sending and receiving of OSPFv3 routing updates on an interface.

Step 8 Click **OK** to save the passive interface configuration.

Step 9 Select **Timer**, and configure the following LSA pacing and SPF calculation timers:

- **Arrival**—Specifies the minimum delay in milliseconds that must pass between acceptance of the same LSA arriving from neighbors. The range is from 0 to 6000,000 milliseconds. The default is 1000 milliseconds.
- **Flood Pacing**—Specifies the time in milliseconds at which LSAs in the flooding queue are paced in between updates. The configurable range is from 5 to 100 milliseconds. The default value is 33 milliseconds.
- **Group Pacing**—Specifies the interval in seconds at which LSAs are collected into a group and refreshed, check summed, or aged. Valid values range from 10 to 1800. The default value is 240.
- **Retransmission Pacing**—Specifies the time in milliseconds at which LSAs in the retransmission queue are paced. The configurable range is from 5 to 200 milliseconds. The default value is 66 milliseconds.
- **LSA Throttle**—Specifics the delay in milliseconds to generate the first occurrence of the LSA. The default value is 0 millisecond. The minimum specifies the minimum delay in milliseconds to originate the same LSA. The default value is 5000 milliseconds. The maximum specifies the maximum delay in milliseconds to originate the same LSA. The default value is 5000 milliseconds.

  Note For LSA throttling, if the minimum or maximum time is less than the first occurrence value, then OSPFv3 automatically corrects to the first occurrence value. Similarly, if the maximum delay specified is less than the minimum delay, then OSPFv3 automatically corrects to the minimum delay value.

- **SPF Throttle**—Specifies the delay in milliseconds to receive a change to the SPF calculation. The default value is 5000 milliseconds. The minimum specifies the delay in milliseconds between the first and second SPF calculations. The default value is 10000 milliseconds. The maximum specifies the maximum wait time in milliseconds for SPF calculations. The default value is 10000 milliseconds.

  Note For SPF throttling, if the minimum or maximum time is less than the first occurrence value, then OSPFv3 automatically corrects to the first occurrence value. Similarly, if the maximum delay specified is less than the minimum delay, then OSPFv3 automatically corrects to the minimum delay value.

Step 10 Click **OK** to save the LSA timer configuration.

Step 11 Select **Non Stop Forwarding**, and check the **Enable graceful-restart helper** check box. This is checked by default. Uncheck this to disable the graceful-restart helper mode on an NSF-aware device.

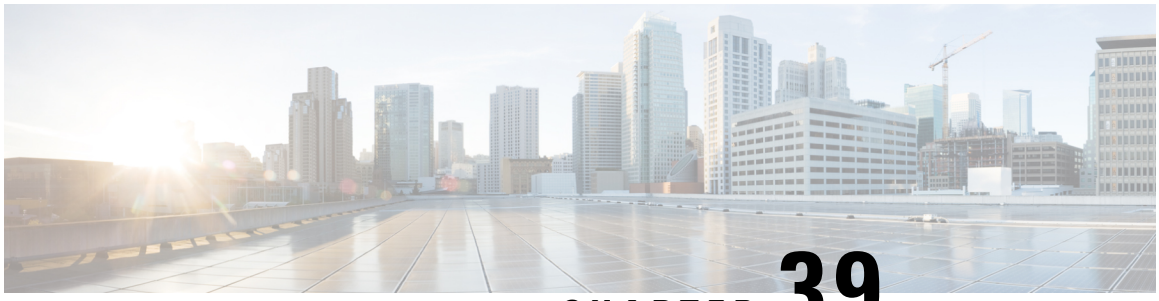
Step 12 Check the **Enable link state advertisement** check box to enable strict link state advertisement checking.

When enabled, it indicates that the helper router will terminate the process of restarting the router if it detects that there is a change to a LSA that would be flooded to the restarting router, or if there is a changed LSA on the retransmission list of the restarting router when the graceful restart process is initiated.

Step 13 Check the **Enable graceful-restart (Use when Spanned Cluster or Failover Configured)** and enter the graceful-restart interval in seconds. The range is 1-1800. The default value is 120 seconds. For a restart interval below 30 seconds, graceful restart will be terminated.

Step 14 Click **OK** to save the graceful restart configuration.

Step 15 Click **Save** on the Routing page to save your changes.



CHAPTER 39

BGP for Firepower Threat Defense

This section describes how to configure the Firepower Threat Defense to route data, perform authentication, and redistribute routing information using the Border Gateway Protocol (BGP).

- About BGP, on page 797
- Requirements and Prerequisites for BGP, on page 800
- Guidelines for BGP, on page 800
- Configure BGP, on page 801

About BGP

BGP is an inter and intra autonomous system routing protocol. An autonomous system is a network or group of networks under a common administration and with common routing policies. BGP is used to exchange routing information for the Internet and is the protocol used between Internet service providers (ISP).

Routing Table Changes

BGP neighbors exchange full routing information when the TCP connection between neighbors is first established. When changes to the routing table are detected, the BGP routers send to their neighbors only those routes that have changed. BGP routers do not send periodic routing updates, and BGP routing updates advertise only the optimal path to a destination network.



Note AS loop detection is done by scanning the full AS path (as specified in the AS_PATH attribute), and checking that the AS number of the local system does not appear in the AS path. By default, EBGP advertises the learned routes to the same peer to prevent additional CPU cycles on the ASA in performing loop checks and to avoid delays in the existing outgoing update tasks.

Routes learned via BGP have properties that are used to determine the best route to a destination, when multiple paths exist to a particular destination. These properties are referred to as BGP attributes and are used in the route selection process:

- Weight—This is a Cisco-defined attribute that is local to a router. The weight attribute is not advertised to neighboring routers. If the router learns about more than one route to the same destination, the route with the highest weight is preferred.

- Local preference—The local preference attribute is used to select an exit point from the local AS. Unlike the weight attribute, the local preference attribute is propagated throughout the local AS. If there are multiple exit points from the AS, the exit point with the highest local preference attribute is used as an exit point for a specific route.
- Multi-exit discriminator—The multi-exit discriminator (MED) or metric attribute is used as a suggestion to an external AS regarding the preferred route into the AS that is advertising the metric. It is referred to as a suggestion because the external AS that is receiving the MEDs may also be using other BGP attributes for route selection. The route with the lower MED metric is preferred.
- Origin—The origin attribute indicates how BGP learned about a particular route. The origin attribute can have one of three possible values and is used in route selection.
  - IGP—The route is interior to the originating AS. This value is set when the network router configuration command is used to inject the route into BGP.
  - EGP—The route is learned via the Exterior Border Gateway Protocol (EBGP).
  - Incomplete—The origin of the route is unknown or learned in some other way. An origin of incomplete occurs when a route is redistributed into BGP.
- AS_path—When a route advertisement passes through an autonomous system, the AS number is added to an ordered list of AS numbers that the route advertisement has traversed. Only the route with the shortest AS_path list is installed in the IP routing table.
- Next hop—The EBGP next-hop attribute is the IP address that is used to reach the advertising router. For EBGP peers, the next-hop address is the IP address of the connection between the peers. For IBGP, the EBGP next-hop address is carried into the local AS.
- Community—The community attribute provides a way of grouping destinations, called communities, to which routing decisions (such as acceptance, preference, and redistribution) can be applied. Route maps are used to set the community attribute. The predefined community attributes are as follows:
  - no-export—Do not advertise this route to EBGP peers.
  - no-advertise—Do not advertise this route to any peer.
  - internet—Advertise this route to the Internet community; all routers in the network belong to it.

When to Use BGP

Customer networks, such as universities and corporations, usually employ an Interior Gateway Protocol (IGP) such as OSPF for the exchange of routing information within their networks. Customers connect to ISPs, and ISPs use BGP to exchange customer and ISP routes. When BGP is used between autonomous systems (AS), the protocol is referred to as External BGP (EBGP). If a service provider is using BGP to exchange routes within an AS, then the protocol is referred to as Interior BGP (IBGP).

BGP can also be used for carrying routing information for IPv6 prefix over IPv6 networks.

BGP Path Selection

BGP may receive multiple advertisements for the same route from different sources. BGP selects only one path as the best path. When this path is selected, BGP puts the selected path in the IP routing table and

propagates the path to its neighbors. BGP uses the following criteria, in the order presented, to select a path for a destination:

- If the path specifies a next hop that is inaccessible, drop the update.
- Prefer the path with the largest weight.
- If the weights are the same, prefer the path with the largest local preference.
- If the local preferences are the same, prefer the path that was originated by BGP running on this router.
- If no route was originated, prefer the route that has the shortest AS_path.
- If all paths have the same AS_path length, prefer the path with the lowest origin type (where IGP is lower than EGP, and EGP is lower than incomplete).
- If the origin codes are the same, prefer the path with the lowest MED attribute.
- If the paths have the same MED, prefer the external path over the internal path.
- If the paths are still the same, prefer the path through the closest IGP neighbor.
- Determine if multiple paths require installation in the routing table for BGP Multipath, on page 799.
- If both paths are external, prefer the path that was received first (the oldest one).
- Prefer the path with the lowest IP address, as specified by the BGP router ID.
- If the originator or router ID is the same for multiple paths, prefer the path with the minimum cluster list length.
- Prefer the path that comes from the lowest neighbor address.

BGP Multipath

BGP Multipath allows installation into the IP routing table of multiple equal-cost BGP paths to the same destination prefix. Traffic to the destination prefix is then shared across all installed paths.

These paths are installed in the table together with the best path for load-sharing. BGP Multipath does not affect best-path selection. For example, a router still designates one of the paths as the best path, according to the algorithm, and advertises this best path to its BGP peers.

In order to be candidates for multipath, paths to the same destination need to have these characteristics equal to the best-path characteristics:

- Weight
- Local preference
- AS-PATH length
- Origin code
- Multi Exit Discriminator (MED)
- One of these:
  - Neighboring AS or sub-AS (before the addition of the BGP Multipaths)
  - AS-PATH (after the addition of the BGP Multipaths)

Some BGP Multipath features put additional requirements on multipath candidates:

- The path should be learned from an external or confederation-external neighbor (eBGP).
- The IGP metric to the BGP next hop should be equal to the best-path IGP metric.

These are the additional requirements for internal BGP (iBGP) multipath candidates:

- The path should be learned from an internal neighbor (iBGP).
- The IGP metric to the BGP next hop should be equal to the best-path IGP metric, unless the router is configured for unequal-cost iBGP multipath.

BGP inserts up to *n* most recently received paths from multipath candidates into the IP routing table, where *n* is the number of routes to install to the routing table, as specified when you configure BGP Multipath. The default value, when multipath is disabled, is 1.

For unequal-cost load balancing, you can also use BGP Link Bandwidth.

Note The equivalent next-hop-self is performed on the best path that is selected among eBGP multipaths before it is forwarded to internal peers.

Requirements and Prerequisites for BGP

Model Support

FTD

FTDv

Supported Domains

Any

User Roles

Admin

Network Admin

Guidelines for BGP

Firewall Mode Guidelines

Does not support transparent firewall mode. BGP is supported only in routed mode.

IPv6 Guidelines

Supports IPv6. Graceful restart is not supported for IPv6 address family.

Additional Guidelines

- The system does not add route entry for the IP address received over PPPoE in the CP route table. BGP always looks into CP route table for initiating the TCP session, hence BGP does not form TCP session.

  Thus, BGP over PPPoE is not supported.

- To avoid adjacency flaps due to route updates being dropped if the route update is larger than the minimum MTU on the link, ensure that you configure the same MTU on the interfaces on both sides of the link.
- The BGP table of the member unit is not synchronized with the control unit table. Only its routing table is synchronized with the control unit routing table.

Configure BGP

To configure BGP, see the following topics:

Procedure

Step 1 Configure BGP Basic Settings, on page 801
Step 2 Configure BGP General Settings, on page 803
Step 3 Configure BGP Neighbor Settings, on page 804
Step 4 Configure BGP Aggregate Address Settings, on page 808
Step 5 Configure BGPv4 Filtering Settings, on page 808

Note The Filtering section is applicable only to IPv4 settings

Step 6 Configure BGP Network Settings, on page 809
Step 7 Configure BGP Redistribution Settings, on page 809
Step 8 Configure BGP Route Injection Settings, on page 810

Configure BGP Basic Settings

You can set many basic settings for BGP.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.
Step 2 Select **Routing**.
Step 3 (For a non-virtual-router-aware device) Select **BGP**.
Step 4 Select the **Enable BGP** check box to enable the BGP routing process.
Step 5 In the **AS Number** field, enter the autonomous system (AS) number for the BGP process. The AS number internally includes multiple autonomous numbers. The AS number can be from 1 to 4294967295 or from 1.0 to 65535.65535. The AS number is a uniquely assigned value, that identifies each network on the Internet.

Step 6 (Optional) Edit the various BGP settings, starting with **General**. The defaults for these settings are appropriate in most cases, but you can adjust them to fit the needs of your network. Click **Edit** (pencil) to edit the settings in the group:

a) In the **Router ID** drop-down list, select Automatic or Manual from the drop-down list. If you choose Automatic, the highest-level IP address on the Firepower Threat Defense device is used as the router ID. To use a fixed router ID, choose Manual and enter an IPv4 address in the**IP Address** field. The default value is Automatic.

b) Enter the **Number of AS numbers in AS_PATH attribute**. An AS _PATH attribute is a sequence of intermediate AS numbers between source and destination routers that form a directed route for packets to travel. Valid values are between 1 and 254. The default value is None.

c) Check the **Log Neighbor Changes** check box to enable logging of BGP neighbor changes (up or down) and resets. This helps in troubleshooting network connectivity problems and measuring network stability. This is enabled by default.

d) Check the **Use TCP Path MTU Discovery** check box to use the Path MTU determining technique to determine the maximum transmission unit (MTU) size on the network path between two IP hosts. This avoids IP fragmentation. This is enabled by default.

e) Check the **Reset session upon Failover** check box to reset the external BGP session immediately upon link failure. This is enabled by default.

f) Check the **Enforce that the first AS is peer's AS for EBGP routes** check box to discard incoming updates received from external BGP peers that do not list their AS number as the first segment in the AS_PATH attribute. This prevents a mis-configured or unauthorized peer from misdirecting traffic by advertising a route as if it was sourced from another autonomous system. This is enabled by default.

g) Check the **Use dot notation for AS number** check box to split the full binary 4-byte AS number into two words of 16 bits each, separated by a dot. AS numbers from 0-65553 are represented as decimal numbers and AS numbers larger than 65535 are represented using the dot notation. This is disabled by default.

h) Click **OK**.

Step 7 (Optional) Edit the **Best Path Selection** section:

a) Enter a value for **Default Local Preference** between 0 and 4294967295. The default value is 100. Higher values indicate higher preference. This preference is sent to all routers and access servers in the local autonomous system.

b) Check the **Allow comparing MED from different neighbors** check box to allow the comparison of Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems. This is disabled by default.

c) Check the **Compare Router ID for identical EBGP paths** check box to compare similar paths received from external BGP peers during the best path selection process and switch the best path to the route with the lowest router ID. This is disabled by default.

d) Check the **Pick the best MED path among paths advertised from the neighboring AS** check box to enable MED comparison among paths learned from confederation peers. The comparison between MEDs is made only if no external autonomous systems are there in the path. This is disabled by default.

e) Check the **Treat missing MED as the least preferred one** check box to consider the missing MED attribute as having a value of infinity, making the path the least desirable; therefore, a path with a missing MED is least preferred. This is disabled by default.

f) Click **OK**.

Step 8 (Optional) Edit the **Neighbor Timers** section:

a) Enter the time interval for which the BGP neighbor remains active after not sending a keepalive message in the **Keepalive interval** field. At the end of this keepalive interval, the BGP peer is declared dead, if no messages are sent. The default value is 60 seconds.

b) Enter the time interval for which the BGP neighbor remains active while a BGP connection is being initiated and configured in the **Hold time** field. The default value is 180 seconds.
c) (Optional) Enter the minimum time interval for which the BGP neighbor remains active while a BGP connection is being initiated and configured in the **Min Hold time** field. Specify a value from 0 to 65535.

Note A hold time of less than 20 seconds increases the possibility of peer flapping.

d) Click **OK**.

Step 9 (Optional) Edit the **Graceful Restart** section:

Note This section is available only when the Firepower Threat Defense device is in failover or spanned cluster mode. This is done so that there is no drop in packets in the traffic flow, when one of the devices in the failover setup fails.

a) Check the **Enable Graceful Restart** checkbox to enable Firepower Threat Defense peers to avoid a routing flap following a switchover.
b) Specify the time duration that Firepower Threat Defense peers will wait to delete stale routes before a BGP open message is received in the **Restart Time** field. The default value is 120 seconds. Valid values are between 1 and 3600 seconds.
c) Enter the time duration that the Firepower Threat Defense will wait before deleting stale routes after an end of record (EOR) message is received from the restarting Firepower Threat Defense in the **Stalepath Time** field. The default value is 360 seconds. Valid values are between 1 and 3600 seconds.
d) Click **OK**.

Step 10 Click **Save**.

Configure BGP General Settings

Configure Route maps, Administrative Route Distances, Synchronisation, Next-hop, and packet forwarding. The defaults for these settings are appropriate in most cases, but you can adjust them to fit the needs of your network.

Procedure

Step 1 On the **Device Management** page, click **Routing**.
Step 2 Choose **BGP** > **IPv4** or **IPv6**.
Step 3 Click **General**.
Step 4 In **General**, update the following sections:

a) In the **Settings** section, enter or select a **Route Map** object and enter a **Scanning Interval** for BGP routers for next-hop validation. Valid values are from 5 to 60 seconds. The default value is 60. Click **OK**.

Note The **Route Map** field is applicable only to IPv4 settings

b) In the **Routes and Synchronization** section, update the following as required, and click **OK**:

- (Optional) **Generate Default Routes** — Select this option to configure default-information originate.

- (Optional) **Summarize subnet routes into network-level routes**— Select this to configure automatic summarization of subnet routes into network-level routes. This check box is applicable only to IPv4 settings.
- (Optional) **Advertise inactive routes**— Select this to advertise routes that are not installed in the routing information base (RIB).
- (Optional) **Synchronise between BGP and IGP system**— Select this to enable synchronization between BGP and your Interior Gateway Protocol (IGP) system. Usually, a BGP speaker does not advertise a route to an external neighbor unless that route is local or exists in the IGP. This feature allows routers and access servers within an autonomous system to have the route before BGP makes it available to other autonomous systems.
- (Optional) **Redistribute IBGP into IGP**— Select this to configure iBGP redistribution into an interior gateway protocol (IGP), such as OSPF.

c) In the **Administrative Route Distances** section, update the following as required, and click **OK**:

- **External** — Enter the administrative distance for external BGP routes. Routes are external when learned from an external autonomous system. The range of values for this argument are from 1 to 255. The default value is 20.
- **Internal** — Enter administrative distance for internal BGP routes. Routes are internal when learned from peer in the local autonomous system. The range of values for this argument are from 1 to 255. The default value is 200.
- **Local** — Enter administrative distance for local BGP routes. Local routes are those networks listed with a network router show command, often as back doors, for the router or for the networks that is being redistributed from another process. The range of values for this argument are from 1 to 255. The default value is 200.

d) In the **Next Hop** section, optionally select the **Enable address tracking** check box to enable BGP next hop address tracking and enter the **Delay Interval** between checks on updated next-hop routes installed in the routing table. Click **OK**.

Note The **Next Hop** section is applicable only to IPv4 settings.

e) In the **Forward Packets over Multiple Paths** section, update the following as required and click **OK**:

- (Optional) **Number of Paths** — Specify the maximum number of Border Gateway Protocol routes that can be installed in a routing table. The range of values are from 1 to 8. The default value is 1.
- (Optional) **IBGP Number of Paths** — Specify the maximum number of parallel internal Border Gateway Protocol (iBGP) routes that can be installed in a routing table. The range of values are from 1 to 8. The default value is 1.

Step 5 Click **Save**.

Configure BGP Neighbor Settings

A BGP router must connect with each of its peers before exchanging updates. These peers are called BGP neighbors. Use Neighbor to define BGP IPv4 or IPv6 neighbors and neighbor settings.

Procedure

Step 1 On the Device Management page, click **Routing**.

Step 2 Choose **BGP** > **IPv4** or **IPv6**.

Step 3 Click **Neighbor**.

Step 4 Click **Add** to define BGP neighbors and neighbor settings.

Step 5 Enter the BGP neighbor **IP address**. This IP address is added to the BGP neighbor table.

Step 6 Choose the BGP neighbor **Interface**.

Note The **Interface** field is only applicable to IPv6 settings.

Step 7 Enter the autonomous system to which the BGP neighbor belongs, in the **Remote AS** field.

Step 8 Select the **Enabled address** check box to enable communication with this BGP neighbor. Further neighbor settings will be configured only if the Enabled address check box is selected.

Step 9 (Optional) Select the **Shutdown administratively** check box to disable a neighbor or peer group.

Step 10 (Optional) Select the **Configure graceful restart** check box to enable configuration of the BGP graceful restart capability for this neighbor. After selecting this option, you must use the Graceful Restart (failover / spanned mode) option to specify whether graceful restart should be enabled or disabled for this neighbor.

Note The graceful restart fields are only applicable to IPv4 settings.

Step 11 (Optional) Enter a **Description** for the BGP neighbor.

Step 12 (Optional) In **Filtering Routes**, use access lists, route maps, prefix lists and AS path filters as required, to distribute BGP Neighbor information. Update the following sections:

a) Enter or Select the appropriate incoming or outgoing **Access List** to distribute BGP neighbor information.

Note Access Lists are only applicable to IPv4 settings.

b) Enter or Select the appropriate incoming or outgoing **Route Maps** to apply a route map to incoming or outgoing routes.

c) Enter or Select the appropriate incoming or outgoing **Prefix List** to distribute BGP neighbor information.

d) Enter or Select the appropriate incoming or outgoing **AS path filter** to distribute BGP neighbor information.

e) Select the **Limit the number of prefixes allowed from the neighbor** to control the number of prefixes that can be received from a neighbor.

- Enter the maximum number of prefixes allowed from a specific neighbor in the **Maximum Prefixes** field.
- Enter the percentage (of maximum) at which the router starts to generate a warning message in the **Threshold Level** field. Valid values are integers between 1 and 100. The default value is 75.

f) Select the **Control prefixes received from the peer** check box to specify additional controls for the prefixes received from a peer. Do one of the following

- Select **Terminate peering when prefix limit is exceeded** to stop the BGP neighbor when the prefix limit is reached. Specify the interval after which the BGP neighbor will restart in the **Restart interval** field.
- Select **Give only warning message when prefix limit is exceeded** to generate a log message when the maximum prefix limit is exceeded. Here, the BGP neighbor will not be terminated.

g) Click **OK**.

Step 13 (Optional) In **Routes**, specify miscellaneous Neighbor route parameter. Proceed to update the following:

a) Enter the minimum interval (in seconds) between the sending of BGP routing updates in the **Advertisment Interval** field. Valid values are between 1 and 600.

b) Select the **Remove private AS numbers from outbound routing updates** to exclude the private AS numbers from being advertised on outbound routes.

c) Select the **Generate default routes** checkbox to allow the local router to send the default route 0.0.0.0 to a neighbor to use as a default route. Enter or Select the route map that allows the route 0.0.0.0 to be injected conditionally in the **Route map** field.

d) To add conditionally advertised routes, click Add Row +. In the Add Advertised Route dialog box, do the following:

1. Add or select a route map in the **Advertise Map** field, that will be advertised if the conditions of the exist map or the non-exist map are met.
2. Select **Exist Map** and choose a route map from the Route Map Object Selector. This route map is compared with the routes in the BGP table, to determine whether the advertise map route is advertised.
3. Select **Non-Exist Map** and choose a route map from the Route Map Object Selector. This route map is compared with the routes in the BGP table, to determine whether the advertise map route is advertised.
4. Click **OK**.

Step 14 In **Timers**, select the **Set Timers for the BGP Peer** check box to set the keepalive frequency, hold time and minimum hold time

- **Keepalive Interval**—Enter the frequency (in seconds) with which the Firepower Threat Defense device sends keepalive messages to the neighbor. Valid values are between 0 and 65535. The default value is 60 seconds.
- **Hold time**—Enter the interval (in seconds) after not receiving a keepalive message that theFirepower Threat Defense device declares a peer dead. Valid values are between 0 and 65535. The default value is 180 seconds.
- **Min hold time**—(Optional) Enter the minimum interval (in seconds) after not receiving a keepalive message that the Firepower Threat Defense device declares a peer dead. Valid values are between 0 and 65535. The default value is 0 seconds.

Note A hold time of less than 20 seconds increases the possibility of peer flapping.

Step 15 In **Advanced**, update the following:

a) (Optional) Select **Enable Authentication** to enable MD5 authentication on a TCP connection between two BGP peers.

1. Choose an encryption type from the **Enable Encryption** drop-down list.
2. Enter a password in the **Password** field. Reenter the password in the **Confirm** field. The password is case-sensitive and can be up to 25 characters long when the service password-encryption command is enabled and up to 81 characters long when the service password-encryption command is not enabled. The string can contain any alphanumeric characters, including spaces.

Note You cannot specify a password in the format number-space-anything. The space after the number can cause authentication to fail.

b) (Optional) Select the **Send Communty attribute to this neighbor** check box to specify that communities attributes should be sent to the BGP neighbor
c) (Optional) Select the **Use FTD as next hop for this neighbor** check box to configure the router as the next-hop for a BGP speaking neighbor or peer group.
d) Select the **Disable Connection Verification** checkbox to disable the connection verification process for eBGP peering sessions that are reachable by a single hop but are configured on a loopback interface or otherwise configured with a non-directly connected IP address. When deselected (default), a BGP routing process will verify the connection of single-hop eBGP peering session (TTL=254) to determine if the eBGP peer is directly connected to the same network segment by default. If the peer is not directly connected to same network segment, connection verification will prevent the peering session from being established.
e) Select **Allow connections with neighbor that is not directly connected** to accept and attempt BGP connections to external peers residing on networks that are not directly connected. (Optional) Enter the time-to-live in the **TTL hops** field. Valid values are between 1 and 255. Alternately, select **Limited number of TTL hops to neighbor**, to secure a BGP peering session. Enter the maximum number of hops that separate eBGP peers in the **TTL hops** field. Valid values are between 1 and 254.
f) (Optional) Select the **Use TCP MTU path discovery** check box to enable a TCP transport session for a BGP session.
g) Choose the TCP connection mode from the **TCP Transport Mode** drop-down list. Options are Default, Active, or Passive.
h) (Optional) Enter a **Weight** for the BGP neighbor connection.
i) Select the **BGP Version** that the Firepower Threat Defense device will accept from the drop-down list. The version can be set to 4-Only to force the software to use only Version 4 with the specified neighbor. The default is to use Version 4 and dynamically negotiate down to Version 2 if requested.

Step 16 Update **Migration**, only if AS migration is considered.

Note The AS migration customization should be removed after transition has been completed.

a) (Optional) Select the **Customize the AS number for routes received from the neighbor** check box to customize the AS_PATH attribute for routes received from an eBGP neighbor.
b) Enter the local autonomous system number in the **Local AS number** field. Valid values are any valid autonomous system number from 1 to 4294967295 or 1.0 to65535.65535.
c) (Optional) Select the **Do not prepend local AS number to routes received from neighbor** check box to prevent the local AS number from being prepended to any routes received from eBGP peer.
d) (Optional) Select the **Replace real AS number with local AS number in routes received from neighbor** check box to replace the real autonomous system number with the local autonomous system number in the eBGP updates. The autonomous system number from the local BGP routing process is not prepended.
e) (Optional) Select the **Accept either real AS number or local AS number in routesreceived from neighbor** check box to configure the eBGP neighbor to establish a peering session using the real autonomous system number (from the local BGP routing process) or by using the local autonomous system number.

Step 17 Click **OK**.

Step 18 Click **Save**.

Configure BGP Aggregate Address Settings

BGP neighbors store and exchange routing information and the amount of routing information increases as more BGP speakers are configured. Route aggregation is the process of combining the attributes of several different routes so that only a single route is advertised. Aggregate prefixes use the classless interdomain routing (CIDR) principle to combine contiguous networks into one classless set of IP addresses that can be summarized in routing tables. As a result fewer routes need to be advertised. Use the Add/Edit Aggregate Address dialog box to define the aggregation of specific routes into one route.

Procedure

Step 1 When editing a Firepower Threat Defense device, click **Routing**.

Step 2 Choose **BGP** > **IPv4** or **IPv6**.

Step 3 Click **Add Aggregate Address**.

Step 4 Enter a value for the aggregate timer (in seconds) in the **Aggregate Timer** field. Valid values are 0 or any value between 6 and 60. The default value is 30.

Step 5 Click **Add** and update the **Add Aggregate Address** dialog:

a) **Network** — Enter an IPv4 address or select the desired network/hosts objects.
b) **Attribute Map** — (Optional) Enter or select the route map used to set the attribute of the aggregate route.
c) **Advertise Map** — (Optional) Enter or select the route map used to select the routes to create AS_SET origin communities.
d) **Suppress Map** — (Optional) Enter or select the route map used to select the routes to be suppressed.
e) **Generate AS set path Information** — (Optional) Select the check box to enable generation of autonomous system set path information.
f) **Filter all routes from updates** — (Optional) Select the check box to filter all more-specific routes from updates.
g) Click **OK**.

What to do next

- For BGPv4 settings, proceed to Configure BGPv4 Filtering Settings, on page 808
- For BGPv6 settings, proceed to Configure BGP Network Settings, on page 809

Configure BGPv4 Filtering Settings

Filtering settings are used to filter routes or networks received in incoming BGP updates. Filtering is used to restrict routing information that the router learns or advertises.

Before you begin

Filtering is only applicable for a BGP IPv4 routing policy.

Procedure

Step 1 On the Device Management page, click **Routing**.

Step 2 Choose **BGP** > **IPv4**.

Step 3 Click **Filtering**.

Step 4 Click **Add** and update the **Add Filter** dialog:

a) **Access List**— Select an access control list that defines which networks are to be received and which are to be suppressed in routing updates.

b) **Direction**— (Optional) Select a direction that specifies if the filter should be applied to inbound updates or outbound updates.

c) **Protocol**— (Optional) Select the routing process for which you want to filter: None, BGP, Connected, OSPF, RIP, or Static.

d) **Process ID**— (Optional) Enter the process ID for the OSPF routing protocol.

e) Click **OK**.

Step 5 Click **Save**.

Configure BGP Network Settings

Network settings are used to add networks that will be advertised by the BGP routing process and route maps that will be examined to filter the networks to be advertised.

Procedure

Step 1 On the **Device Management** page, click **Routing**.

Step 2 Choose **BGP** > **IPv4** or **IPv6**.

Step 3 Click **Networks**.

Step 4 Click **Add** and update the **Add Networks** dialog:

a) **Network**— Enter the network to be advertised by the BGP routing processes.

b) (Optional) **Route Map**— Enter or select a route map that should be examined to filter the networks to be advertised. If not specified, all networks are redistributed.

c) Click **OK**.

Step 5 Click **Save**.

Configure BGP Redistribution Settings

Redistribution settings allow you to define the conditions for redistributing routes from another routing domain into BGP.

Procedure

Step 1 On the **Device Management** page, click **Routing**.

Step 2 Choose **BGP** > **IPv4** or **IPv6**.

Step 3 Click **Redistribution**.

Step 4 Click **Add** and update the **Add Redistribution** dialog:

a) **Source Protocol**— Select the protocol from which you want to redistribute routes into the BGP domain from the Source Protocol drop-down list.

b) **Process ID**— Enter the identifier for the selected source protocol. Applies to the OSPF protocol.

c) **Metric**— (Optional) Enter a metric for the redistributed route.

d) **Route Map**— Enter or select a route map that should be examined to filter the networks to be redistributed. If not specified, all networks are redistributed.

e) **Match**— The conditions used for redistributing routes from one routing protocol to another. The routes must match the selected condition to be redistributed. You can choose one or more of the following match conditions. These options are enabled only when OSPF is chosen as the Source Protocol.

- Internal
- External 1
- External 2
- NSSA External 1
- NSSA External 2

f) Click **OK**.

Configure BGP Route Injection Settings

Route Injection settings allow you to define the routes to be conditionally injected into the BGP routing table.

Procedure

Step 1 On the **Device Management** page, click **Routing**.

Step 2 Choose **BGP** > **IPv4** or **IPv6**.

Step 3 Click **Route Injection**.

Step 4 Click **Add** and update the **Add Route Injection** dialog:

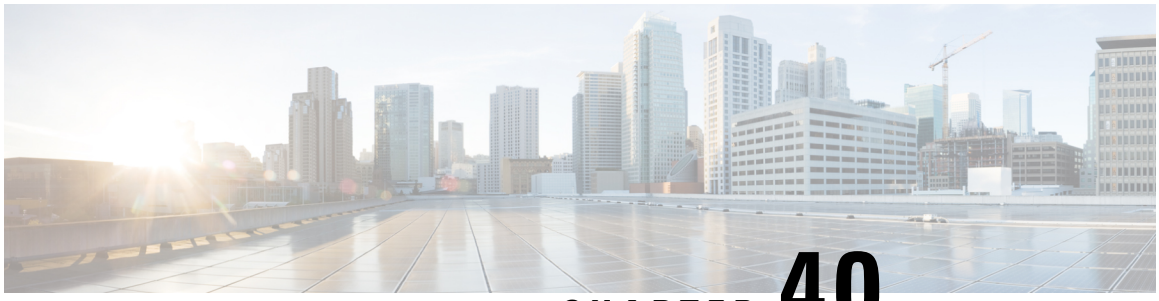
a) **Inject Map**— Enter or select the route map that specifies the prefixes to inject into the local BGP routing table.

b) **Exist Map**— Enter or select the route map containing the prefixes that the BGP speaker will track.

c) **Injected routes will inherit the attributes of the aggregate route**— Select this to configure the injected route to inherit attributes of the aggregate route.

d) Click **OK**.

Step 5 Click **Save**.



CHAPTER 40

RIP for Firepower Threat Defense

This chapter describes how to configure the Firepower Threat Defense to route data, perform authentication, and redistribute routing information, using the Routing Information Protocol (RIP).

- About RIP, on page 813
- Requirements and Prerequisites for RIP, on page 815
- Guidelines for RIP, on page 815
- Configure RIP, on page 816

About RIP

The Routing Information Protocol, or RIP, as it is more commonly called, is one of the most enduring of all routing protocols. RIP has four basic components: routing update process, RIP routing metrics, routing stability, and routing timers. Devices that support RIP send routing-update messages at regular intervals and when the network topology changes. These RIP packets include information about the networks that the devices can reach, as well as the number of routers or gateways that a packet must travel through to reach the destination address. RIP generates more traffic than OSPF, but is easier to configure.

RIP is a distance-vector routing protocol that uses hop count as the metric for path selection. When RIP is enabled on an interface, the interface exchanges RIP broadcasts with neighboring devices to dynamically learn about and advertise routes.

The Firepower Threat Defense device supports both RIP Version 1 and RIP Version 2. RIP Version 1 does not send the subnet mask with the routing update. RIP Version 2 sends the subnet mask with the routing update and supports variable-length subnet masks. Additionally, RIP Version 2 supports neighbor authentication when routing updates are exchanged. This authentication ensures that the Firepower Threat Defense device receives reliable routing information from a trusted source.

RIP has advantages over static routes because the initial configuration is simple, and you do not need to update the configuration when the topology changes. The disadvantage to RIP is that there is more network and processing overhead than in static routing.

Routing Update Process

RIP sends routing-update messages at regular intervals and when the network topology changes. When a router receives a routing update that includes changes to an entry, it updates its routing table to reflect the new route. The metric value for the path is increased by 1, and the sender is indicated as the next hop. RIP routers maintain only the best route (the route with the lowest metric value) to a destination. After updating

its routing table, the router immediately begins transmitting routing updates to inform other network routers of the change. These updates are sent independently of the regularly scheduled updates that RIP routers send.

RIP Routing Metric

RIP uses a single routing metric (hop count) to measure the distance between the source and a destination network. Each hop in a path from source to destination is assigned a hop count value, which is typically 1. When a router receives a routing update that contains a new or changed destination network entry, the router adds 1 to the metric value indicated in the update and enters the network in the routing table. The IP address of the sender is used as the next hop.

RIP Stability Features

RIP prevents routing loops from continuing indefinitely by implementing a limit on the number of hops allowed in a path from the source to a destination. The maximum number of hops in a path is 15. If a router receives a routing update that contains a new or changed entry, and if increasing the metric value by 1 causes the metric to be infinity (that is, 16), the network destination is considered unreachable. The downside of this stability feature is that it limits the maximum diameter of a RIP network to less than 16 hops.

RIP includes a number of other stability features that are common to many routing protocols. These features are designed to provide stability despite potentially rapid changes in network topology. For example, RIP implements the split horizon and hold-down mechanisms to prevent incorrect routing information from being propagated.

RIP Timers

RIP uses numerous timers to regulate its performance. Following are the timer stages for RIP:

- Update—The routing-update timer is the interval between periodic routing updates. This is how often the device sends routing updates. Generally, it is set to 30 seconds, with a small random amount of time added whenever the timer is reset. This is done to help prevent congestion, which could result from all routers simultaneously attempting to update their neighbors.
- Invalid—Each routing table entry has a route-timeout timer associated with it. This is the number of seconds since the device received the last valid update. When the route-timeout timer expires, the route is marked invalid but is retained in the table until the route-flush timer expires. Once this timer expires, the route goes into holddown. The default is 180 seconds (3 minutes).
- Holddown—The holddown period is the number of seconds the system waits before accepting any new updates for the route that is in holddown (that is, routes that have been marked invalid). The default is 180 seconds (3 minutes).
- Flush—The route-flush timer is the number of seconds since the system received the last valid update until the route is discarded and removed from the routing table. The default is 240 seconds (4 minutes).

As an example, when the interface on an adjacent router goes down, the system no longer receives routing updates from the adjacent router. At this time, the Invalid and Flush timers start increasing. In the first 180 seconds, nothing will happen. After 180 seconds, the invalid timer expires, making the route invalid, and the Holddown timer starts and holds the route for another 60 seconds. If there is still no update regarding the interface status on the adjacent router (that is, it is still down), then the route enters into the Flush state where in total the system has waited for 240 seconds from the last update (180 seconds for the Invalid timer and 60 seconds for Holddown timer), and the system flushes the route. Even if the adjacent routers interface comes

up immediately, the system does not accept a routing update until the Holddown timer completes the remaining 120 seconds.

Requirements and Prerequisites for RIP

Model Support

FTD

FTDv

Supported Domains

Any

User Roles

Admin

Network Admin

Guidelines for RIP

IPv6 Guidelines

Does not support IPv6.

Additional Guidelines

The following information applies to RIP Version 2 only:

- If using neighbor authentication, the authentication key and key ID must be the same on all neighbor devices that provide RIP Version 2 updates to the interface.
- With RIP Version 2, the Firepower Threat Defense device transmits and receives default route updates using the multicast address 224.0.0.9. In passive mode, it receives route updates at that address.
- When RIP Version 2 is configured on an interface, the multicast address 224.0.0.9 is registered on that interface. When a RIP Version 2 configuration is removed from an interface, that multicast address is unregistered.

Limitations

- The Firepower Threat Defense device cannot pass RIP updates between interfaces.
- RIP Version 1 does not support variable-length subnet masks.
- RIP has a maximum hop count of 15. A route with a hop count greater than 15 is considered unreachable.
- RIP convergence is relatively slow compared to other routing protocols.
- You can only enable a single RIP process on the Firepower Threat Defense device.

Configure RIP

RIP is a distance-vector routing protocol that uses hop count as the metric for path selection.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Select **Routing**.

Step 3 Select **RIP** from the table of contents.

Step 4 Select the **Enable RIP** check box to configure the RIP settings.

Step 5 Select the RIP versions for sending and receiving RIP updates from the **RIP Version** drop-down list.

Step 6 (Optional) Select the **Generate Default Route** check box to generate a default route for distribution, based on the route map that you specify.

a) Specify a route map name to use for generating default routes, in the **Route Map** field.
The default route 0.0.0.0/0 is generated for distribution over a certain interface , when the route map, specified in the **Route Map** field, is present.

Step 7 When Send and Receive Version 2 is the chosen RIP Version, the **Enable Auto Summary** option is available. When the **Enable Auto Summary** checkbox is checked, automatic route summarization is enabled. Disable automatic summarization if you must perform routing between disconnected subnets. When automatic summarization is disabled, subnets are advertised.

Note RIP Version 1 always uses automatic summarization—you cannot disable it.

Step 8 Click **Networks**. Define one or more networks for RIP routing. Enter IP address(es), or enter or select the desired Network/Hosts objects. There is no limit to the number of networks you can add to the security appliance configuration. Any interface that belongs to a network defined by this command, will participate in the RIP routing process. The RIP routing updates will be sent and received only through interfaces on the specified networks. Also, if the network of an interface is not specified, the interface will not be advertised in any RIP updates.

Note RIP only supports IPv4 objects.

Step 9 (Optional) Click **Passive Interface**. Use this option to specify passive interfaces on the appliance, and by extension the active interfaces. The device listens for RIP routing broadcasts on passive interfaces, using that information to populate its routing tables, but does not broadcast routing updates on passive interfaces. Interfaces that are not designated as passive, receive and send updates.

Step 10 Click **Redistribution** to manage redistribution routes. These are the routes that are being redistributed from other routing processes into the RIP routing process.

a) Click **Add** to specify redistribution routes.

b) Select the routing protocol to redistribute into the RIP routing process, in the **Protocol** drop-down list.

Note For the OSPF protocol, specify a process ID. Similarly, specify an AS path for BGP. When you choose the Connected option in the **Protocol** drop-down list, you can redistribute, directly connected networks into the RIP routing process.

c) (Optional) If you are redistributing OSPF routes into the RIP routing process, you can select specific types of OSPF routes to redistribute in the **Match** drop-down list . Ctrl-click to select multiple types:

- Internal – Routes internal to the autonomous system (AS) are redistributed.
- External 1 – Type 1 routes external to the AS are redistributed.
- External 2 – Type 2 routes external to the AS are redistributed.
- NSSA External 1 – Type 1 routes external to a not-so-stubby area (NSSA) are redistributed.
- NSSA External 2 – Type 2 routes external to an NSSA are redistributed

Note The default is match Internal, External 1, and External 2

d) Select the RIP metric type to apply to the redistributed routes in the **Metric** drop-down list. The two choices are:

- Transparent – Use the current route metric
- Specified Value – Assign a specific metric value. Enter a specific value from 0-16, in the **Metric Value** field.
- None – No metric is specified. Do not use any metric value, to apply to redistributed routes.

e) (Optional) Enter the name of a route map that must be satisfied, in the **Route Map** field before the route can be redistributed into the RIP routing process. Routes are redistributed only if IP address matches an allow statement in the route map address list.

f) Click **OK**.

Step 11 (Optional) Click **Filtering** to manage filters for the RIP policy. In this section, filters are used to prevent routing updates through an interface, control the advertising of routes in routing updates, control the processing of routing updates and filtering sources of routing updates.

a) Click **Add** to add RIP filters.

b) Select the type of traffic to be filtered - Inbound or Outbound in the **Traffic Direction** field.

Note If traffic direction is inbound, you can only define an Interface filter.

c) Specify whether the filter is based on an Interface or a Route, by selecting appropriate in the **Filter On** field. If you select Interface, enter or Select the name of the interface on which routing updates are to be filtered. If you select Route, choose the route type:

- Static – Only static routes are filtered.
- Connected – Only connected routes are filtered.
- OSPF – Only OSPFv2 routes discovered by the specified OSPF process are filtered. Enter the Process ID of the OSPF process to be filtered.
- BGP – Only BGPv4 routes discovered by the specified BGP process are filtered. Enter the AS path of the BGP process to be filtered.

d) In the **Access List** field, enter or select the name of one or more access control lists (ACLs) that define the networks to be allowed or removed from RIP route advertisements.

e) Click **OK**.

Step 12 (Optional) Click **Broadcast** to add or edit interface configurations. Using Broadcastf, you can override the global RIP versions to send or receive per interface. You can also define the authentication parameters per interface if you want to implement authentication to ensure valid RIP updates.

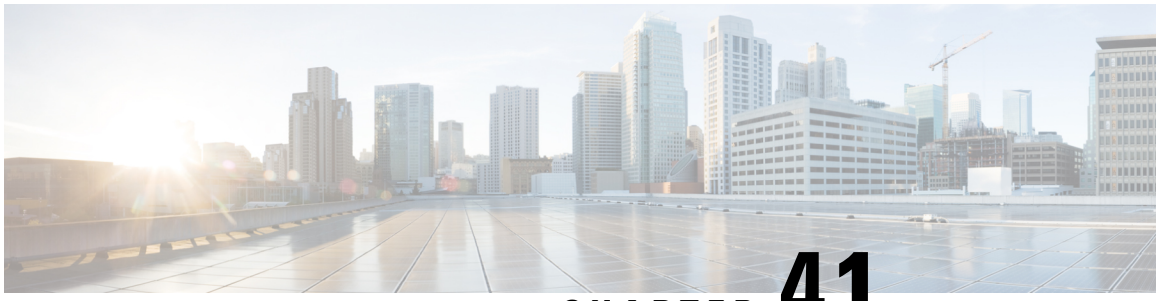
a) Click **Add** to add interface configurations.
b) Enter or Select an interface defined on this appliance in the **Interface** field.
c) In the Send option, select the appropriate boxes to specify sending updates using the RIP **Version 1**, **Version 2**, or both. These options let you override, for the specified interface, the global Send versions specified .
d) In the Receive option, select the appropriate boxes to specify accepting updates using the RIP **Version 1**, **Version 2**, or both. These options let you override, for the specified interface, the global Receive versions specified .
e) Select the **Authentication** used on this interface for RIP broadcasts.

   - None – No authentication
   - MD5 – Employ MD5
   - Clear Text – Employ clear-text authentication

   If you choose MD5 or Clear Text, you must also provide the following authentication parameters.

   - Key ID – The ID of the authentication key. Valid values are from 0 to 255.
   - Key – The key used by the chosen authentication method. Can contain up to 16 characters
   - Confirm – Enter the authentication key again, to confirm

f) Click **OK**.



CHAPTER 41

Multicast Routing for Firepower Threat Defense

This chapter describes how to configure the Firepower Threat Defense device to use the multicast routing protocol.

- About Multicast Routing, on page 819
- Requirements and Prerequisites for Multicast Routing, on page 823
- Guidelines for Multicast Routing, on page 823
- Configure IGMP Features, on page 824
- Configure PIM Features, on page 829
- Configure Multicast Routes, on page 835
- Configure Multicast Boundary Filters, on page 835

About Multicast Routing

Multicast routing is a bandwidth-conserving technology that reduces traffic by simultaneously delivering a single stream of information to thousands of corporate recipients and homes. Applications that take advantage of multicast routing include videoconferencing, corporate communications, distance learning, and distribution of software, stock quotes, and news.

Multicast routing protocols deliver source traffic to multiple receivers without adding any additional burden on the source or the receivers while using the least network bandwidth of any competing technology. Multicast packets are replicated in the network by FTD device enabled with Protocol Independent Multicast (PIM) and other supporting multicast protocols, which results in the most efficient delivery of data to multiple receivers possible.

The FTD device supports both stub multicast routing and PIM multicast routing. However, you cannot configure both concurrently on a single FTD device.

Note The UDP and non-UDP transports are both supported for multicast routing. However, the non-UDP transport has no FastPath optimization.

IGMP Protocol

IP hosts use the Internet Group Management Protocol (IGMP) to report their group memberships to directly-connected multicast routers. IGMP is used to dynamically register individual hosts in a multicast

group on a particular LAN. Hosts identify group memberships by sending IGMP messages to their local multicast router. Under IGMP, routers listen to IGMP messages and periodically send out queries to discover which groups are active or inactive on a particular subnet.

IGMP uses group addresses (Class D IP address) as group identifiers. Host group address can be in the range of 224.0.0.0 to 239.255.255.255. The address 224.0.0.0 is never assigned to any group. The address 224.0.0.1 is assigned to all systems on a subnet. The address 224.0.0.2 is assigned to all routers on a subnet.



Note When you enable multicast routing on the Firepower Threat Defense device, IGMP Version 2 is automatically enabled on all interfaces.

Query Messages to Multicast Groups

The Firepower Threat Defense device sends query messages to discover which multicast groups have members on the networks attached to the interfaces. Members respond with IGMP report messages indicating that they want to receive multicast packets for specific groups. Query messages are addressed to the all-systems multicast group, which has an address of 224.0.0.1, with a time-to-live value of 1.

These messages are sent periodically to refresh the membership information stored on the Firepower Threat Defense device. If the Firepower Threat Defense device discovers that there are no local members of a multicast group still attached to an interface, it stops forwarding multicast packets for that group to the attached network, and it sends a prune message back to the source of the packets.

By default, the PIM designated router on the subnet is responsible for sending the query messages. By default, they are sent once every 125 seconds.

When changing the query response time, by default, the maximum query response time advertised in IGMP queries is 10 seconds. If the Firepower Threat Defense device does not receive a response to a host query within this amount of time, it deletes the group.

Stub Multicast Routing

Stub multicast routing provides dynamic host registration and facilitates multicast routing. When configured for stub multicast routing, the FTD device acts as an IGMP proxy agent. Instead of fully participating in multicast routing, the FTD device forwards IGMP messages to an upstream multicast router, which sets up delivery of the multicast data. When configured for stub multicast routing, the FTD device cannot be configured for PIM sparse or bidirectional mode. You must enable PIM on the interfaces participating in IGMP stub multicast routing.

The FTD device supports both PIM-SM and bidirectional PIM. PIM-SM is a multicast routing protocol that uses the underlying unicast routing information base or a separate multicast-capable routing information base. It builds unidirectional shared trees rooted at a single Rendezvous Point (RP) per multicast group and optionally creates shortest-path trees per multicast source.

PIM Multicast Routing

Bidirectional PIM is a variant of PIM-SM that builds bidirectional shared trees connecting multicast sources and receivers. Bidirectional trees are built using a Designated Forwarder (DF) election process operating on each link of the multicast topology. With the assistance of the DF, multicast data is forwarded from sources to the Rendezvous Point (RP), and therefore along the shared tree to receivers, without requiring source-specific state. The DF election takes place during RP discovery and provides a default route to the RP.



Note If the FTD device is the PIM RP, use the untranslated outside address of the FTD device as the RP address.

PIM Source Specific Multicast Support

The FTD device does not support PIM Source Specific Multicast (SSM) functionality and related configuration. However, the FTD device allows SSM-related packets to pass through unless it is placed as a last-hop router.

SSM is classified as a data delivery mechanism for one-to-many applications such as IPTV. The SSM model uses a concept of "channels" denoted by an (S,G) pair, where S is a source address and G is an SSM destination address. Subscribing to a channel is achieved by using a group management protocol such as IGMPv3. SSM enables a receiving client, once it has learned about a particular multicast source, to receive multicast streams directly from the source rather than receiving it from a shared Rendezvous Point (RP). Access control mechanisms are introduced within SSM providing a security enhancement not available with current sparse or sparse-dense mode implementations.

PIM-SSM differs from PIM-SM in that it does not use an RP or shared trees. Instead, information on source addresses for a multicast group is provided by the receivers through the local receivership protocol (IGMPv3) and is used to directly build source-specific trees.

Multicast Bidirectional PIM

Multicast bidirectional PIM is useful for networks that have many sources and receivers talking to each other simultaneously and where each participant can become both the source and receiver of multicast traffic, such as in videoconferencing, Webex meetings, and group chat. When PIM bidirectional mode is used, the RP only creates the (*,G) entry for the shared tree. There is no (S,G) entry. This conserves resources on the RP because state tables for each (S,G) entry are not maintained.

In PIM sparse mode, traffic only flows down the shared tree. In PIM bidirectional mode, traffic flows up and down the shared tree.

PIM bidirectional mode also does not use the PIM register/register-stop mechanism to register sources to the RP. Each source can begin sending to the source at any time. When the multicast packets arrive at the RP, they are forwarded down the shared tree (if there are receivers) or dropped (when there are no receivers). However, there is no way for the RP to tell the source to stop sending multicast traffic.

Design-wise you must think about where to place the RP in your network because it should be somewhere in the middle between the sources and receivers in the network.

PIM bidirectional mode has no Reverse Path Forwarding (RPF) check. Instead it uses the concept of a Designated Forwarder (DF) to prevent loops. This DF is the only router on the segment that is allowed to send multicast traffic to the RP. If there is only one router per segment that forwards multicast traffic, there will be no loops. The DF is chosen using the following mechanism:

- The router with the lowest metric to the RP is the DF.
- If the metric is equal, then the router with the highest IP address becomes the DF.

PIM Bootstrap Router (BSR)

PIM Bootstrap Router (BSR) is a dynamic Rendezvous Point (RP) selection model that uses candidate routers for RP function and for relaying the RP information for a group. The RP function includes RP discovery and provides a default route to the RP. It does this by configuring a set of devices as candidate BSRs (C-BSR) which participate in a BSR election process to choose a BSR amongst themselves. Once the BSR is chosen, devices that are configured as candidate Rendezvous Points (C-RP) start sending their group mapping to the elected BSR. The BSR then distributes the group-to-RP mapping information to all the other devices down the multicast tree through BSR messages that travel from PIM router to PIM router on a per-hop basis.

This feature provides a means of dynamically learning RPs, which is very essential in large complex networks where an RP can periodically go down and come up.

PIM Bootstrap Router (BSR) Terminology

The following terms are frequently referenced in the PIM BSR configuration:

- Bootstrap Router (BSR) — A BSR advertises Rendezvous Point (RP) information to other routers with PIM on a hop-by-hop basis. Among multiple Candidate-BSRs, a single BSR is chosen after an election process. The primary purpose of this Bootstrap router is to collect all Candidate-RP (C-RP) announcements in to a database called the RP-set and to periodically send this out to all other routers in the network as BSR messages (every 60 seconds).
- Bootstrap Router (BSR) messages — BSR messages are multicast to the All-PIM-Routers group with a TTL of 1. All PIM neighbors that receive these messages retransmit them (again with a TTL of 1) out of all interfaces except the one in which the messages were received. BSR messages contain the RP-set and the IP address of the currently active BSR. This is how C-RPs know where to unicast their C-RP messages.
- Candidate Bootstrap Router (C-BSR) — A device that is configured as a candidate-BSR participates in the BSR election mechanism. A C-BSR with highest priority is elected as the BSR. The highest IP address of the C-BSR is used as a tiebreaker. The BSR election process is preemptive, for example if a new C-BSR with a higher priority comes up, it triggers a new election process.
- Candidate Rendezvous Point (C-RP) — An RP acts as a meeting place for sources and receivers of multicast data. A device that is configured as a C-RP periodically advertises the multicast group mapping information directly to the elected BSR through unicast. These messages contain the Group-range, C-RP address, and a hold time. The IP address of the current BSR is learned from the periodic BSR messages that are received by all routers in the network. In this way, the BSR learns about possible RPs that are currently up and reachable.



Note The FTD device does not act as a C-RP, even though the C-RP is a mandatory requirement for BSR traffic. Only routers can act as a C-RP. So, for BSR testing functionality, you must add routers to the topology.

- BSR Election Mechanism — Each C-BSR originates Bootstrap messages (BSMs) that contain a BSR Priority field. Routers within the domain flood the BSMs throughout the domain. A C-BSR that hears about a higher-priority C-BSR than itself suppresses its sending of further BSMs for some period of time. The single remaining C-BSR becomes the elected BSR, and its BSMs inform all the other routers in the domain that it is the elected BSR.

Multicast Group Concept

Multicast is based on the concept of a group. An arbitrary group of receivers expresses an interest in receiving a particular data stream. This group does not have any physical or geographical boundaries—the hosts can be located anywhere on the Internet. Hosts that are interested in receiving data flowing to a particular group must join the group using IGMP. Hosts must be a member of the group to receive the data stream.

Multicast Addresses

Multicast addresses specify an arbitrary group of IP hosts that have joined the group and want to receive traffic sent to this group.

Clustering

Multicast routing supports clustering. In Spanned EtherChannel clustering, the control unit sends all multicast routing packets and data packets until fast-path forwarding is established. After fast-path forwarding is established, data units may forward multicast data packets. All data flows are full flows. Stub forwarding flows are also supported. Because only one unit receives multicast packets in Spanned EtherChannel clustering, redirection to the control unit is common.

Requirements and Prerequisites for Multicast Routing

Model Support

FTD

FTDv

Supported Domains

Any

User Roles

Admin

Network Admin

Guidelines for Multicast Routing

Firewall Mode

Supported only in routed firewall mode. Transparent firewall mode is not supported.

IPv6

Does not support IPv6.

Multicast Group

The range of addresses between 224.0.0.0 and 224.0.0.255 is reserved for the use of routing protocols and other topology discovery or maintenance protocols, such as gateway discovery and group membership reporting. Hence, Internet multicast routing from address range 224.0.0/24 is not supported; IGMP group is not created when enabling multicast routing for the reserved addresses.

Clustering

In clustering, for IGMP and PIM, this feature is only supported on the primary unit.

Additional Guidelines

- You must configure an access control or prefilter rule on the inbound security zone to allow traffic to the multicast host, such as 224.1.2.3. However, you cannot specify a destination security zone for the rule, or it cannot be applied to multicast connections during initial connection validation.
- You cannot disable an interface with PIM configured on it. If you have configured PIM on the interface (see #unique_1199), disabling the multicast routing and PIM does not remove the PIM configuration. You must remove (delete) the PIM configuration to disable the interface.
- PIM/IGMP Multicast routing is not supported on interfaces in a traffic zone.
- Do not configure FTD to simultaneously be a Rendezvous Point (RP) and a First Hop Router.
- HSRP standby IP address does not participate in PIM neighborship. Thus, if the RP router IP is routed through a HSRP standby IP address, the multicast routing does not work in FTD. Hence for the multicast traffic to pass through successfully, ensure that the route for the RP address is not the HSRP standby IP address, instead, configure the route address to an interface IP address.

Configure IGMP Features

IP hosts use IGMP to report their group memberships to directly-connected multicast routers. IGMP is used to dynamically register individual hosts in a multicast group on a particular LAN. Hosts identify group memberships by sending IGMP messages to their local multicast router. Under IGMP, routers listen to IGMP messages and periodically send out queries to discover which groups are active or inactive on a particular subnet.

This section describes how to configure optional IGMP settings on a per-interface basis.

Procedure

Step 1 Enable Multicast Routing, on page 825
Step 2 #unique_1202.
Step 3 #unique_1203.
Step 4 #unique_1204.
Step 5 #unique_1205.

Enable Multicast Routing

Enabling multicast routing on the Firepower Threat Defense device, enables IGMP and PIM on all interfaces by default. IGMP is used to learn whether members of a group are present on directly attached subnets. Hosts join multicast groups by sending IGMP report messages. PIM is used to maintain forwarding tables to forward multicast datagrams.

Note Only the UDP transport layer is supported for multicast routing.

The following list shows the maximum number of entries for specific multicast tables. Once these limits are reached, any new entries are discarded.

- MFIB—30,000
- IGMP Groups—30,000
- PIM Routes—72,000

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **IGMP**.

Step 3 Check the **Enable Multicast Routing** check box.

Checking this check box enables IP multicast routing on the Firepower Threat Defense device. Unchecking this check box disables IP multicast routing. By default, multicast is disabled. Enabling multicast routing enables multicast on all interfaces.

You can disable multicast on a per-interface basis. This is useful if you know that there are no multicast hosts on a specific interface and you want to prevent the Firepower Threat Defense device from sending host query messages on that interface.

Configure IGMP Protocol

You can configure IGMP parameters per interface, such as the forward interface, query messages, and time intervals.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **IGMP**.

Step 3 On **Protocol**, click **Add** or **Edit**.

Use the **Add IGMP parameters** dialog box to add new IGMP parameters to the Firepower Threat Defense device. Use the **Edit IGMP parameters** dialog box to change existing parameters.

Step 4 Configure the following options:

- **Interface**—From the drop-down list, select the interface for which you want to configure IGMP protocol.
- **Enable IGMP**—Check the check box to enable IGMP.

  Note Disabling IGMP on specific interfaces is useful if you know that there are no multicast hosts on a specific interface and you want to prevent the Firepower Threat Defense device from sending host query messages on that interface.

- **Forward Interface**—From the drop-down list, select the specific interface from which you want to forward IGMP messages.

  This configures the Firepower Threat Defense device to act as an IGMP proxy agent and forward IGMP messages from hosts connected on one interface to an upstream multicast router on another interface.

- **Version**—Choose IGMP Version 1 or 2.

  By default, the Firepower Threat Defense device runs IGMP Version 2, which enables several additional features.

  Note All multicast routers on a subnet must support the same version of IGMP. The Firepower Threat Defense device does not automatically detect Version 1 routers and switch to Version 1. However, you can have a mix of IGMP Version 1 and 2 hosts on the subnet; the Firepower Threat Defense device running IGMP Version 2 works correctly when IGMP Version 1 hosts are present.

- **Query Interval**—The interval in seconds at which the designated router sends IGMP host-query messages. The range is 1 to 3600. The default is 125.

  Note If the Firepower Threat Defense device does not hear a query message on an interface for the specified timeout value, then the Firepower Threat Defense device becomes the designated router and starts sending the query messages.

- **Response Time**—The interval in seconds before the Firepower Threat Defense device deletes the group. The range is 1 to 25. The default is 10.

  If the Firepower Threat Defense device does not receive a response to a host query within this amount of time, it deletes the group.

- **Group Limit**—The maximum number of hosts that can join on an interface. The range is 1 to 500. The default is 500.

  You can limit the number of IGMP states resulting from IGMP membership reports on a per-interface basis. Membership reports exceeding the configured limits are not entered in the IGMP cache, and traffic for the excess membership reports is not forwarded

- **Query Timeout**—The period of time in seconds before which the Firepower Threat Defense device takes over as the requester for the interface after the previous requester has stopped. The range is 60 to 300. The default is 255.

Step 5 Click **OK** to save the IGMP protocol configuration.

Configure IGMP Access Groups

You can control access to multicast groups by using access control lists.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **Access Group**.

Step 3 On **Access Group**, click **Add** or **Edit**.

Use the **Add IGMP Access Group parameters** dialog box to add new IGMP access groups to the Access Group table. Use the **Edit IGMP Access Group parameters** dialog box to change existing parameters.

Step 4 Configure the following options:

a) From the **Interface** drop-down list, select the interface with which the access group is associated. You cannot change the associated interface when you are editing an existing access group.

b) Click one of the following:

- **Standard Access List**— From the **Standard Access List** drop-down list, select the standard ACL or click **Add** ( ) to create a new standard ACL. See Configure Standard ACL Objects, on page 469 for the procedure.
- **Extended Access List**—From the **Extended Access List** drop-down list, select the extended ACL or click **Add** ( ) to create a new extended ACL. See Configure Extended ACL Objects, on page 468 for the procedure.

Step 5 Click **OK** to save the access group configuration.

Configure IGMP Static Groups

Sometimes a group member cannot report its membership in the group or there may be no members of a group on the network segment, but you still want multicast traffic for that group to be sent to that network segment. You can have multicast traffic for that group sent to the segment by configuring a statically joined IGMP group. With this method, the Firepower Threat Defense device does not accept the packets itself, but only forwards them. Therefore, this method allows fast switching. The outgoing interface appears in the IGMP cache, but this interface is not a member of the multicast group. When you configure IGMP static groups, ensure that the Firepower Threat Defense is configured as the destination router on the interface.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **IGMP**.

Step 3 On **Static Group**, click **Add** or **Edit**.

Use the **Add IGMP Static Group parameters** dialog box to statically assign a multicast group to an interface. Use the **Edit IGMP Static Group parameters** dialog box to change existing static group assignments.

Step 4 Configure the following options:

- From the **Interface** drop-down list, select the interface to which you want to statically assign a multicast group. If you are editing an existing entry, you cannot change the value.
- From the **Multicast Groups** drop-down list, select the multicast group to which you want to assign the interface, or click **Add** ( ) to create a new multicast group. See Creating Network Objects for the procedure.

Step 5 Click **OK** to save the static group configuration.

Configure IGMP Join Groups

You can configure an interface to be a member of a multicast group. Configuring the Firepower Threat Defense device to join a multicast group causes upstream routers to maintain multicast routing table information for that group and keep the paths for that group active. When you configure IGMP join groups, ensure that the Firepower Threat Defense is configured as the destination router on the interface.

Note See #unique_1204 if you want to forward multicast packets for a specific group to an interface without the Firepower Threat Defense device accepting those packets as part of the group.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **IGMP**.

Step 3 On **Join Group**, click **Add** or **Edit**.

Use the **Add IGMP Join Group parameters** dialog box to configure the Firepower Threat Defense device to be a member of a multicast group. Use the **Edit IGMP Join Group parameters** dialog box to change existing parameters.

Step 4 Configure the following options:

- From the **Interface** drop-down list, select the interface you want to be a member of a multicast group. If you are editing an existing entry, you cannot change the value.
- From the **Join Group** drop-down list, select the multicast group to which you want to assign the interface, or click **Plus** to create a new multicast group. See Creating Network Objects for the procedure.

Configure PIM Features

Routers use PIM to maintain forwarding tables to use for forwarding multicast diagrams. When you enable multicast routing on the Firepower Threat Defense device, PIM and IGMP are automatically enabled on all interfaces.

Note PIM is not supported with PAT. The PIM protocol does not use ports, and PAT only works with protocols that use ports.

This section describes how to configure optional PIM settings.

Procedure

Step 1 #unique_1199
Step 2 #unique_1211
Step 3 #unique_1212
Step 4 #unique_1213
Step 5 #unique_1214
Step 6 #unique_1215
Step 7 Configure Multicast Boundary Filters, on page 835

Configure PIM Protocol

You can enable or disable PIM on a specific interface.

You can also configure the Designated Router (DR) priority. The DR is responsible for sending PIM register, join, and prune messages to the RP. When there is more than one multicast router on a network segment, choosing the DR is based on the DR priority. If multiple devices have the same DR priority, then the device with the highest IP address becomes the DR. By default, the FTD device has a DR priority of 1.

Router query messages are used to choose the PIM DR. The PIM DR is responsible for sending router query messages. By default, router query messages are sent every 30 seconds. Additionally, every 60 seconds, the FTD device sends PIM join or prune messages.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.
Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.
Step 3 On **Protocol**, click **Add** or **Edit**.

Use the **Add PIM parameters** dialog box to add new PIM parameters to the interface. Use the **Edit PIM parameters** dialog box to change existing parameters.

Step 4 Configure the following options:

- **Interface**—From the drop-down list, select the interface for which you want to configure PIM protocol.
- **Enable PIM**—Check the check box to enable PIM.
- **DR Priority**—The value for the DR for the selected interface. The router with the highest DR priority on the subnet becomes the designated router. Valid values range from 0 to 4294967294. The default DR priority is 1. Setting this value to 0 makes the FTD device interface ineligible to become the designated router.
- **Hello Interval**—The interval in seconds at which the interface sends PIM hello messages. The range is 1 to 3600. The default is 30.
- **Join Prune Interval**—The interval in seconds at which the interface sends PIM join and prune advertisements. The range is 10 to 600. The default is 60.

Step 5 Click **OK** to save the PIM protocol configuration.

Configure PIM Neighbor Filters

You can define the routers that can become PIM neighbors. By filtering the routers that can become PIM neighbors, you can do the following:

- Prevent unauthorized routers from becoming PIM neighbors.
- Prevent attached stub routers from participating in PIM.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 On **Neighbor Filter**, click **Add** or **Edit**.

Use the **Add PIM Neighbor Filter** dialog box to add new PIM neighbor filters to the interface. Use the **Edit PIM Neighbor Filter** dialog box to change existing parameters.

Step 4 Configure the following options:

- From the **Interface** drop-down list, select the interface to which you want to add a PIM neighbor filter.
- **Standard Access List**— From the **Standard Access List** drop-down list, select a standard ACL or click **Add** ( ) to create a new standard ACL. See Configure Standard ACL Objects, on page 469 for the procedure.

  Note Choosing **Allow** on the **Add Standard Access List Entry** dialog box lets the multicast group advertisements pass through the interface. Choosing **Block** prevents the specified multicast group advertisements from passing through the interface. When a multicast boundary is configured on an interface, all multicast traffic is prevented from passing through the interface unless permitted with a neighbor filter entry.

Step 5 Click **OK** to save the PIM neighbor filter configuration.

Configure PIM Bidirectional Neighbor Filters

A PIM bidirectional neighbor filter is an ACL that defines the neighbor devices that can participate in the Designated Forwarder (DF) election. If a PIM bidirectional neighbor filter is not configured for an interface, there are no restrictions. If a PIM bidirectional neighbor filter is configured, only those neighbors permitted by the ACL can participate in the DF election process.

Bidirectional PIM allows multicast routers to keep reduced state information. All of the multicast routers in a segment must be bidirectionally enabled to elect a DF.

When a PIM bidirectional neighbor filter is enabled, the routers that are permitted by the ACL are considered to be bidirectionally capable. Therefore, the following is true:

- If a permitted neighbor does not support bidirectional mode, then the DF election does not occur.
- If a denied neighbor supports bidirectional mode, then the DF election does not occur.
- If a denied neighbor does not support bidirectional mode, the DF election can occur.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Choose **Multicast Routing** > **PIM**.

Step 3 On **Bidirectional Neighbor Filter**, click **Add** or **Edit**.

Use the **Add PIM Bidirectional Neighbor Filter** dialog box to create ACL entries for the PIM bidirectional neighbor filter ACL. Use the **Edit PIM Bidirectional Neighbor Filter** dialog box to change existing parameters.

Step 4 Configure the following options:

- From the **Interface** drop-down list, select the interface to which you want to configure the PIM bidirectional neighbor filter ACL entry.
- **Standard Access List**— From the **Standard Access List** drop-down list, select a standard ACL or click **Add** () to create a new standard ACL. See Configure Standard ACL Objects, on page 469 for the procedure.

  Note Choosing **Allow** on the **Add Standard Access List Entry** dialog box lets the specified devices participate in the DR election process. Choosing **Block** prevents the specified devices from participating in the DR election process.

Step 5 Click **OK** to save the PIM bidirectional neighbor filter configuration.

Configure PIM Rendezvous Points

You can configure the Firepower Threat Defense device to serve as a RP to more than one group. The group range specified in the ACL determines the PIM RP group mapping. If an ACL is not specified, then the RP for the group is applied to the entire multicast group range (224.0.0.0/4). See Multicast Bidirectional PIM, on page 821 for more information about bidirectional PIM.

The following restrictions apply to RPs:

- You cannot use the same RP address twice.
- You cannot specify All Groups for more than one RP.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 On **Rendezvous Points**, click **Add** or **Edit**.

Use the **Add Rendezvous Point** dialog box to create a new entry to the Rendezvous Point table. Use the **Edit Rendezvous Point** dialog box to change existing parameters.

Step 4 Configure the following options:

- From the **Rendezvous Point IP address** drop-down list, select the IP address that you want to add as an RP or click **Add** ( ) to create a new network object. See Creating Network Objects for the procedure.
- Check the **Use bi-directional forwarding** check box if the specified multicast groups are to operate in bidirectional mode. In bidirectional mode, if the Firepower Threat Defense device receives a multicast packet and has no directly connected members or PIM neighbors present, it sends a prune message back to the source.
- Choose **Use this RP for all Multicast Groups** to to use the specified RP for all multicast groups on the interface.
- Choose the **Use this RP for all Multicast Groups as specified below** to designate the multicast groups to use with the specified RP and then from the **Standard Access List** drop-down list, choose a standard ACL or click **Add** ( ) to create a new standard ACL. See Configure Standard ACL Objects, on page 469 for the procedure.

Step 5 Click **OK** to save the rendezvous point configuration.

Configure PIM Route Trees

By default, PIM leaf routers join the shortest-path tree immediately after the first packet arrives from a new source. This method reduces delay, but requires more memory than the shared tree. You can configure whether or not the Firepower Threat Defense device should join the shortest-path tree or use the shared tree, either for all multicast groups or only for specific multicast addresses.

The shortest-path tree is used for any group that is not specified in the Multicast Groups table. The Multicast Groups table displays the multicast groups to use with the shared tree. The table entries are processed from the top down. You can create an entry that includes a range of multicast groups, but excludes specific groups within that range by placing deny rules for the specific groups at the top of the table and the permit rule for the range of multicast groups below the deny statements.

Note This behavior is known as Shortest Path Switchover (SPT). We recommend that you always use the Shared Tree option.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 On **Route Tree**, select the path for the route tree:

- Click **Shortest Path** to use the shortest-path tree for all multicast groups.
- Click **Shared Tree** to use the shared tree for all multicast groups.
- Click **Shared tree for below mentioned group** to designate the groups specified in the Multicast Groups table, and then from the **Standard Access List** drop-down list, select a standard ACL or click **Add** ( ) to create a new standard ACL. See Configure Standard ACL Objects, on page 469 for the procedure.

Step 4 Click **OK** to save the route tree configuration.

Configure PIM Request Filters

When the Firepower Threat Defense device is acting as an RP, you can restrict specific multicast sources from registering with it to prevent unauthorized sources from registering with the RP. You can define the multicast sources from which the Firepower Threat Defense device will accept PIM register messages.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 On **Request Filter**, define the multicast sources that are allowed to register with the Firepower Threat Defense device when it acts as an RP:

- From the **Filter PIM register messages using:** drop-down list select **None**, **Access List**, or **Route Map**.
- If you choose **Access List** from the drop-down list, select an extended ACL or click **Add** ( ) to create a new extended ACL. See Configure Extended ACL Objects, on page 468 for the procedure.

Note In the **Add Extended Access List Entry** dialog box, select **Allow** from the drop-down list o create a rule that allows the specified source of the specified multicast traffic to register with the Firepower Threat Defense device, or select **Block** to create a rule that prevents the specified source of the specified multicast traffic from registering with the Firepower Threat Defense device.

- If you choose **Route Map**, select a route map from the **Route Map** drop-down list, or click **Add** ( ) to create a new route map. See Creating Network Objects for the procedure.

Step 4 Click **OK** to save the request filter configuration.

Configure the Firepower Threat Defense Device as a Candidate Bootstrap Router

You can configure the Firepower Threat Defense device as a candidate BSR.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 On **Bootstrap Router**, check the **Configure this FTD as a Candidate Bootstrap Router (C-BSR)** check box to perform the C-BSR setup.

a) From the **Interface** drop-down list, select the interface on the Firepower Threat Defense device from which the BSR address is derived to make it a candidate.

This interface must be enabled with PIM.

b) In the **Hash mask length** field, enter the length of a mask (32 bits maximum) that is to be ANDed with the group address before the hash function is called. All groups with the same seed hash (correspond) to the same RP. For example, if this value is 24, only the first 24 bits of the group addresses matter. This fact allows you to get one RP for multiple groups. The range is 0 to 32.

c) In the **Priority** field, enter the priority of the candidate BSR. The BSR with the larger priority is preferred. If the priority values are the same, the router with the larger IP address is the BSR. The range is 0 to 255. The default value is 0.

Step 4 (Optional) Click **Add** ( ) to select an interface on which no PIM BSR messages will be sent or received in the **Configure this FTD as a Border Bootstrap Router (BSR)** section.

- From the **Interface** drop-down list, select the interface on which no PIM BSR messages will be sent or received.

  RP or BSR advertisements are filtered effectively isolating two domains of RP information exchange.

- Check the **Enable Border BSR** check box to enable BSR.

Step 5 Click **OK** to save the bootstrap router configuration.

Configure Multicast Routes

Configuring static multicast routes lets you separate multicast traffic from unicast traffic. For example, when a path between a source and destination does not support multicast routing, the solution is to configure two multicast devices with a GRE tunnel between them and to send the multicast packets over the tunnel.

When using PIM, the Firepower Threat Defense device expects to receive packets on the same interface where it sends unicast packets back to the source. In some cases, such as bypassing a route that does not support multicast routing, you may want unicast packets to take one path and multicast packets to take another.

Static multicast routes are not advertised or redistributed.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **Multicast Routes** > **Add or Edit**.

Use the **Add Multicast Route Configuration** dialog box to add a new multicast route to the Firepower Threat Defense device. Use the **Edit Multicast Route Configuration** dialog box to change an existing multicast route.

Step 3 From the **Source Network** drop-down box, select an existing network or click **Add** ( ) to add a new one. See Creating Network Objects for the procedure.

Step 4 To configure an interface to forward the route, click **Interface** and configure the following options:

- From the **Source Interface** drop-down list, select the incoming interface for the multicast route.
- From the **Output Interface/Dense** drop-down list, select the destination interface that the route is forwarded through.
- In the **Distance** field, enter the distance of the multicast route. The range is 0 to 255.

Step 5 To configure an RPF address to forward the route, click **Address** and configure the following options:

- In the **RPF Address** field, enter the IP address for the multicast route.
- In the **Distance** field, enter the distance of the multicast route The range is 0 to 255.

Step 6 Click **OK** to save the multicast routes configuration.

Configure Multicast Boundary Filters

Address scoping defines domain boundary filters so that domains with RPs that have the same IP address do not leak into each other. Scoping is performed on the subnet boundaries within large domains and on the boundaries between the domain and the Internet.

You can set up an administratively scoped boundary filter on an interface for multicast group addresses. IANA has designated the multicast address range from 239.0.0.0 to 239.255.255.255 as the administratively scoped

addresses. This range of addresses can be reused in domains administered by different organizations. The addresses would be considered local, not globally unique.

A standard ACL defines the range of affected addresses. When a boundary filter is set up, no multicast data packets are allowed to flow across the boundary from either direction. The boundary filter allows the same multicast group address to be reused in different administrative domains.

You can configure, examine, and filter Auto-RP discovery and announcement messages at the administratively scoped boundary. Any Auto-RP group range announcements from the Auto-RP packets that are denied by the boundary ACL are removed. An Auto-RP group range announcement is permitted and passed by the boundary filter only if all addresses in the Auto-RP group range are permitted by the boundary ACL. If any address is not permitted, the entire group range is filtered and removed from the Auto-RP message before the Auto-RP message is forwarded.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firepower Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **Multicast Boundary Filter**, and then click **Add** or **Edit**.

Use the **Add Multicast Boundary Filter** dialog box to add new multicast boundary filters to the Firepower Threat Defense device. Use the **Edit Multicast Boundary Filter** dialog box to change existing parameters.

You can configure a multicast boundary for administratively scoped multicast addresses. A multicast boundary restricts multicast data packet flows and enables reuse of the same multicast group address in different administrative domains. When a multicast boundary is defined on an interface, only the multicast traffic permitted by the filter ACL passes through the interface.

Step 3 From the **Interface** drop-down list, choose the interface for which you are configuring the multicast boundary filter ACL.

Step 4 From the **Standard Access List** drop-down list, choose the standard ACL you want to use, or click **Add** ( ) to create a new standard ACL. See Configure Standard ACL Objects, on page 469 for the procedure.

Step 5 Check the **Remove any Auto-RP group range announcement from the Auto-RP packets that are denied by the boundary** check box to filter Auto-RP messages from sources denied by the boundary ACL. If this check box is not checked, all Auto-RP messages are passed.

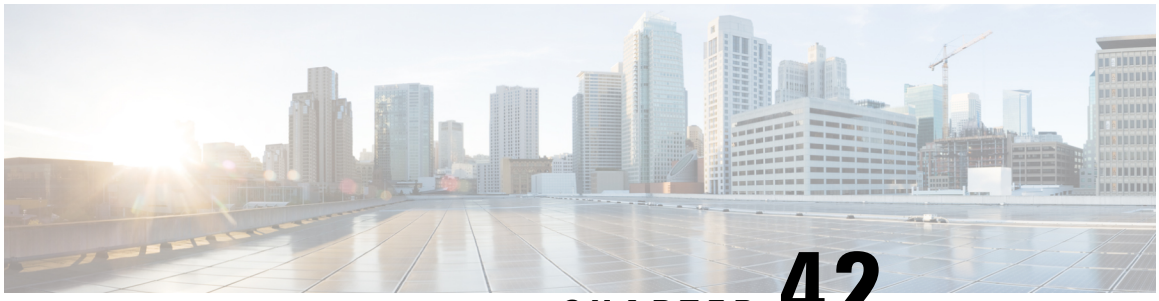
Step 6 Click **OK** to save the multicast boundary filter configuration.



PART XI

Firepower Threat Defense VPN

- VPN Overview for Firepower Threat Defense, on page 839
- Site-to-Site VPNs for Firepower Threat Defense, on page 851



CHAPTER 42

VPN Overview for Firepower Threat Defense

A virtual private network (VPN) connection establishes a secure tunnel between endpoints over a public network such as the Internet.

This chapter applies to Site-to-site VPNs on Firepower Threat Defense devices. It describes the Internet Protocol Security (IPsec), the Internet Security Association and Key Management Protocol (ISAKMP, or IKE) and SSL standards that are used to build site-to-site VPNs.

Site-to-site VPNs on 7000 and 8000 Series devices, referred to as Gateway VPNs or Firepower VPNs in the Firepower Management Center are described in Gateway VPNs, on page 1165.

- VPN Types, on page 839
- VPN Basics, on page 840
- VPN Packet Flow, on page 842
- VPN Licensing, on page 842
- How Secure Should a VPN Connection Be?, on page 842
- VPN Topology Options, on page 847

VPN Types

The Firepower Management Center supports the following types of VPN connections:

- Site-to-site VPNs on Firepower Threat Defense devices.

  A site-to-site VPN connects networks in different geographic locations. You can create site-to-site IPsec connections between managed devices, and between managed devices and other Cisco or third-party peers that comply with all relevant standards. These peers can have any mix of inside and outside IPv4 and IPv6 addresses. Site-to-site tunnels are built using the Internet Protocol Security (IPsec) protocol suite and IKEv1 or IKEv2. After the VPN connection is established, the hosts behind the local gateway can connect to the hosts behind the remote gateway through the secure VPN tunnel.

- Site-to-site VPNs on 7000 and 8000 Series devices.

  These site-to-site VPNs are referred to as Gateway VPNs or Firepower VPNs in the Firepower Management Center. See Gateway VPNs, on page 1165, for information on this type of VPN connection.

VPN Basics

Tunneling makes it possible to use a public TCP/IP network, such as the Internet, to create secure connections between remote users and private corporate networks. Each secure connection is called a tunnel.

IPsec-based VPN technologies use the Internet Security Association and Key Management Protocol (ISAKMP, or IKE) and IPsec tunneling standards to build and manage tunnels. ISAKMP and IPsec accomplish the following:

- Negotiate tunnel parameters.
- Establish tunnels.
- Authenticate users and data.
- Manage security keys.
- Encrypt and decrypt data.
- Manage data transfer across the tunnel.
- Manage data transfer inbound and outbound as a tunnel endpoint or router.

A device in a VPN functions as a bidirectional tunnel endpoint. It can receive plain packets from the private network, encapsulate them, create a tunnel, and send them to the other end of the tunnel where they are unencapsulated and sent to their final destination. It can also receive encapsulated packets from the public network, unencapsulate them, and send them to their final destination on the private network.

After the site-to-site VPN connection is established, the hosts behind the local gateway can connect to the hosts behind the remote gateway through the secure VPN tunnel. A connection consists of the IP addresses and hostnames of the two gateways, the subnets behind them, and the method the two gateways use to authenticate to each other.

Internet Key Exchange (IKE)

Internet Key Exchange (IKE) is a key management protocol that is used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and to automatically establish IPsec security associations (SAs).

The IKE negotiation comprises two phases. Phase 1 negotiates a security association between two IKE peers, which enables the peers to communicate securely in Phase 2. During Phase 2 negotiation, IKE establishes SAs for other applications, such as IPsec. Both phases use proposals when they negotiate a connection.

An IKE policy is a set of algorithms that two peers use to secure the IKE negotiation between them. IKE negotiation begins by each peer agreeing on a common (shared) IKE policy. This policy states which security parameters protect subsequent IKE negotiations. For IKE version 1 (IKEv1), IKE policies contain a single set of algorithms and a modulus group. Unlike IKEv1, in an IKEv2 policy, you can select multiple algorithms and modulus groups from which peers can choose during the Phase 1 negotiation. It is possible to create a single IKE policy, although you might want different policies to give higher priority to your most desired options. For site-to-site VPNs, you can create a single IKE policy.

To define an IKE policy, specify:

- A unique priority (1 to 65,543, with 1 the highest priority).
- An encryption method for the IKE negotiation, to protect the data and ensure privacy.

- A Hashed Message Authentication Codes (HMAC) method (called integrity algorithm in IKEv2) to ensure the identity of the sender, and to ensure that the message has not been modified in transit.
- For IKEv2, a separate pseudorandom function (PRF) used as the algorithm to derive keying material and hashing operations required for the IKEv2 tunnel encryption. The options are the same as those used for the hash algorithm.
- A Diffie-Hellman group to determine the strength of the encryption-key-determination algorithm. The device uses this algorithm to derive the encryption and hash keys.
- An authentication method, to ensure the identity of the peers.
- A limit to the time the device uses an encryption key before replacing it.

When IKE negotiation begins, the peer that starts the negotiation sends all of its policies to the remote peer, and the remote peer searches for a match with its own policies, in priority order. A match between IKE policies exists if they have the same encryption, hash (integrity and PRF for IKEv2), authentication, and Diffie-Hellman values, and an SA lifetime less than or equal to the lifetime in the policy sent. If the lifetimes are not identical, the shorter lifetime—From the remote peer policy—Applies. By default, the FMC deploys an IKEv1 policy at the lowest priority for all VPN endpoints to ensure a successful negotiation.

IPsec

IPsec is one of the most secure methods for setting up a VPN. IPsec provides data encryption at the IP packet level, offering a robust security solution that is standards-based. With IPsec, data is transmitted over a public network through tunnels. A tunnel is a secure, logical communication path between two peers. Traffic that enters an IPsec tunnel is secured by a combination of security protocols and algorithms.

An IPsec Proposal policy defines the settings required for IPsec tunnels. An IPsec proposal is a collection of one or more crypto-maps that are applied to the VPN interfaces on the devices. A crypto map combines all the components required to set up IPsec security associations, including:

- A proposal (or transform set) is a combination of security protocols and algorithms that secure traffic in an IPsec tunnel. During the IPsec security association (SA) negotiation, peers search for a proposal that is the same at both peers. When it is found, it is applied to create an SA that protects data flows in the access list for that crypto map, protecting the traffic in the VPN. There are separate IPsec proposals for IKEv1 and IKEv2. In IKEv1 proposals (or transform sets), for each parameter, you set one value. For IKEv2 proposals, you can configure multiple encryption and integration algorithms for a single proposal.
- A crypto map, combines all components required to set up IPsec security associations (SA), including IPsec rules, proposals, remote peers, and other parameters that are necessary to define an IPsec SA. When two peers try to establish an SA, they must each have at least one compatible crypto map entry.

  Dynamic crypto map policies are used in site-to-site VPNs when an unknown remote peer tries to start an IPsec security association with the local hub. The hub cannot be the initiator of the security association negotiation. Dynamic crypto-policies allow remote peers to exchange IPsec traffic with a local hub even if the hub does not know the remote peer's identity. A dynamic crypto map policy essentially creates a crypto map entry without all the parameters configured. The missing parameters are later dynamically configured (as the result of an IPsec negotiation) to match a remote peer's requirements.

  Dynamic crypto map policies apply only in a hub-and-spoke and full-mesh VPN topologies. In a point-to-point or full mesh VPN topology, you can apply only static crypto map policies. Emulate the use of dynamic crypto-maps in a point-to-point topology by creating a hub-and-spoke topology with two devices. Specify a dynamic IP address for the spoke and enable dynamic crypto-maps on this topology.



Note Simultaneous IKEv2 dynamic crypto map is not supported for the same interface for both remote access and site-to-site VPNs on Firepower Threat Defense (FTD).

VPN Packet Flow

On a Firepower Threat Defense device, by default no traffic is allowed to pass through access-control without explicit permission. VPN tunnel traffic as well, is not relayed to the endpoints until it has passed through Snort. Incoming tunnel packets are decrypted before being sent to the Snort process. Snort processes outgoing packets before encryption.

Access Control identifying the protected networks for each endpoint node of a VPN tunnel determines which traffic is allowed to pass through the Firepower Threat Defense device and reach the endpoints.

In addition, the system does not send tunnel traffic to the public source when the tunnel is down.

VPN Licensing

There is no specific licensing for enabling Firepower Threat Defense VPN, it is available by default.

The Firepower Management Center determines whether to allow or block the usage of strong crypto on a Firepower Threat Defense device based on attributes provided by the smart licensing server.

This is controlled by whether you selected the option to allow export-controlled functionality on the device when you registered with Cisco Smart License Manager. If you are using the evaluation license, or you did not enable export-controlled functionality, you cannot use strong encryption.

How Secure Should a VPN Connection Be?

Because a VPN tunnel typically traverses a public network, most likely the Internet, you need to encrypt the connection to protect the traffic. You define the encryption and other security techniques to apply using IKE polices and IPsec proposals.

If your device license allows you to apply strong encryption, there is a wide range of encryption and hash algorithms, and Diffie-Hellman groups, from which to choose. However, as a general rule, the stronger the encryption that you apply to the tunnel, the worse the system performance. Find a balance between security and performance that provides sufficient protection without compromising efficiency.

We cannot provide specific guidance on which options to choose. If you operate within a larger corporation or other organization, there might already be defined standards that you need to meet. If not, take the time to research the options.

The following topics explain the available options.

Deciding Which Encryption Algorithm to Use

When deciding which encryption algorithms to use for the IKE policy or IPsec proposal, your choice is limited to algorithms supported by the devices in the VPN.

For IKEv2, you can configure multiple encryption algorithms. The system orders the settings from the most secure to the least secure and negotiates with the peer using that order. For IKEv1, you can select a single option only.

For IPsec proposals, the algorithm is used by the Encapsulating Security Protocol (ESP), which provides authentication, encryption, and anti-replay services. ESP is IP protocol type 50. In IKEv1 IPsec proposals, the algorithm name is prefixed with ESP-.

If your device license qualifies for strong encryption, you can choose from the following encryption algorithms. If you are not qualified for strong encryption, you can select DES only.

- AES-GCM—(IKEv2 only.) Advanced Encryption Standard in Galois/Counter Mode is a block cipher mode of operation providing confidentiality and data-origin authentication, and provides greater security than AES. AES-GCM offers three different key strengths: 128-, 192-, and 256-bit keys. A longer key provides higher security but a reduction in performance. GCM is a mode of AES that is required to support NSA Suite B. NSA Suite B is a set of cryptographic algorithms that devices must support to meet federal standards for cryptographic strength. .
- AES-GMAC—(IKEv2 IPsec proposals only.) Advanced Encryption Standard Galois Message Authentication Code is a block cipher mode of operation providing only data-origin authentication. It is a variant of AES-GCM that allows data authentication without encrypting the data. AES-GMAC offers three different key strengths: 128-, 192-, and 256-bit keys.
- AES—Advanced Encryption Standard is a symmetric cipher algorithm that provides greater security than DES and is computationally more efficient than 3DES. AES offers three different key strengths: 128-, 192-, and 256-bit keys. A longer key provides higher security but a reduction in performance.
- 3DES—Triple DES, which encrypts three times using 56-bit keys, is more secure than DES because it processes each block of data three times with a different key. However, it uses more system resources and is slower than DES.
- DES—Data Encryption Standard, which encrypts using 56-bit keys, is a symmetric secret-key block algorithm. If your license account does not meet the requirements for export controls, this is your only option. It is faster than 3DES and uses less system resources, but it is also less secure. If you do not need strong data confidentiality, and if system resources or speed is a concern, choose DES.
- Null, ESP-Null—Do not use. A null encryption algorithm provides authentication without encryption. This is typically used for testing purposes only. However, it does not work at all on many platforms, including virtual and the Firepower 2100.

Deciding Which Hash Algorithms to Use

In IKE policies, the hash algorithm creates a message digest, which is used to ensure message integrity. In IKEv2, the hash algorithm is separated into two options, one for the integrity algorithm, and one for the pseudo-random function (PRF).

In IPsec proposals, the hash algorithm is used by the Encapsulating Security Protocol (ESP) for authentication. In IKEv2 IPsec Proposals, this is called the integrity hash. In IKEv1 IPsec proposals, the algorithm name is prefixed with ESP-, and there is also an -HMAC suffix (which stands for “hash method authentication code”).

For IKEv2, you can configure multiple hash algorithms. The system orders the settings from the most secure to the least secure and negotiates with the peer using that order. For IKEv1, you can select a single option only.

You can choose from the following hash algorithms.

- SHA (Secure Hash Algorithm)—Standard SHA (SHA1) produces a 160-bit digest. SHA is more resistant to brute-force attacks than MD5. However, it is also more resource intensive than MD5. For implementations that require the highest level of security, use the SHA hash algorithm.

  The following SHA-2 options, which are even more secure, are available for IKEv2 configurations. Choose one of these if you want to implement the NSA Suite B cryptography specification.

  - SHA256—Specifies the Secure Hash Algorithm SHA 2 with the 256-bit digest.
  - SHA384—Specifies the Secure Hash Algorithm SHA 2 with the 384-bit digest.
  - SHA512—Specifies the Secure Hash Algorithm SHA 2 with the 512-bit digest.

- MD5 (Message Digest 5)—Produces a 128-bit digest. MD5 uses less processing time for an overall faster performance than SHA, but it is considered to be weaker than SHA.
- Null or None (NULL, ESP-NONE)—(IPsec Proposals only.) A null Hash Algorithm; this is typically used for testing purposes only. However, you should choose the null integrity algorithm if you select one of the AES-GCM/GMAC options as the encryption algorithm. Even if you choose a non-null option, the integrity hash is ignored for these encryption standards.

Deciding Which Diffie-Hellman Modulus Group to Use

You can use the following Diffie-Hellman key derivation algorithms to generate IPsec security association (SA) keys. Each group has a different size modulus. A larger modulus provides higher security, but requires more processing time. You must have a matching modulus group on both peers.

If you select AES encryption, to support the large key sizes required by AES, you should use Diffie-Hellman (DH) Group 5 or higher. IKEv1 policies do not support all of the groups listed below.

To implement the NSA Suite B cryptography specification, use IKEv2 and select one of the elliptic curve Diffie-Hellman (ECDH) options: 19, 20, or 21. Elliptic curve options and groups that use 2048-bit modulus are less exposed to attacks such as Logjam.

For IKEv2, you can configure multiple groups. The system orders the settings from the most secure to the least secure and negotiates with the peer using that order. For IKEv1, you can select a single option only.

- 1—Diffie-Hellman Group 1: 768-bit modulus. DH group 1 is considered insecure, please do not use it.
- 2—Diffie-Hellman Group 2: 1024-bit modular exponential (MODP) group. This option is no longer considered good protection.
- 5—Diffie-Hellman Group 5: 1536-bit MODP group. Formerly considered good protection for 128-bit keys, this option is no longer considered good protection.
- 14—Diffie-Hellman Group 14: 2048-bit modular exponential (MODP) group. Considered good protection for 192-bit keys.
- 19—Diffie-Hellman Group 19: National Institute of Standards and Technology (NIST) 256-bit elliptic curve modulo a prime (ECP) group.

- 20—Diffie-Hellman Group 20: NIST 384-bit ECP group.
- 21—Diffie-Hellman Group 21: NIST 521-bit ECP group.
- 24—Diffie-Hellman Group 24: 2048-bit MODP group with 256-bit prime order subgroup. This option is no longer recommended.

Deciding Which Authentication Method to Use

Preshared keys and digital certificates are the methods of authentication available for VPNs.

Site-to-site, IKEv1 and IKEv2 VPN connections can use both options.

Preshared keys allow for a secret key to be shared between two peers and used by IKE during the authentication phase. The same shared key must be configured at each peer or the IKE SA cannot be established.

Digital certificates use RSA key pairs to sign and encrypt IKE key management messages. Certificates provide non-repudiation of communication between two peers, meaning that it can be proved that the communication actually took place. When using this authentication method, you need a Public Key Infrastructure (PKI) defined where peers can obtain digital certificates from a Certification Authority (CA). CAs manage certificate requests and issue certificates to participating network devicesproviding centralized key management for all of the participating devices.

Preshared keys do not scale well, using a CA improves the manageability and scalability of your IPsec network. With a CA, you do not need to configure keys between all encrypting devices. Instead, each participating device is registered with the CA, and requests a certificate from the CA. Each device that has its own certificate and the public key of the CA can authenticate every other device within a given CA's domain.

Pre-shared Keys

Preshared keys allow for a secret key to be shared between two peers. The key is used by IKE in the authentication phase. The same shared key must be configured on each peer, or the IKE SA cannot be established.

To configure the pre-shared keys, choose whether you will use a manual or automatically generated key, and then speicify the key in the IKEv1/IKEv2 options. Then, when your configuration is deployed, the key is configured on all devices in the topology.

PKI Infrastructure and Digital Certificates

Public Key Infrastructure

A PKI provides centralized key management for participating network devices. It is a defined set of policies, procedures, and roles that support *public key cryptography* by generating, verifying, and revoking *public key certificates* commonly known as *digital certificates*.

In public key cryptography, each endpoint of a connection has a key pair consisting of both a public and a private key. The key pairs are used by the VPN endpoints to sign and encrypt messages. The keys act as complements, and anything encrypted with one of the keys can be decrypted with the other, securing the data flowing over the connection.

Generate a general purpose RSA or ECDSA key pair, used for both signing and encryption, or you generate separate key pairs for each purpose. Separate signing and encryption keys help to reduce exposure of the keys. SSL uses a key for encryption but not signing, however, IKE uses a key for signing but not encryption. By using separate keys for each, exposure of the keys is minimized.

Digital Certificates or Identify Certificates

When you use Digital Certificates as the authentication method for VPN connections, peers are configured to obtain digital certificates from a Certificate Authority (CA). CAs are trusted authorities that "sign" certificates to verify their authenticity, thereby guaranteeing the identity of the device or user.

CA servers manage public CA certificate requests and issue certificates to participating network devices as part of a Public Key Infrastructure (PKI), this activity is called Certificate Enrollment. These digital certificates, also called identity certificates contain:

- The digital identification of the owner for authentication, such as name, serial number, company, department, or IP address.
- A public key needed to send and receive encrypted data to the certificate owner.
- The secure digital signature of a CA.

Certificates also provide non-repudiation of communication between two peers, meaning that it they prove that the communication actually took place.

Certificate Enrollment

Using a PKI improves the manageability and scalability of your VPN since you do not have to configure pre-shared keys between all the encrypting devices. Instead, you individually *enroll* each participating device with a CA server, which is explicitly trusted to validate identities and create an identity certificate for the device. When this has been accomplished, each participating peer sends their identity certificate to the other peer to validate their identities and establish encrypted sessions with the public keys contained in the certificates. See Certificate Enrollment Objects, on page 456for details on enrolling Firepower Threat Defense devices.

Certificate Authority Certificates

In order to validate a peer's certificate, each participating device must retrieve the CA's certificate from the server. A CA certificate is used to sign other certificates. It is self-signed and called a root certificate. This certificate contains the public key of the CA, used to decrypt and validate the CA's digital signature and the contents of the received peer's certificate. The CA certificate may be obtained by:

- Using the Simple Certificate Enrollment Protocol (SCEP) to retrieve the CA's certificate from the CA server
- Manually copying the CA's certificate from another participating device

Trustpoints

Once enrollment is complete, a trustpoint is created on the managed device. It is the object representation of a CA and associated certificates. A trustpoint includes the identity of the CA, CA-specific parameters, and an association with a single enrolled identity certificate.

PKCS#12 File

A PKCS#12, or PFX, file holds the server certificate, any intermediate certificates, and the private key in one encrypted file. This type of file may be imported directly into a device to create a trustpoint.

Revocation Checking

A CA may also revoke certificates for peers that no longer participate in you network. Revoked certificates are either managed by an Online Certificate Status Protocol (OCSP) server or are listed in a certificate revocation list (CRL) stored on an LDAP server. A peer may check these before accepting a certificate from another peer.

VPN Topology Options

When you create a new VPN topology you must, at minimum, give it a unique name, specify a topology type, and select the IKE version. You can select from three types of topologies, each containing a group of VPN tunnels:

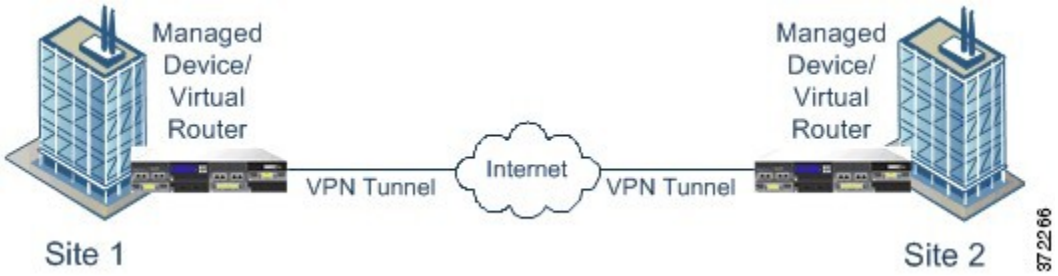
- Point-to-point (PTP) topologies establish a VPN tunnel between two endpoints.
- Hub and Spoke topologies establish a group of VPN tunnels connecting a hub endpoint to a group of spoke endpoints.
- Full Mesh topologies establish a group of VPN tunnels among a set of endpoints.

Define a pre-shared key for VPN authentication manually or automatically, there is no default key. When choosing automatic, the Firepower Management Center generates a pre-shared key and assigns it to all the nodes in the topology.

Point-to-Point VPN Topology

In a point-to-point VPN topology, two endpoints communicate directly with each other. You configure the two endpoints as peer devices, and either device can start the secured connection.

The following diagram displays a typical point-to-point VPN topology.

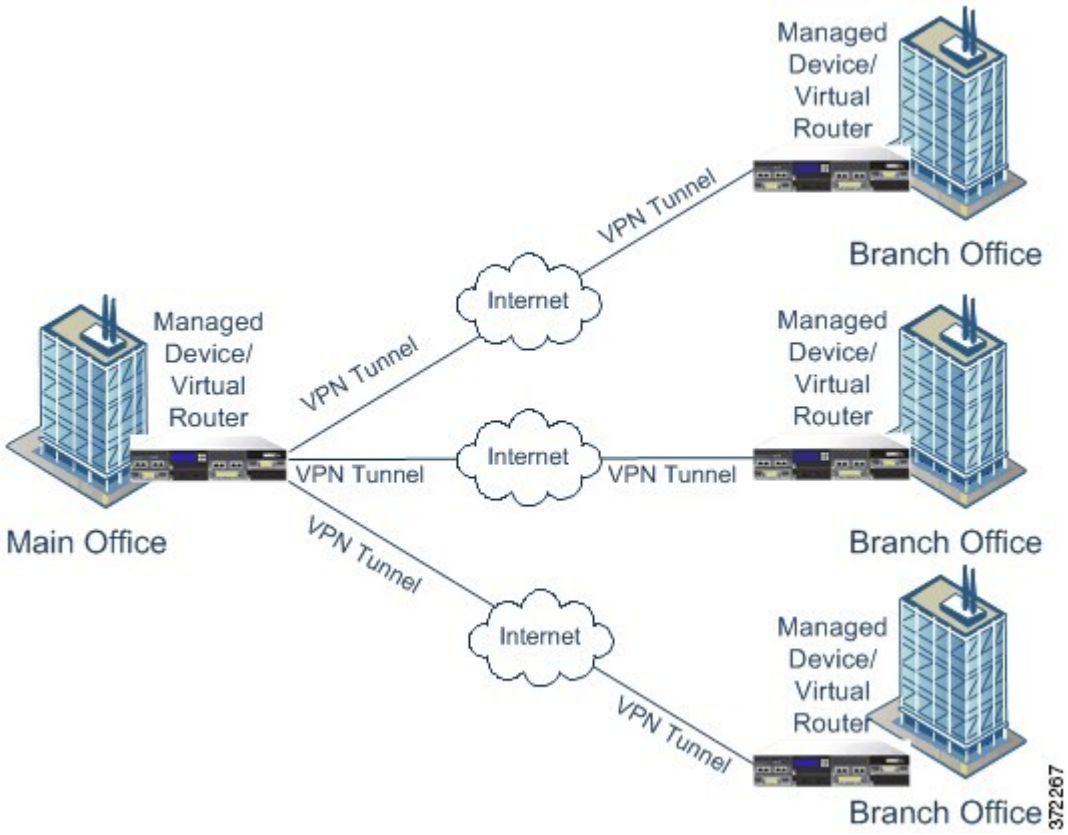


Hub and Spoke VPN Topology

In a Hub and Spoke VPN topology, a central endpoint (hub node) connects with multiple remote endpoints (spoke nodes). Each connection between the hub node and an individual spoke endpoint is a separate VPN tunnel. The hosts behind any of the spoke nodes can communicate with each other through the hub node.

The Hub and Spoke topology commonly represent a VPN that connects an organization's main and branch office locations using secure connections over the Internet or other third-party network. These deployments provide all employees with controlled access to the organization's network. Typically, the hub node is located at the main office. Spoke nodes are located at branch offices and start most of the traffic.

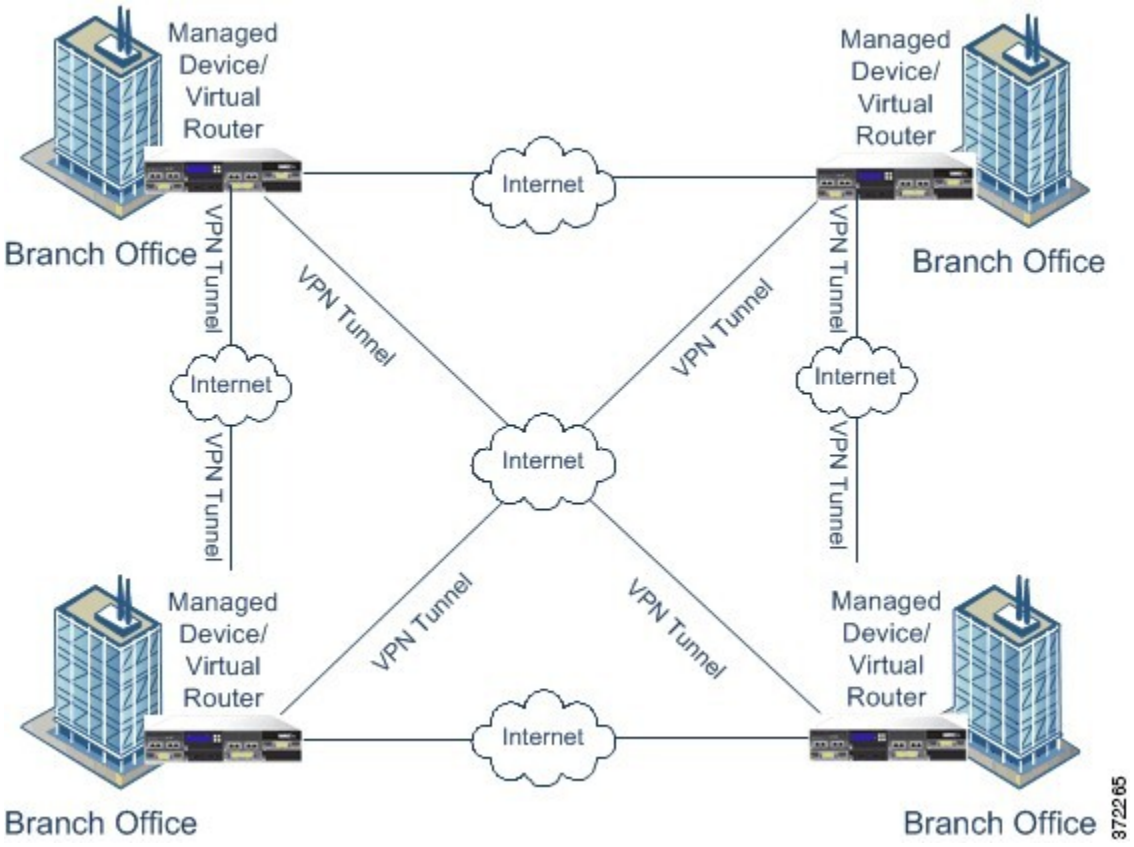
The following diagram displays a typical Hub and Spoke VPN topology.



Full Mesh VPN Topology

In a Full Mesh VPN topology, all endpoints can communicate with every other endpoint by an individual VPN tunnel. This topology offers redundancy so that when one endpoint fails, the remaining endpoints can still communicate with each other. It commonly represents a VPN that connects a group of decentralized branch office locations. The number of VPN-enabled managed devices you deploy in this configuration depends on the level of redundancy you require.

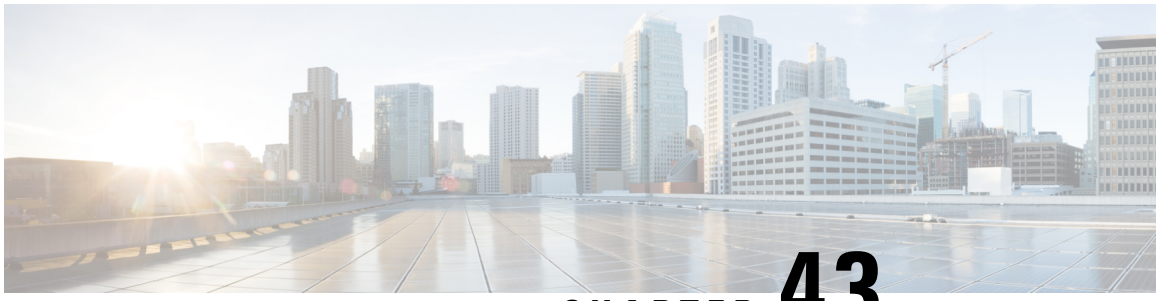
The following diagram displays a typical Full Mesh VPN topology.



Implicit Topologies

In addition to the three main VPN topologies, other more complex topologies can be created as combinations of these topologies. They include:

- Partial mesh—A network in which some devices are organized in a full mesh topology, and other devices form either a hub-and-spoke or a point-to-point connection to some of the fully meshed devices. A partial mesh does not provide the level of redundancy of a full mesh topology, but it is less expensive to implement. Partial mesh topologies are used in peripheral networks that connect to a fully meshed backbone.
- Tiered hub-and-spoke—A network of hub-and-spoke topologies in which a device can behave as a hub in one or more topologies and a spoke in other topologies. Traffic is permitted from spoke groups to their most immediate hub.
- Joined hub-and-spoke—A combination of two topologies (hub-and-spoke, point-to-point, or full mesh) that connect to form a point-to-point tunnel. For example, a joined hub-and-spoke topology could comprise two hub-and-spoke topologies, with the hubs acting as peer devices in a point-to-point topology.



CHAPTER 43

Site-to-Site VPNs for Firepower Threat Defense

- About Firepower Threat Defense Site-to-site VPNs, on page 851
- Requirements and Prerequisites for Site-to-Site VPN, on page 853
- Managing Firepower Threat Defense Site-to-site VPNs, on page 853
- Configuring Firepower Threat Defense Site-to-site VPNs, on page 854
- Monitoring Firepower Threat Defense VPNs, on page 863

About Firepower Threat Defense Site-to-site VPNs

Firepower Threat Defense site-to-site VPN supports the following features:

- Both IPsec IKEv1 & IKEv2 protocols are supported.
- Certificates and automatic or manual preshared keys for authentication.
- IPv4 & IPv6. All combinations of inside and outside are supported.
- Static and Dynamic Interfaces.
- Support for both Firepower Management Center and Firepower Threat Defense HA environments.
- VPN alerts when the tunnel goes down.
- Tunnel statistics available using the Firepower Threat Defense Unified CLI.

VPN Topology

To create a new site-to-site VPN topology you must, at minimum, give it a unique name, specify a topology type, choose the IKE version that is used for IPsec IKEv1 or IKEv2, or both. Also, determine your authentication method. Once configured, you deploy the topology to Firepower Threat Defense devices. The Firepower Management Center configures site-to-site VPNs on Firepower Threat Defense devices only.

You can select from three types of topologies, containing one or more VPN tunnels:

- Point-to-point (PTP) deployments establish a VPN tunnel between two endpoints.
- Hub and Spoke deployments establish a group of VPN tunnels connecting a hub endpoint to a group of spoke nodes.
- Full Mesh deployments establish a group of VPN tunnels among a set of endpoints.

IPsec and IKE

In the Firepower Management Center, site-to-site VPNs are configured based on IKE policies and IPsec proposals that are assigned to VPN topologies. Policies and proposals are sets of parameters that define the characteristics of a site-to-site VPN, such as the security protocols and algorithms that are used to secure traffic in an IPsec tunnel. Several policy types may be required to define a full configuration image that can be assigned to a VPN topology.

Authentication

For authentication of VPN connections, configure a preshared key in the topology, or a trustpoint on each device. Preshared keys allow for a secret key, used during the IKE authentication phase, to be shared between two peers. A trustpoint includes the identity of the CA, CA-specific parameters, and an association with a single enrolled identity certificate.

Extranet Devices

Each topology type can include Extranet devices, devices that you do not manage in Firepower Management Center. These include:

- Cisco devices that Firepower Management Center supports, but for which your organization is not responsible. Such as spokes in networks managed by other organizations within your company, or a connection to a service provider or partner's network.
- Non-Cisco devices. You cannot use Firepower Management Center to create and deploy configurations to non-Cisco devices.

Add non-Cisco devices, or Cisco devices not managed by the Firepower Management Center, to a VPN topology as "Extranet" devices. Also specify the IP address of each remote device.

Firepower Threat Defense Site-to-site VPN Guidelines and Limitations

- A VPN connection can only be made across domains by using an extranet peer for the endpoint not in the current domain.
- A VPN topology cannot be moved between domains.
- Network objects with a 'range' option are not supported in VPN
- Firepower Threat Defense VPNs are only be backed up using the Firepower Management backup.
- The Firepower Threat Defense VPNs do not currently support PDF export and policy comparison.
- There is no per-tunnel or per-device edit option for Firepower Threat Defense VPNs, only the whole topology can be edited.
- Firepower Threat Defense VPNs are not supported in clustered environment.
- Tunnel status is not updated in realtime, but at an interval of 5 minutes in the Firepower Management Center.

Requirements and Prerequisites for Site-to-Site VPN

Model Support

FTD

Supported Domains

Leaf

User Roles

Admin

Managing Firepower Threat Defense Site-to-site VPNs

Procedure

Step 1 For certificate authentication for your VPNs, you must prepare the devices by allocating trustpoints as described in Firepower Threat Defense Certificate-Based Authentication, on page 481.

Step 2 Select **Devices > VPN > Site To Site** to manage your Firepower Threat Defense Site-to-site VPN configurations and deployments. Choose from the following:

- Add—To create a new VPN topology, click **Add** ( ) **Add VPN** > **Firepower Threat Defense Device**, and continue as instructed in Configuring Firepower Threat Defense Site-to-site VPNs, on page 854:

  Note VPNs topologies can be created only on leaf domains.

- Edit—To modify the settings of an existing VPN topology, click **Edit** ( ). Modifying is similar to configuring, continue as instructed above.

  Note You cannot edit the topology type after you initially save it. To change the topology type, delete the topology and create a new one.

  Two users should **not** edit the same topology simultaneously; however, the web interface does not prevent simultaneous editing.

- Delete—To delete a VPN deployment, click **Delete** ( ).
- View VPN status—This status applies to Firepower VPNs ONLY. Currently, no status is displayed for Firepower Threat Defense VPNs. To determine the status of the Firepower Threat Defense VPNs, see Monitoring Firepower Threat Defense VPNs, on page 863.
- Deploy—Click **Deploy**; see Deploy Configuration Changes, on page 346.

Note Some VPN settings are validated only during deployment. Be sure to verify that your deployment was successful.

Configuring Firepower Threat Defense Site-to-site VPNs

Procedure

Step 1 Choose **Devices > VPN > Site To Site**.Then **Add VPN > Firepower Threat Defense Device**, or edit a listed VPN Topology. .

Step 2 Enter a unique **Topology Name**. We recommend naming your topology to indicate that it is a Firepower Threat Defense VPN, and its topology type.

Step 3 Choose the **Network Topology** for this VPN.

Step 4 Choose the IKE versions to use during IKE negotiations. **IKEv1** or **IKEv2**.

Default is IKEv2. Select either or both options as appropriate; select IKEv1 if any device in the topology does not support IKEv2. You can also configure backup peer for point-to-point extranet VPNs. For more information, see Firepower Threat Defense VPN Endpoint Options, on page 855.

Step 5 Required: Add Endpoints for this VPN deployment by clicking **Add** () for each node in the topology.

Configure each endpoint field as described in Firepower Threat Defense VPN Endpoint Options, on page 855.

- For Point to point, configure **Node A** and **Node B**.
- For Hub and Spoke, configure a **Hub Node** and **Spoke Nodes**
- For Full Mesh, configure multiple **Nodes**

Step 6 (Optional) Specify non-default IKE options for this deployment as described in Firepower Threat Defense VPN IKE Options, on page 856

Step 7 (Optional) Specify non-default IPsec options for this deployment as described in Firepower Threat Defense VPN IPsec Options, on page 858

Step 8 (Optional) Specify non-default Advanced options for this deployment as described in Firepower Threat Defense Advanced Site-to-site VPN Deployment Options, on page 860.

Step 9 Click **Save**.
The endpoints are added to your configuration.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Note Some VPN settings are validated only during deployment. Be sure to verify that your deployment was successful.

Firepower Threat Defense VPN Endpoint Options

Navigation Path

Devices > VPN > Site To Site. Then **Add VPN > Firepower Threat Defense Device**, or edit a listed VPN Topology. Open the **Endpoint** tab.

Fields

Device

Choose an endpoint node for your deployment:

- A Firepower Threat Defense device managed by this Firepower Management Center.
- A Firepower Threat Defense high availability container managed by this Firepower Management Center.
- An **Extranet** device, any device (Cisco or third-party) not managed by this Firepower Management Center.

Device Name

For Extranet devices only, provide a name for this device. We recommend naming it such that it is identifiable as an un-managed device.

Interface

If you chose a managed device as your endpoint, choose an interface on that managed device.

IP Address

- If you choose an extranet device, a device **not** managed by the Firepower Management Center, specify an IP address for the endpoint.
- If you chose a managed device as an endpoint, choose a single IPv4 address or multiple IPv6 addresses from the drop-down list (these are the addresses already assigned to this interface on this managed device).
- All endpoints in a topology must have the same IP addressing scheme. IPv4 tunnels can carry IPv6 traffic and vice-versa. The Protected Networks define which addressing scheme the tunneled traffic will use.
- If the managed device is a high-availability container, choose from a list of interfaces.

This IP is Private

Check the check box if the endpoint resides behind a firewall with network address translation (NAT).

Note Use this option only when the peer is managed by the same Firepower Management Center and do not use this option if peer is from extranet.

Public IP address

If you checked the **This IP is Private** check box, specify a public IP address for the firewall. If the endpoint is a responder, specify this value.

Connection Type

Specify the allowed negotiation as bidirectional, answer-only, or originate-only. Supported combinations for the connection type are:

Table 68: Connection Type Supported Combinations

Remote Node	Central Node
Originate-Only	Answer-Only
Bi-Directional	Answer-Only
Bi-Directional	Bi-Directional

Protected Networks

Defines a list of networks protected by this VPN endpoint. Click **Add** ( ) to select from available Network Objects or add Network Objects inline. See Creating Network Objects, on page 403. Access Control Lists will be generated from the choices made here.

- **Subnet/IP Address (Network)**—VPN endpoints cannot have the same IP address and protected networks in a VPN endpoint pair cannot overlap. If a list of protected networks for an endpoint contains one or more IPv4 or IPv6 entries, the other endpoint's protected network must have at least one entry of the same type (that is, IPv4 or IPv6). If it does not, then the other endpoint's IP address must be of the same type and must not overlap with the entries in the protected network. (Use /32 CIDR address blocks for IPv4 and /128 CIDR address blocks for IPv6.) If both of these checks fail, the endpoint pair is invalid.

Note **Reverse Route Injection is enabled** by default in Firepower Management Center.

-

Firepower Threat Defense VPN IKE Options

For the versions of IKE you have chosen for this topology, specify the **IKEv1/IKEv2 Settings**.

Note Settings in this dialog apply to the entire topology, all tunnels, and all managed devices.

Navigation Path

Devices > VPN > Site To Site. Then **Add VPN > Firepower Threat Defense Device**, or edit a listed VPN Topology. Open the **IKE** tab.

Fields

Policy

Choose a predefined IKEv1 or IKEv2 policy object or create a new one to use. For details, see Firepower Threat Defense IKE Policies, on page 473

Authentication Type

Site-to-site VPN supports two authentication methods, pre-shared key and certificate. For an explanation of the two methods, see Deciding Which Authentication Method to Use, on page 845.

Note In a VPN topology that supports IKEv1, the **Authentication Method** specified in the chosen IKEv1 Policy object becomes the default in the IKEv1 **Authentication Type** setting. These values must match, otherwise, your configuration will error.

- **Pre-shared Automatic Key**—The Management Center automatically defines the pre-shared key that is used for this VPN. Specify the **Pre-shared Key Length**, the number of characters in the key, 1-27.
- **Pre-shared Manual Key**—Manually assign the pre-shared key that is used for this VPN. Specify the **Key** and then re-enter it in **Confirm Key** to confirm.
- **Certificate**—When you use certificates as the authentication method for VPN connections, peers obtain digital certificates from a CA server in your PKI infrastructure, and trade them to authenticate each other.

  In the **Certificate** field, select a pre-configured certificate enrollment object. This enrollment object is used to generate a trustpoint with the same name on the managed device. The certificate enrollment object should be associated with and installed on the device, post which the enrollment process is complete, and then a trustpoint is created.

  A trustpoint is a representation of a CA or identity pair. A trustpoint includes the identity of the CA, CA-specific configuration parameters, and an association with one enrolled identity certificate.

  Before you select this option, note the following:

  - Ensure you have enrolled a certificate enrollment object on all the endpoints in the topology—A certificate enrollment object contains the Certification Authority (CA) server information and enrollment parameters that are required for creating Certificate Signing Requests (CSRs) and obtaining Identity Certificates from the specified CA. Certificate Enrollment Objects are used to enroll your managed devices into your PKI infrastructure, and create trustpoints (CA objects) on devices that support VPN connections. For instructions on creating a certificate enrollment object, see Adding Certificate Enrollment Objects, on page 458, and for instructions on enrolling the object on the endpoints see one of the following as applicable:
    - Installing a Certificate Using Self-Signed Enrollment , on page 483
    - Installing a Certificate Using SCEP Enrollment, on page 484
    - Installing a Certificate Using Manual Enrollment, on page 485

- Installing a Certificate by Importing a PKCS12 File, on page 486

Note For a site-to-site VPN topology, ensure that the same certificate enrollment object is enrolled in all the endpoints in the topology. For further details, see the table below.

- Refer the following table to understand the enrollment requirement for different scenarios. Some of the scenarios require you to override the certificate enrollment object for specific devices. See Managing Object Overrides, on page 401 to understand how to override objects.

Certificate Enrollment Types	Device identity certificate for all endpoints is from the same CA		Device identity certificate for all endpoints is from different CAs
	Device-specific parameters are NOT specified in the certificate enrollment object	Device-specific parameters are specified in the certificate enrollment object	
Manual	No override required	Override required	Override required
SCEP	No override required	Override required	Override required
PKCS	Override required	Override required	Override required
Self-signed	Not applicable	Not applicable	Not applicable

- Understand the VPN certificate limitations mentioned in Firepower Threat Defense VPN Certificate Guidelines and Limitations, on page 482.

Note If you use a Windows Certificate Authority (CA), the default Application Policies extension is **IP security IKE intermediate**. If you are using this default setting, you must select the **Ignore IPsec Key Usage** option in the Advanced Settings section on the **Key** tab in the PKI Certificate Enrollment dialog box for the object you select. Otherwise, the endpoints cannot complete the site-to-site VPN connection.

Firepower Threat Defense VPN IPsec Options

Note Settings in this dialog apply to the entire topology, all tunnels, and all managed devices.

Crypto-Map Type

A crypto map combines all the components required to set up IPsec security associations (SA). When two peers try to establish an SA, they must each have at least one compatible crypto map entry. The

proposals defined in the crypto map entry are used in the IPsec security negotiation to protect the data flows specified by that crypto map's IPsec rules. Choose static or dynamic for this deployment's crypto-map:

- **Static**—Use a static crypto map in a point-to-point or full mesh VPN topology.
- **Dynamic**—Dynamic crypto-maps essentially create a crypto map entry without all the parameters configured. The missing parameters are later dynamically configured (as the result of an IPsec negotiation) to match a remote peer's requirements.

  Dynamic crypto map policies apply only in a hub-and-spoke VPN configuration. In a point-to-point or full mesh VPN topology, you can apply only static crypto map policies. Emulate the use of dynamic crypto-maps in a point-to-point topology by creating a hub-and-spoke topology with two devices. Specify a dynamic IP address for the spoke, and enable dynamic crypto map on this topology.

IKEv2 Mode

For IPsec IKEv2 only, specify the encapsulation mode for applying ESP encryption and authentication to the tunnel. This determines what part of the original IP packet has ESP applied.

- **Tunnel mode**—(default) Encapsulation mode is set to tunnel mode. Tunnel mode applies ESP encryption and authentication to the entire original IP packet (IP header and data), hiding the ultimate source and destination addresses and becoming the payload in a new IP packet.

  The major advantage of tunnel mode is that the end systems do not need to be modified to receive the benefits of IPsec. This mode allows a network device, such as a router, to act as an IPsec proxy. That is, the router performs encryption on behalf of the hosts. The source router encrypts packets and forwards them along the IPsec tunnel. The destination router decrypts the original IP datagram and forwards it onto the destination system. Tunnel mode also protects against traffic analysis; with tunnel mode, an attacker can only determine the tunnel endpoints and not the true source and destination of the tunneled packets, even if they are the same as the tunnel endpoints.

- **Transport preferred**— Encapsulation mode is set to transport mode with an option to fallback to tunnel mode if the peer does not support it. In Transport mode only the IP payload is encrypted, and the original IP headers are left intact. Therefore, the admin must select a protected network that matches the VPN interface IP address.

  This mode has the advantages of adding only a few bytes to each packet and allowing devices on the public network to see the final source and destination of the packet. With transport mode, you can enable special processing (for example, QoS) on the intermediate network based on the information in the IP header. However, the Layer 4 header is encrypted, which limits examination of the packet.

- **Transport required**— Encapsulation mode is set to transport mode only, falling back to tunnel mode is not allowed. If the endpoints cannot successfully negotiate transport mode, due to one endpoint not supporting it, the VPN connection is not made.

Proposals

Click **Edit** ( ) to specify the proposals for your chosen IKEv1 or IKEv2 method. Select from the available **IKEv1 IPsec Proposals** or **IKEv2 IPsec Proposals** objects, or create and then select a new one. See Configure IKEv1 IPsec Proposal Objects, on page 477 and Configure IKEv2 IPsec Proposal Objects, on page 477 for details.

Enable Security Association (SA) Strength Enforcement

Enabling this option ensures that the encryption algorithm used by the child IPsec SA is not stronger (in terms of the number of bits in the key) than the parent IKE SA.

Enable Reverse Route Injection

Reverse Route Injection (RRI) enables static routes to be automatically inserted into the routing process for those networks and hosts protected by a remote tunnel endpoint.

Enable Perfect Forward Secrecy

Whether to use Perfect Forward Secrecy (PFS) to generate and use a unique session key for each encrypted exchange. The unique session key protects the exchange from subsequent decryption, even if the entire exchange was recorded and the attacker has obtained the preshared or private keys used by the endpoint devices. If you select this option, also select the Diffie-Hellman key derivation algorithm to use when generating the PFS session key in the Modulus Group list.

Modulus Group

The Diffie-Hellman group to use for deriving a shared secret between the two IPsec peers without transmitting it to each other. A larger modulus provides higher security but requires more processing time. The two peers must have a matching modulus group. For a full explanation of the options, see Deciding Which Diffie-Hellman Modulus Group to Use, on page 844.

Lifetime (seconds)

The number of seconds a security association exists before expiring. The default is 28,800 seconds.

Lifetime (kbytes)

The volume of traffic (in kilobytes) that can pass between IPsec peers using a given security association before it expires. The default is 4,608,000 kilobytes. Infinite data is not allowed.

ESPv3 Settings

Validate incoming ICMP error messages

Choose whether to validate ICMP error messages received through an IPsec tunnel and destined for an interior host on the private network.

Enable 'Do Not Fragment' Policy

Define how the IPsec subsystem handles large packets that have the do-not-fragment (DF) bit set in the IP header.

Policy

- Copy DF bit—Maintains the DF bit.
- Clear DF bit—Ignores the DF bit.
- Set DF bit—Sets and uses the DF bit.

Enable Traffic Flow Confidentiality (TFC) Packets

Enable dummy TFC packets that mask the traffic profile which traverses the tunnel. Use the **Burst**, **Payload Size**, and **Timeout** parameters to generate random length packets at random intervals across the specified SA.

Firepower Threat Defense Advanced Site-to-site VPN Deployment Options

The following sections describes the advanced options you can specify in your S2S VPN deployment. These settings apply to the entire topology, all tunnels, and all managed devices.

Firepower Threat Defense VPN Advanced IKE Options

Advanced > IKE > ISAKAMP Settings

IKE Keepalive

Enable or disables IKE Keepalives. Or set to EnableInfinite specifying that the device never starts keepalive monitoring itself.

Threshold
Specifies the IKE keep alive confidence interval. This is the number of seconds allowing a peer to idle before beginning keepalive monitoring. The minimum and default is 10 seconds; the maximum is 3600 seconds.

Retry Interval
Specifies number of seconds to wait between IKE keep alive retries. The default is 2 seconds, the maximum is 10 seconds.

Identity Sent to Peers:
Choose the identity that the peers will use to identify themselves during IKE negotiations:

- **autoOrDN**(default)—Determines IKE negotiation by connection type: IP address for preshared key, or Cert DN for certificate authentication (not supported).
- **ipAddress**—Uses the IP addresses of the hosts exchanging ISAKMP identity information.
- **hostname**—Uses the fully qualified domain name of the hosts exchanging ISAKMP identity information. This name comprises the hostname and the domain name.



Note Enable or disable this option for all your VPN connections.

Enable Aggressive Mode
Available only in a hub-and-spoke VPN topology. Select this negotiation method for exchanging key information if the IP address is not known and DNS resolution might not be available on the devices. Negotiation is based on hostname and domain name.

Advanced > IKE > IVEv2 Security Association (SA) Settings

More session controls are available for IKE v2 that limit the number of open SAs. By default, there is no limit to the number of open SAs:

Cookie Challenge
Whether to send cookie challenges to peer devices in response to SA initiate packets, which can help thwart denial of service (DoS) attacks. The default is to use cookie challenges when 50% of the available SAs are in negotiation. Select one of these options:

- Custom:
- Never (default)
- Always

Threshold to Challenge Incoming Cookies
The percentage of the total allowed SAs that are in-negotiation. This triggers cookie challenges for any future SA negotiations. The range is zero to 100%.

Number of SAs Allowed in Negotiation
Limits the maximum number of SAs that can be in negotiation at any time. If used with Cookie Challenge, configure the cookie challenge threshold lower than this limit for an effective cross-check.

Maximum number of SAs Allowed
Limits the number of allowed IKEv2 connections. Default is unlimited.

Enable Notification on Tunnel Disconnect
Allows an administrator to enable or disable the sending of an IKE notification to the peer when an inbound packet that is received on an SA does not match the traffic selectors for that SA. Sending this notification is disabled by default.

Do not allow device reboot until all sessions are terminated
Check to enable waiting for all active sessions to voluntarily terminate before the system reboots. This is disabled by default.

Firepower Threat Defense VPN Advanced IPsec Options

Advanced > IPsec > IPsec Settings

Enable Fragmentation Before Encryption
This option lets traffic travel across NAT devices that do not support IP fragmentation. It does not impede the operation of NAT devices that do support IP fragmentation.

Path Maximum Transmission Unit Aging
Check to enable PMTU (Path Maximum Transmission Unit) Aging, the interval to Reset PMTU of an SA (Security Association)

Value Reset Interval
Enter the number of minutes at which the PMTU value of an SA (Security Association) is reset to its original value. Valid range is 10 to 30 minutes, default is unlimited.

Firepower Threat Defense Advanced Site-to-site VPN Tunnel Options

Navigation Path

Devices > VPN > Site To Site, then select **Add VPN > Firepower Threat Defense Device**, or edit a listed VPN Topology. Open the **Advanced** tab, and select **Tunnel** in the navigation pane.

Tunnel Options

Only available for Hub and Spoke, and Full Mesh topologies. This section will not display for Point to Point configurations.

- **Enable Spoke to Spoke Connectivity through Hub**—Disabled by default. Choosing this field enables the devices on each end of the spokes to extend their connection through the hub node to the other device.

NAT Settings

- **Keepalive Messages Traversal** —Elect whether to enable NAT keepalive message traversal. NAT traversal keepalive is used for the transmission of keepalive messages when there is a device (middle device) located between a VPN-connected hub and spoke, and that device performs NAT on the IPsec flow.

  If you select this option, configure the **Interval**, in seconds, between the keepalive signals sent between the spoke and the middle device to indicate that the session is active. The value can be from 5 to 3600 seconds. The default is 20 seconds.

Enabling Sysopt Connection Permit-vpn Option

When you want to bypass the inspection of decrypted traffic, follow these steps to enable the *sysopt connection permit-vpn* option. However, the VPN filter ACL and authorization ACL downloaded from AAA server are still applied to VPN traffic.

Procedure

Step 1 On your FMC web interface, choose to **Objects** > **Object Management** > **FlexConfig** > **Text Object** > **Add Text Object**.

Step 2 Create a text object variable (for example, **vpnSysVar**), a single entry with value **sysopt**.

Step 3 Go to **Objects** > **Object Management** > **FlexConfig** > **FlexConfig Object** > **Add FlexConfig Object**.

Step 4 Create FlexConfig object with the command **connection permit-vpn**.

Step 5 Insert the text object variable in FlexConfig object at the start of the command, (Example, **$vpnSysVar connection permit-vpn**), and click **Save**.

Step 6 Apply the FlexConfig object Type as **Append** and select Deployment as **Everytime**.

Step 7 Go to **Devices** > **FlexConfig** and edit existing policy or create a new one.

Step 8 Add the newly created FlexConfig object, and click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Monitoring Firepower Threat Defense VPNs

Monitor Firepower Threat Defense VPN activity in the following ways:

- System Messages

  The Message Center is the place to start your monitoring. This feature allows you to view messages that are continually generated about system activities and status. To open the Message Center, click in the **System Status** icon, located to the immediate right of the **Deploy** button in the main menu. See System Messages, on page 311 for details on using the Message Center.

- VPN Health Events

  These events are displayed along with other system events under **System** > **Health** > **Events** > **VPN Status**. See Health Monitoring, on page 275 for details on viewing system health events.

- System Logs

  Currently the Firepower Management Center does not have the capability to read the Firepower Threat Defense VPN syslogs. These syslogs need to be forwarded to a third-party server for analysis and archiving. See About Configuring Syslog, on page 962 for details on configuring syslog servers and viewing the system logs.

- Unified CLI Commands, see *Command Reference for Firepower Threat Defense*

Use the show, clear, and debug commands on the Firepower Threat Defense device to monitor and troubleshoot VPN activity.

Monitoring Guidelines

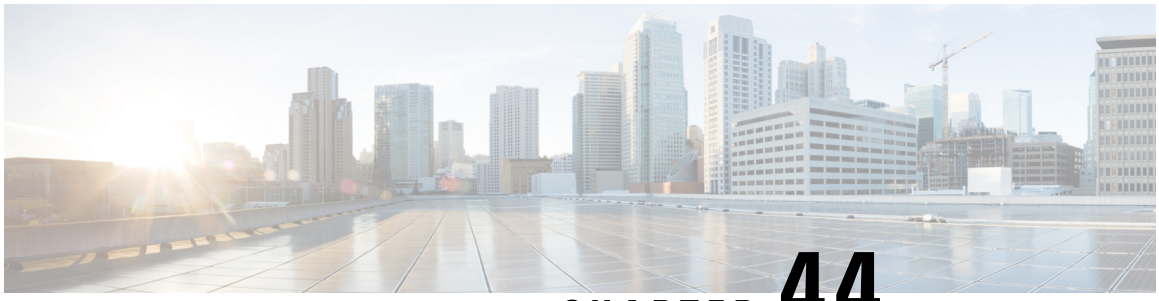
- If more than 300 Firepower Threat Defense devices are configured in the Firepower Management Center, event handling issues may arise.
- Event loss between the Firepower Threat Defense device and the Firepower Management Center is possible if the connection is broken.



PART XII

Appliance Platform Settings

- System Configuration, on page 867
- Platform Settings Policies, on page 927
- Platform Settings for Classic Devices, on page 931
- Platform Settings for Firepower Threat Defense, on page 949



CHAPTER 44

System Configuration

The following topics explain how to configure system configuration settings on Firepower Management Centers and managed devices:

- Requirements and Prerequisites for the System Configuration, on page 868
- About System Configuration, on page 868
- Appliance Information, on page 870
- HTTPS Certificates, on page 871
- External Database Access Settings, on page 877
- Database Event Limits, on page 878
- Management Interfaces, on page 881
- Shut Down or Restart, on page 889
- Remote Storage Management, on page 890
- Change Reconciliation, on page 893
- Policy Change Comments, on page 895
- Access List, on page 895
- Audit Logs, on page 897
- Audit Log Certificate, on page 900
- Dashboard Settings, on page 904
- DNS Cache, on page 905
- Email Notifications, on page 906
- Language Selection, on page 907
- Login Banners, on page 907
- SNMP Polling, on page 908
- Security Certifications Compliance, on page 909
- Time and Time Synchronization, on page 913
- Session Timeouts, on page 917
- Vulnerability Mapping, on page 918
- Remote Console Access Management, on page 918
- REST API Preferences, on page 924
- VMware Tools and Virtual Systems, on page 925

Requirements and Prerequisites for the System Configuration

Model Support

FMC

Some settings also apply to 7000 & 8000 Series devices.

Supported Domains

Global

User Roles

Admin

About System Configuration

System configuration settings apply to either a Firepower Management Center or a Classic managed device (7000 and 8000 Series, ASA FirePOWER, NGIPSv):

- For the Firepower Management Center these configuration settings are part of a "local" system configuration. Note that system configuration on the Firepower Management Center is specific to a single system, and changes to a FMC's system configuration affect only that system.
- For a Classic managed device, you apply a configuration from the Firepower Management Center as part of a platform settings policy. You create a shared policy to configure a subset of the system configuration settings, appropriate for managed devices, that are likely to be similar across a deployment.

Tip For 7000 and 8000 Series devices, you can perform limited system configuration tasks from the local web interface, such as console configuration and remote management. These are not the same configurations that you apply to a 7000 or 8000 Series device using a platform settings policy.

Navigating the Firepower Management Center System Configuration

The system configuration identifies basic settings for a Firepower Management Center.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Use the navigation panel to choose configurations to change; see Table 69: System Configuration Settings , on page 869 for more information.

System Configuration Settings

Note that for managed devices, many of these configurations are handled by a *platform settings* policy applied from the FMC; see Platform Settings Policies, on page 927. For 7000/8000 series devices, you can also log into the local web interface for non-policy based system configurations; see Local System Configuration for 7000/8000 Series Devices, on page 941.

Table 69: System Configuration Settings

Setting	Description
Access Control Preferences	Configure the system to prompt users for a comment when they add or modify an access control policy; see Policy Change Comments, on page 895.
Access List	Control which computers can access the system on specific ports; see Access List, on page 895.
Audit Log	Configure the system to send an audit log to an external host; see Audit Logs, on page 897.
Audit Log Certificate	Configure the system to secure the channel when streaming the audit log to an external host; see Audit Log Certificate, on page 900 .
Change Reconciliation	Configure the system to send a detailed report of changes to the system over the last 24 hours; see Change Reconciliation, on page 893.
Console Configuration	Configure console access via VGA or serial port, or via Lights-Out Management (LOM); see Remote Console Access Management, on page 918.
Dashboard	Enable Custom Analysis widgets on the dashboard; see Dashboard Settings, on page 904.
Database	Specify the maximum number of each type of event that the Firepower Management Center can store; see Database Event Limits, on page 878.
DNS Cache	Configure the system to resolve IP addresses automatically on event view pages; see DNS Cache, on page 905.
Email Notification	Configure a mail host, select an encryption method, and supply authentication credentials for email-based notifications and reporting; see Email Notifications, on page 906.
External Database Access	Enable external read-only access to the database, and provide a client driver to download; see External Database Access Settings, on page 877.
HTTPS Certificate	Request an HTTPS server certificate, if needed, from a trusted authority and upload certificates to the system; see HTTPS Certificates, on page 871.
Information	View current information about the appliance and edit the display name; see Appliance Information, on page 870.
Intrusion Policy Preferences	Configure the system to prompt users for a comment when they modify an intrusion policy; see Policy Change Comments, on page 895.
Language	Specify a different language for the web interface; see Language Selection, on page 907.
Login Banner	Create a custom login banner that appears when users log in; see Login Banners, on page 907.

Setting	Description
Management Interfaces	Change options such as the IP address, hostname, and proxy settings of the appliance; see Management Interfaces, on page 881.
Network Analysis Policy Preferences	Configure the system to prompt users for a comment when they modify a network analysis policy; see Policy Change Comments, on page 895.
Process	Shut down, reboot, or restart Firepower processes; see Shut Down or Restart, on page 889.
Remote Storage Device	Configure remote storage for backups and reports; see Remote Storage Management, on page 890.
REST API Preferences	Enable or disable access to the Firepower Management Center via the Firepower REST API; see REST API Preferences, on page 924.
Shell Timeout	Configure the amount of idle time, in minutes, before a user's login session times out due to inactivity; see Session Timeouts, on page 917.
SNMP	Enable Simple Network Management Protocol (SNMP) polling; see SNMP Polling, on page 908.
Time	View and change the current time setting; see Time and Time Synchronization, on page 913.
Time Synchronization	Manage time synchronization on the system; see Time and Time Synchronization, on page 913.
UCAPL/CC Compliance	Enable compliance with specific requirements set out by the United States Department of Defense; see Enable Security Certifications Compliance, on page 911.
VMware Tools	Enable and use VMware Tools on a Firepower Management Center Virtual; see VMware Tools and Virtual Systems, on page 925.
Vulnerability Mapping	Map vulnerabilities to a host IP address for any application protocol traffic received or sent from that address; see Vulnerability Mapping, on page 918.

Related Topics

About Platform Settings for Classic Devices, on page 931

Appliance Information

The **System > Configuration** page of the web interface includes the information listed in the table below. Unless otherwise noted, all fields are read-only.

Note See also the **Help > About** page, which includes similar but slightly different information.

Field	Description
Name	A descriptive name you assign to the FMCappliance. Although you can use the host name as the name of the appliance, entering a different name in this field does not change the host name.
Product Model	The model name of the appliance.

Field	Description
Serial Number	The serial number of the appliance.
Software Version	The version of the software currently installed on the appliance.
Prohibit Packet Transfer to the Firepower Management Center	Specifies whether the managed device sends packet data with events, allowing the data to be stored on the Firepower Management Center. This setting is available on the local web interface on 7000 and 8000 Series devices.
Operating System	The operating system currently running on the appliance.
Operating System Version	The version of the operating system currently running on the appliance.
IPv4 Address	The IPv4 address of the default (`eth0`) management interface. If IPv4 management is disabled, this field indicates that.
IPv6 Address	The IPv6 address of the default (`eth0`) management interface. If IPv6 management is disabled, this field indicates that.
Current Policies	The system-level policies currently deployed. If a policy has been updated since it was last deployed, the name of the policy appears in italics.
Model Number	The appliance-specific model number stored on the internal flash drive. This number may be important for troubleshooting.

HTTPS Certificates

Secure Sockets Layer (SSL)/TLS certificates enable Firepower Management Centers and 7000 and 8000 Series devices to establish an encrypted channel between the system and a web browser. A default certificate is included with all Firepower devices, but it is not generated by a certificate authority (CA) trusted by any globally known CA. For this reason, consider replacing it with a custom certificate signed by a globally known or internally trusted CA.



Caution The FMC supports 4096-bit HTTPS certificates. If the certificate used by the FMC was generated using a public server key larger than 4096 bits, you will not be able to log in to the FMC web interface. If this happens, contact Cisco TAC.

Default HTTPS Server Certificates

If you use the default server certificate provided with an appliance, do not configure the system to require a valid HTTPS client certificate for web interface access because the default server certificate is not signed by the CA that signs your client certificate.

The default server certificate provided with an appliance expires 20 years from when it was first generated.

Custom HTTPS Server Certificates

You can use the Firepower Management Center web interface to generate a server certificate request based on your system information and the identification information you supply. You can use that request to sign a certificate if you have an internal certificate authority (CA) installed that is trusted by your browser. You can also send the resulting request to a certificate authority to request a server certificate. After you have a signed certificate from a certificate authority (CA), you can import it.

HTTPS Server Certificate Requirements

When you use HTTPS certificates to secure the connection between your web browser and the Firepower appliance web interface, you must use certificates that comply with the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile (RFC 5280). When you import a server certificate to the appliance, the system rejects the certificate if it does not comply with version 3 (X.509 v3) of that standard.

Before importing an HTTPS server certificate, be certain it includes the following fields:

Certificate Field	Description
Version	Version of the encoded certificate. Use version 3. See RFC 5280, section 4.1.2.1.
Serial number	A positive integer assigned to the certificate by the issuing CA. Issuer and serial number together uniquely identify the certificate. See RFC 5280, section 4.1.2.2.
Signature	Identifier for the algorithm used by the CA to sign the certificate. Must match the signatureAlgorithm field. See RFC 5280, section 4.1.2.3.
Issuer	Identifies the entity that signed and issued the certificate. See RFC 5280, section 4.1.2.4.
Validity	Interval during which the CA warrants that it will maintain information about the status of the certificate. See RFC 5280, section 4.1.2.5.
Subject	Identifies the entitity associated with the public key stored in the subject public key field; must be an X.500 disinguished name (DN). See RFC 5280, section 4.1.2.6.
Subject Public Key Info	Public key and an identifier for its algorithm. See RFC 5280, section 4.1.2.7.
Authority Key Identifier	Provides a means of identifying the public key corresponding to the private key used to sign a certificate. See RFC 5280, section 4.2.1.1.
Subject Key Identifier	Provides a means of identifying certificates that contain a particular public key. See RFC 5280, section 4.2.1.2.

Certificate Field	Description
Key Usage	Defines the purpose of the key contained in the certificates. See RFC 5280, section 4.2.1.3.
Basic Constraints	Identifies whether the certificate Subject is a CA, and the maximum depth of validation certification paths that include this certificate. See RFC 5280, section 4.2.1.9. For server certificates used in Firepower appliances, use `critical CA:FALSE`.
Extended Key Usage extension	Indicates one or more purposes for which the certified public key may be used, in addition to or in place of the basic purposes indicated in the Key Usage extension. See RFC 5280, section 4.2.1.12. Be certain you import certificates that can be used as server certificates.
signatureAlgorithm	Identifier for the algorithm the CA used to sign the certificate. Must match the Signature field. See RFC 5280, section 4.1.1.2.
signatureValue	Digital signature. See RFC 5280, section 4.1.1.3.

HTTPS Client Certificates

You can restrict access to the Firepower System web server using client browser certificate checking. When you enable user certificates, the web server checks that a user's browser client has a valid user certificate selected. That user certificate must be generated by the same trusted certificate authority that is used for the server certificate. The browser cannot load the web interface under any of the following circumstances:

- The user selects a certificate in the browser that is not valid.
- The user selects a certificate in the browser that is not generated by the certificate authority that signed the server certificate.
- The user selects a certificate in the browser that is not generated by a certificate authority in the certificate chain on the device.

To verify client browser certificates, configure the system to use the online certificate status protocol (OCSP) or load one or more certificate revocation lists (CRLs). Using the OCSP, when the web server receives a connection request it communicates with the certificate authority to confirm the client certificate's validity before establishing the connection. If you configure the server to load one or more CRLs, the web server compares the client certificate against those listed in the CRLs. If a user selects a certificate that is listed in a CRL as a revoked certificate, the browser cannot load the web interface.

Note If you choose to verify certificates using CRLs, the system uses the same CRLs to validate both client browser certificates and audit log server certificates.

Viewing the Current HTTPS Server Certificate

You can only view server certificates for the appliance you are logged in to.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **HTTPS Certificate**.

Generating an HTTPS Server Certificate Signing Request

If you install a certificate that is not signed by a globally known or internally trusted CA, the user's browser displays a security warning when they try to connect to the web interface.

A certificate signing request (CSR) is unique to the appliance or device from which you generated it. You cannot generate a CSR for multiple devices from a single appliance. Although all fields are optional, we recommend entering values for the following: CN, Organization, Organization Unit, City/Locality, State/Province, Country/Region.

The key generated for the certificate request is in Base-64 encoded PEM format.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **HTTPS Certificate**.

Step 3 Click **Generate New CSR**.

Step 4 Enter a country code in the **Country Name (two-letter code)** field.

Step 5 Enter a state or province postal abbreviation in the **State or Province** field.

Step 6 Enter a **Locality or City**.

Step 7 Enter an **Organization** name.

Step 8 Enter an **Organizational Unit (Department)** name.

Step 9 Enter the fully qualified domain name of the server for which you want to request a certificate in the **Common Name** field.

Note Enter the fully qualified domain name of the server exactly as it should appear in the certificate in the **Common Name** field. If the common name and the DNS hostname do not match, you receive a warning when connecting to the appliance.

Step 10 Click **Generate**.

Step 11 Open a text editor.

Step 12 Copy the entire block of text in the certificate request, including the `BEGIN CERTIFICATE REQUEST` and `END CERTIFICATE REQUEST` lines, and paste it into a blank text file.

Step 13 Save the file as *`servername`*`.csr`, where *servername* is the name of the server where you plan to use the certificate.

Step 14 Click **Close**.

What to do next

- Submit the certificate request to the certificate authority.
- When you receive the signed certificate, import it to the Firepower Management Center; see Importing HTTPS Server Certificates, on page 875.

Importing HTTPS Server Certificates

If the signing authority that generated the certificate requires you to trust an intermediate CA, you must also supply a certificate chain (or certificate path).

If you require client certificates, accessing an appliance via the web interface will fail when the server certificate does not meet either of the following criteria:

- The certificate is signed by the same CA that signed the client certificate.
- The certificate is signed by a CA that has signed an intermediate certificate in the certificate chain.

Caution The Firepower Management Center supports 4096-bit HTTPS certificates. If the certificate used by the Firepower Management Center was generated using a public server key larger than 4096 bits, you will not be able to log in to the FMC web interface. For more information about updating HTTPS Certificates to Version 6.0.0, see "Update Management Center HTTPS Certificates to Version 6.0" in *Firepower System Release Notes, Version 6.0*. If you generate or import an HTTPS Certificate and cannot log in to the FMC web interface, contact Support.

Before you begin

- Generate a certificate signing request; see Generating an HTTPS Server Certificate Signing Request, on page 874.
- Upload the CSR file to the certificate authority where you want to request a certificate, or use the CSR to create a self-signed certificate.
- Confirm that the certificate meets the requirements described in HTTPS Server Certificate Requirements, on page 872.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **HTTPS Certificate**.

Step 3 Click **Import HTTPS Server Certificate**.

Note You cannot import an encrypted HTTPS certificate.

Step 4 Open the server certificate in a text editor, copy the entire block of text, including the `BEGIN CERTIFICATE` and `END CERTIFICATE` lines. Paste this text into the **Server Certificate** field.

Step 5 Whether you must supply a **Private Key** depends on how you generated the Certificate Signing Request:

- If you generated the Certificate Signing Request using the Firepower Management Center web interface (as described in Generating an HTTPS Server Certificate Signing Request, on page 874), the system already has the private key and you need not enter one here.
- If you generated the Certificate Signing Request using some other means, you must supply the private key here. Open the private key file and copy the entire block of text, include the `BEGIN RSA PRIVATE KEY` and `END RSA PRIVATE KEY` lines. Paste this text into the **Private Key** field.

Step 6 Open any required intermediate certificates, copy the entire block of text for each, and paste it into the **Certificate Chain** field. If you received a root certificate, paste it here. If you received an intermediate certificate, paste it below the root certificate. In both cases, copy the entire block of text, including the `BEGIN CERTIFICATE` and `END CERTIFICATE` lines.

Step 7 Click **Save**.

Requiring Valid HTTPS Client Certificates

Use this procedure to require users connecting to the FMC web interface to supply a user certificate. The system supports validating HTTPS client certificates using either OCSP or imported CRLs in Privacy-enhanced Electronic Mail (PEM) format.

If you choose to use CRLs, to ensure that the list of revoked certificates stays current, you can create a scheduled task to update the CRLs. The system displays the most recent refresh of the CRLs.

Note To access the web interface after enabling client certificates, you **must** have a valid client certificate present in your browser (or a CAC inserted in your reader).

Before you begin

- Import a server certificate signed by the same certificate authority that signed the client certificate to be used for the connection; see Importing HTTPS Server Certificates, on page 875.
- Import the server certificate chain if needed; see Importing HTTPS Server Certificates, on page 875.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **HTTPS Certificate**.

Step 3 Choose **Enable Client Certificates**. If prompted, select the appropriate certificate from the drop-down list.

Step 4 You have three options:

- To verify client certificates using one or more CRLS, select **Enable Fetching of CRL** and continue with Step 5.
- To verify client certificates using OCSP, select **Enable OCSP** and skip to Step 7.

- To accept client certificates without checking for revocation, skip to Step 8.

Step 5 Enter a valid URL to an existing CRL file and click **Add CRL**. Repeat to add up to 25 CRLs.

Step 6 Click **Refresh CRL** to load the current CRL or CRLs from the specified URL or URLs.

Note Enabling fetching of the CRL creates a scheduled task to regularly update the CRL or CRLs. Edit the task to set the frequency of the update.

Step 7 Verify that the client certificate is signed by the certificate authority loaded onto the appliance and the server certificate is signed by a certificate authority loaded in the browser certificate store. (These should be the same certificate authority.)

Caution Saving a configuration with enabled client certificates, with no valid client certificate in your browser certificate store, disables all web server access to the appliance. Make sure that you have a valid client certificate installed before saving settings.

Step 8 Click **Save**.

Related Topics

Configuring Certificate Revocation List Downloads, on page 182

External Database Access Settings

You can configure the Firepower Management Center to allow read-only access to its database by a third-party client. This allows you to query the database using SQL using any of the following:

- industry-standard reporting tools such as Actuate BIRT, JasperSoft iReport, or Crystal Reports
- any other reporting application (including a custom application) that supports JDBC SSL connections
- the Cisco-provided command-line Java application called RunQuery, which you can either run interactively or use to obtain comma-separated results for a single query

Use the Firepower Management Center's system configuration to enable database access and create an access list that allows selected hosts to query the database. Note that this access list does not also control appliance access.

You can also download a package that contains the following:

- RunQuery, the Cisco-provided database query tool
- InstallCert, a tool you can use to retrieve and accept the SSL certificate from the Firepower Management Center you want to access
- the JDBC driver you must use to connect to the database

See the *Firepower System Database Access Guide* for information on using the tools in the package you downloaded to configure database access.

Enabling External Access to the Database

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **External Database Access**.

Step 3 Select the **Allow External Database Access** check box.

Step 4 Enter an appropriate value in the **Server Hostname** field. Depending on your third-party application requirements, this value can be either the fully qualified domain name (FQDN), IPv4 address, or IPv6 address of the Firepower Management Center.

Note In an FMC high availability setup, enter only the active peer details. We do not recommend entering details of the standby peer.

Step 5 Next to **Client JDBC Driver**, click **Download** and follow your browser's prompts to download the `client.zip` package.

Step 6 To add database access for one or more IP addresses, click **Add Hosts**. An **IP Address** field appears in the **Access List** field.

Step 7 In the **IP Address** field, enter an IP address or address range, or `any`.

Step 8 Click **Add**.

Step 9 Click **Save**.

Tip If you want to revert to the last saved database settings, click **Refresh**.

Related Topics

Firepower System IP Address Conventions, on page 15

Database Event Limits

To manage disk space, the FMC periodically prunes the oldest intrusion events, audit records, Security Intelligence data, and URL filtering data from the event database. For each event type, you can specify how many records the FMC retains after pruning; never rely on the event database containing more records of any type than the retention limit configured for that type. To improve performance, tailor the event limits to the number of events you regularly work with. You can optionally choose to receive email notifications when pruning occurs. For some event types, you can disable storage.

To manually delete individual events, use the event viewer. You can also manually purge the database; see Data Storage, on page 197.

Configuring Database Event Limits

Before you begin

- If you want to receive email notifications when events are pruned from the Firepower Management Center's database, you must configure an email server; see Configuring a Mail Relay Host and Notification Address, on page 906.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Choose **Database**.

Step 3 For each of the databases, enter the number of records you want to store.

For information on how many records each database can maintain, see Database Event Limits, on page 879.

Step 4 Optionally, in the **Data Pruning Notification Address** field, enter the email address where you want to receive pruning notifications.

Step 5 Click **Save**.

Database Event Limits

The following table lists the minimum and maximum number of records for each event type that you can store on a Firepower Management Center.

Table 70: Database Event Limits

Event Type	Upper Limit	Lower Limit
Intrusion events	10 million (FMC Virtual) 20 million (FMC750) 30 million (FMC1000, FMC1500,) 60 million (FMC2000,FMC2500) 150 million (FMC3500) 300 million (FMC4000, FMC4500)	10,000
Discovery events	10 million (FMC Virtual) 20 million (FMC2000, FMC2500, FMC4000, FMC4500)	Zero (disables storage)

Event Type	Upper Limit	Lower Limit
Connection events Security Intelligence events	50 million (FMC Virtual, FMC750) 100 million (FMC1000, FMC1500,) 300 million (FMC2000, FMC2500) 500 million (FMC3500) 1 billion (FMC4000, FMC4500) Limit is shared between connection events and Security Intelligence events. The sum of the configured maximums cannot exceed this limit.	Zero (disables storage) Setting **Maximum Connection Events** to zero immediately purges existing connection events. Note that disabling connection event storage on the Firepower Management Center does not affect connection summaries or correlation. The system still uses connection event information for features like traffic profiles, correlation policies, and dashboard displays.
Connection summaries (aggregated connection events)	50 million (FMC Virtual, FMC750) 100 million (FMC1000, FMC1500,) 300 million (FMC2000, FMC2500) 500 million (FMC3500) 1 billion (FMC4000, FMC4500)	Zero (disables storage)
Correlation events and compliance white list events	1 million (FMC Virtual) 2 million (FMC2000, FMC2500, FMC4000, FMC4500)	One
Malware events	10 million (FMC Virtual) 20 million (FMC2000,FMC2500, FMC4000, FMC4500)	10,000
File events	10 million (FMC Virtual) 20 million (FMC2000, FMC2500, FMC4000, FMC4500)	Zero (disables storage)
Health events	1 million	Zero (disables storage)
Audit records	100,000	One
Remediation status events	10 million	One
White list violation history	a 30-day history of violations	One day’s history
User activity (user events)	10 million	One
User logins (user history)	10 million	One

Event Type	Upper Limit	Lower Limit
Intrusion rule update import log records	1 million	One

Management Interfaces

After setup, you can change the management network settings, including adding more management interfaces, hostname, search domains, DNS servers, and HTTP proxy on the FMC.

About FMC Management Interfaces

By default, the FMC manages all devices on a single management interface. You can also perform initial setup on the management interface and log into the FMC on this interface as an administrator. The management interface is also used to communicate with the Smart Licensing server, to download updates, and to perform other management functions.

For information about device management interfaces, see About Device Management Interfaces, on page 221.

Management Interfaces on the FMC

The FMC uses the eth0 interface for initial setup, HTTP access for administrators, management of devices, as well as other management functions such as licensing and updates.

You can also configure additional management interfaces on the same network, or on different networks. When the FMC manages large numbers of devices, adding more management interfaces can improve throughput and performance. You can also use these interfaces for all other management functions. You might want to use each management interface for particular functions; for example, you might want to use one interface for HTTP administrator access and another for device management.

For device management, the management interface carries two separate traffic channels: the *management traffic channel* carries all internal traffic (such as inter-device traffic specific to managing the device), and the *event traffic channel* carries all event traffic (such as web events). You can optionally configure a separate event-only interface on the FMC to handle event traffic; you can configure only one event interface. Event traffic can use a large amount of bandwidth, so separating event traffic from management traffic can improve the performance of the FMC. For example, you can assign a 10 GigabitEthernet interface to be the event interface, if available, while using 1 GigabitEthernet interfaces for management. You might want to configure an event-only interface on a completely secure, private network while using the regular management interface on a network that includes Internet access, for example. You can also use both management and event interfaces on the same network if the goal is only to take advantage of increased throughput. Managed devices will send management traffic to the FMC management interface and event traffic to the FMCs event-only interface. If the managed device cannot reach the event-only interface, then it will fall back to sending events to the management interface.

Note All management interfaces support HTTP administrator access as controlled by your Access List configuration. Conversely, you cannot restrict an interface to *only* HTTP access; management interfaces always support device management (management traffic, event traffic, or both).

Note Only the eth0 interface supports DHCP IP addressing. Other management interfaces only support static IP addresses.

Management Interface Support Per FMC Model

See the hardware installation guide for your model for the management interface locations.

See the following table for supported management interfaces on each FMC model.

Table 71: Management Interface Support on the FMC

Model	Management Interfaces
MC750, MC1500, MC3500	eth0 (Default) eth1
MC2000, MC4000	eth0 (Default) eth1 eth2 eth3
MC1000	eth0 (Default) eth1
MC2500, MC4500	eth0 (Default) eth1 eth2 eth3
Firepower Management Center Virtual	eth0 (Default)

Network Routes on FMC Management Interfaces

Management interfaces (including event-only interfaces) support only static routes to reach remote networks. When you set up your FMC, the setup process creates a default route to the gateway IP address that you specify. You cannot delete this route; you can only modify the gateway address.

The default route always uses the lowest-numbered management interface (e.g. eth0).

At least one static route is recommended per management interface to access remote networks. We recommend placing each interface on a separate network to avoid potential routing problems, including routing problems from other devices to the FMC. If you do not experience problems with interfaces on the same network, then be sure to configure static routes correctly. For example, on the FMC both eth0 and eth1 are on the same network, but you want to manage a different group of devices on each interface. The default gateway is 192.168.45.1. If you want eth1 to manage devices on the remote 10.6.6.0/24 destination network, you can create a static route for 10.6.6.0/24 through eth1 with the same gateway of 192.168.45.1. Traffic to 10.6.6.0/24 will hit this route before it hits the default route, so eth1 will be used as expected.

If you want to use two FMC interfaces to manage remote devices that are on the same network, then static routing on the FMC may not scale well, because you need separate static routes per device IP address.

Another example includes separate management and event-only interfaces on both the FMC and the managed device. The event-only interfaces are on a separate network from the management interfaces. In this case, add a static route through the event-only interface for traffic destined for the remote event-only network, and vice versa.

NAT Environments

Network address translation (NAT) is a method of transmitting and receiving network traffic through a router that involves reassigning the source or destination IP address. The most common use for NAT is to allow private networks to communicate with the internet. Static NAT performs a 1:1 translation, which does not pose a problem for FMC communication with devices, but port address translation (PAT) is more common. PAT lets you use a single public IP address and unique ports to access the public network; these ports are dynamically assigned as needed, so you cannot initiate a connection to a device behind a PAT router.

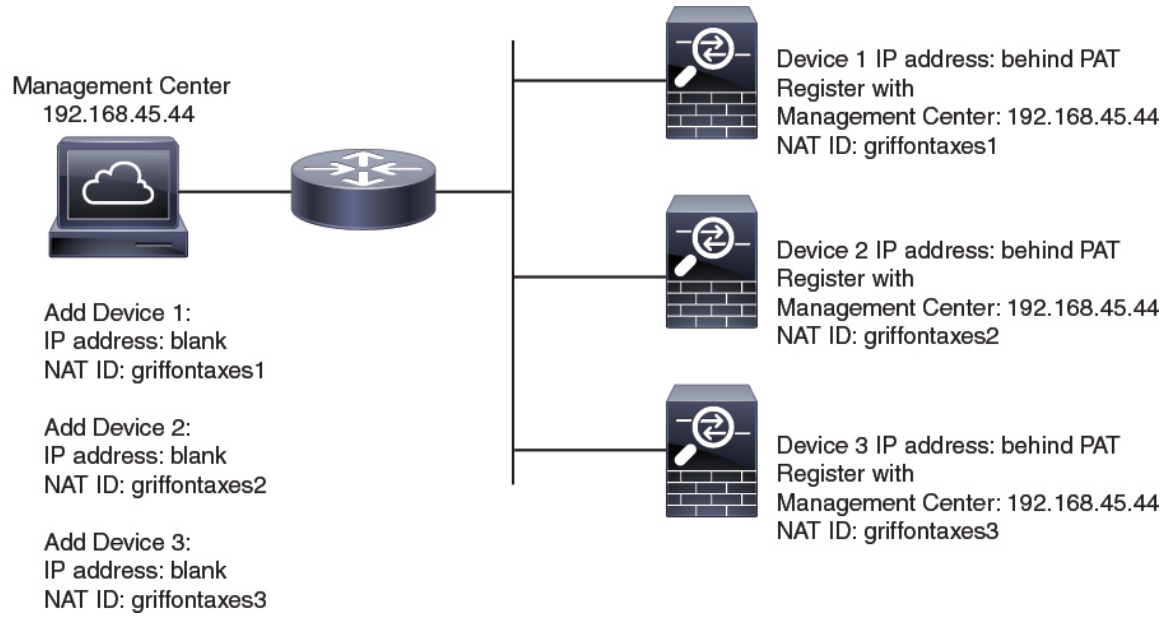
Normally, you need both IP addresses (along with a registration key) for both routing purposes and for authentication: the FMC specifies the device IP address when you add a device, and the device specifies the FMC IP address. However, if you only know one of the IP addresses, which is the minimum requirement for routing purposes, then you must also specify a unique NAT ID on both sides of the connection to establish trust for the initial communication and to look up the correct registration key. The FMC and device use the registration key and NAT ID (instead of IP addresses) to authenticate and authorize for initial registration.

For example, you add a device to the FMC, and you do not know the device IP address (for example, the device is behind a PAT router), so you specify only the NAT ID and the registration key on the FMC; leave the IP address blank. On the device, you specify the FMC IP address, the same NAT ID, and the same registration key. The device registers to the FMC's IP address. At this point, the FMC uses the NAT ID instead of IP address to authenticate the device.

Although the use of a NAT ID is most common for NAT environments, you might choose to use the NAT ID to simplify adding many devices to the FMC. On the FMC, specify a unique NAT ID for each device you want to add while leaving the IP address blank, and then on each device, specify both the FMC IP address and the NAT ID. Note: The NAT ID must be unique per device.

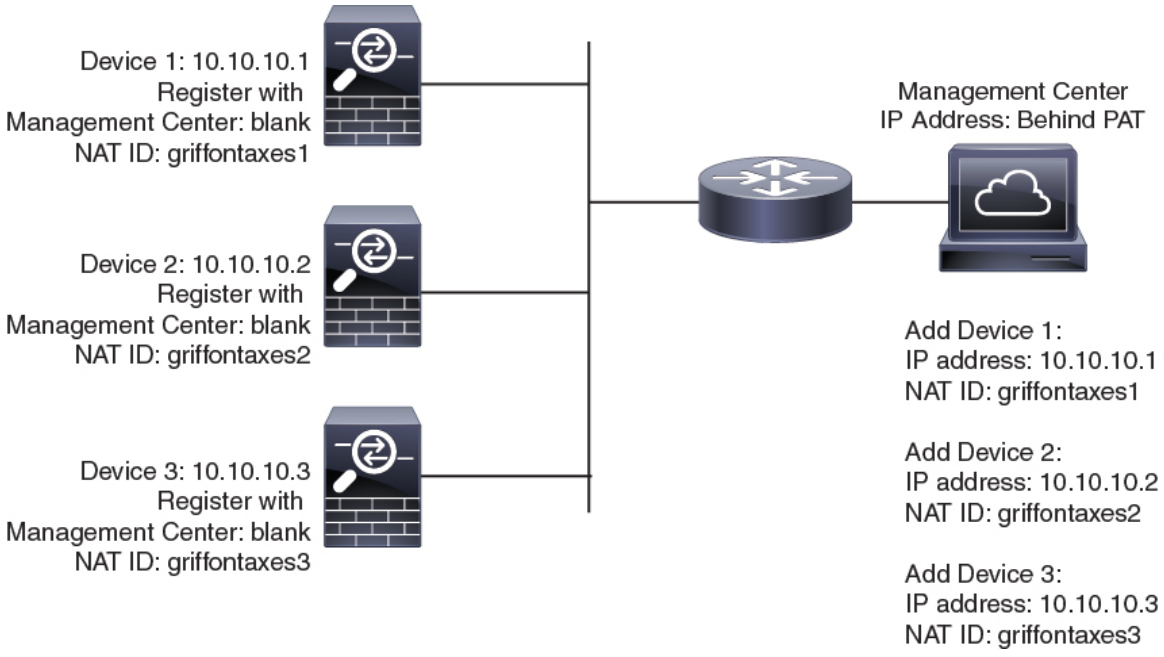
The following example shows three devices behind a PAT IP address. In this case, specify a unique NAT ID per device on both the FMC and the devices, and specify the FMC IP address on the devices.

Figure 26: NAT ID for Managed Devices Behind PAT



The following example shows the FMC behind a PAT IP address. In this case, specify a unique NAT ID per device on both the FMC and the devices, and specify the device IP addresses on the FMC.

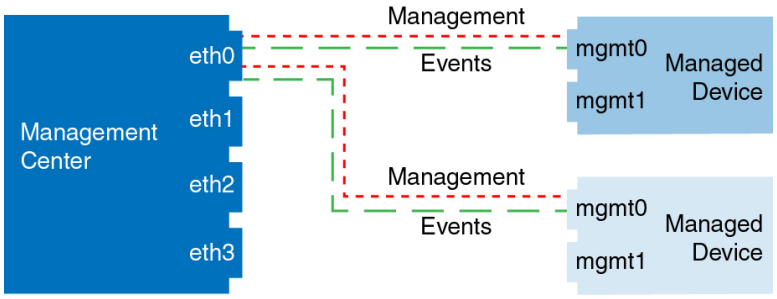
Figure 27: NAT ID for FMC Behind PAT



Management and Event Traffic Channel Examples

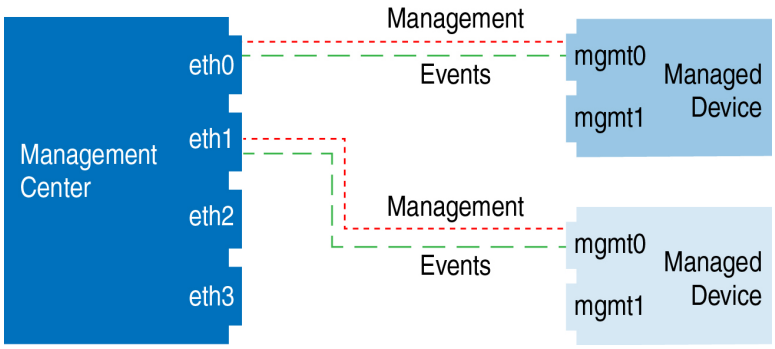
The following example shows the Firepower Management Center and managed devices using only the default management interfaces.

Figure 28: Single Management Interface on the Firepower Management Center



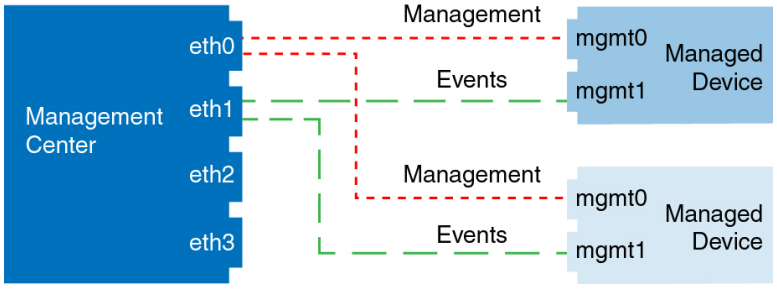
The following example shows the Firepower Management Center using separate management interfaces for devices; and each managed device using 1 management interface.

Figure 29: Mutliple Management Interfaces on the Firepower Management Center



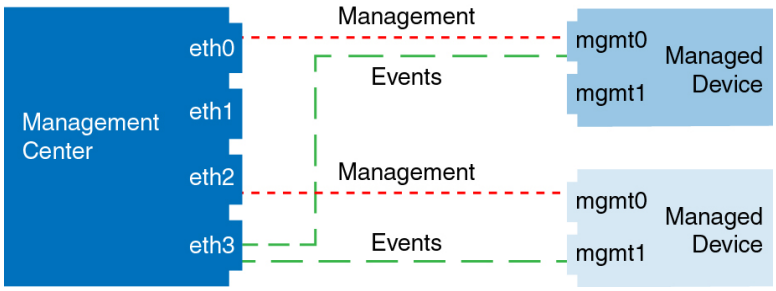
The following example shows the Firepower Management Center and managed devices using a separate event interface.

Figure 30: Separate Event Interface on the Firepower Management Center and Managed Devices



The following example shows a mix of multiple management interfaces and a separate event interface on the Firepower Management Center and a mix of managed devices using a separate event interface, or using a single management interface.

Figure 31: Mixed Management and Event Interface Usage



Modify FMC Management Interfaces

Caution Do NOT push the FMC deployments over a VPN tunnel that is terminating directly on the Firepower Threat Defense. Pushing the FMC deployments can potentially inactivate the tunnel and disconnect the FMC and the Firepower Threat Defense.

Recovering the device from this situation can be very disruptive and require executing the disaster recovery procedure. This procedure resets the Firepower Threat Defense configuration to factory defaults by changing manager from FMC to local and configuring the device from beginning. For more information, see Deploying the FMC Policy Configuration over VPN Tunnel, on page 344.

Modify the management interface settings on the Firepower Management Center. You can optionally enable additional management interfaces or configure an event-only interface.

Caution Be careful when making changes to the management interface to which you are connected; if you cannot re-connect because of a configuration error, you need to access the FMC console port to re-configure the network settings in the Linux shell. You must contact Cisco TAC to guide you in this operation.

Note If you change the FMC IP address, then see the following tasks to ensure device management connectivity depending on how you added the device to the FMC:

- **IP address—No action.** If you added the device to the FMC using a reachable device IP address, then the management connection will be reestablished automatically after several minutes even though the IP address identified on the FTD is the old IP address. **Note**: If you specified a device IP address that is unreachable, then you must contact Cisco TAC, who can advise you how to restore connectivity for your devices.
- **NAT ID only—Contact Cisco TAC.** If you added the device using only the NAT ID, then the connection cannot be reestablished. In this case, you must contact Cisco TAC, who can advise you how to restore connectivity for your devices.

Before you begin

- For information about how device management works, see About Device Management Interfaces, on page 221.
- If you use a proxy:
  - Proxies that use NT LAN Manager (NTLM) authentication are not supported.
  - If you use or will use Smart Licensing, the proxy FQDN cannot have more than 64 characters.

Procedure

Step 1 Choose **System** > **Configuration**, and then choose **Management Interfaces**.

Step 2 In the **Interfaces** area, click **Edit** next to the interface that you want to configure.

All available interfaces are listed in this section. You cannot add more interfaces.

You can configure the following options on each management interface:

- **Enabled**—Enable the management interface. Do **not** disable the default eth0 management interface. Some processes require the eth0 interface.
- **Channels**—Configure an event-only interface; you can configure only one event interface on the FMC. To do so, uncheck the **Management Traffic** check box, and leave the **Event Traffic** check box checked. You can optionally disable **Event Traffic** for the management interface(s). In either case, the device will try to send events to the event-only interface, and if that interface is down, it will send events on the management interface even if you disable the event channel. You cannot disable both event and management channels on an interface.
- **Mode**—Specify a link mode. Note that any changes you make to auto-negotiation are ignored for GigabitEthernet interfaces.
- **MDI/MDIX**—Set the **Auto-MDIX** setting.
- **MTU**—Set the maximum transmission unit (MTU). The default is 1500. The range within which you can set the MTU can vary depending on the model and interface type.

  Because the system automatically trims 18 bytes from the configured MTU value, any value below 1298 does not comply with the minimum IPv6 MTU setting of 1280, and any value below 594 does not comply with the minimum IPv4 MTU setting of 576. For example, the system automatically trims a configured value of 576 to 558.
- **IPv4 Configuration**—Set the IPv4 IP address. Choose:
  - **Static**—Manually enter the **IPv4 Management IP** address and **IPv4 Netmask**.
  - **DHCP**—Set the interface to use DHCP (eth0 only).
  - **Disabled**—Disable IPv4. Do **not** disable both IPv4 and IPv6.
- **IPv6 Configuration**—Set the IPv6 IP address. Choose:
  - **Static**—Manually enter the **IPv6 Management IP** address and **IPv6 Prefix Length**.
  - **DHCP**—Set the interface to use DHCPv6 (eth0 only).

- **Router Assigned**—Enable stateless autoconfiguration.
- **Disabled**—Disable IPv6. Do **not** disable both IPv4 and IPv6.

Step 3 In the **Routes** area, edit a static route by clicking **Edit** (), or add a route by clicking **Add** ().

View the route table by clicking .

You need a static route for each additional interface to reach remote networks. For more information about when new routes are needed, see Network Routes on FMC Management Interfaces, on page 882.

Note For the default route, you can change only the gateway IP address. The default route always uses the eth0 interface.

You can configure the following settings for a static route:

- **Destination**—Set the destination address of the network to which you want to create a route.
- **Netmask** or **Prefix Length**—Set the netmask (IPv4) or prefix length (IPv6) for the network.
- **Interface**—Set the egress management interface.
- **Gateway**—Set the gateway IP address.

Step 4 In the **Shared Settings** area, set network parameters shared by all interfaces.

Note If you selected **DHCP** for the eth0 interface, you cannot manually specify some shared settings derived from the DHCP server.

You can configure the following shared settings:

- **Hostname**—Set the FMC hostname. The hostname must start and end with a letter or digit, and have only letters, digits, or a hyphen. If you change the hostname, reboot the FMC if you want the new hostname reflected in syslog messages. Syslog messages do not reflect a new hostname until after a reboot.
- **Domains**—Set the search domain(s) for the FMC, separated by commas. These domains are added to hostnames when you do not specify a fully-qualified domain name in a command, for example, **ping system**. The domains are used only on the management interface, or for commands that go through the management interface.
- **Primary DNS Server**, **Secondary DNS Server**, **Tertiary DNS Server**—Set the DNS servers to be used in order of preference.
- **Remote Management Port**—Set the remote management port for communication with managed devices. The FMC and managed devices communicate using a two-way, SSL-encrypted communication channel, which by default is on port 8305.

  Note Cisco **strongly** recommends that you keep the default settings for the remote management port, but if the management port conflicts with other communications on your network, you can choose a different port. If you change the management port, you must change it for **all** devices in your deployment that need to communicate with each other.

Step 5 In the **Proxy** area, configure HTTP proxy settings.

The FMC is configured to directly-connect to the internet on ports TCP/443 (HTTPS) and TCP/80 (HTTP). You can use a proxy server, to which you can authenticate via HTTP Digest.

See proxy requirements in the prerequisites to this topic.

a) Check the **Enabled** check box.
b) In the **HTTP Proxy** field, enter the IP address or fully-qualified domain name of your proxy server.

   See requirements in the prerequisites to this topic.

c) In the **Port** field, enter a port number.
d) Supply authentication credentials by choosing **Use Proxy Authentication**, and then provide a **User Name** and **Password**.

Step 6 Click **Save**.

Step 7 If you change the FMC IP address, then see If you change the FMC IP address, then see the following tasks to ensure device management connectivity depending on how you added the device to the FMC:

- **IP address—No action.** If you added the device to the FMC using a reachable device IP address, then the management connection will be reestablished automatically after several minutes even though the IP address identified on the FTD is the old IP address. **Note**: If you specified a device IP address that is unreachable, then you must contact Cisco TAC, who can advise you how to restore connectivity for your devices.
- **NAT ID only—Contact Cisco TAC.** If you added the device using only the NAT ID, then the connection cannot be reestablished. In this case, you must contact Cisco TAC, who can advise you how to restore connectivity for your devices.

Shut Down or Restart

Use the web interface to control the shut down and restart of processes on the FMC. You can:

- Shut down: Initiate a graceful shutdown of the appliance.

Caution Do **not** shut off Firepower appliances using the power button; it may cause a loss of data. Using the web interface (or CLI) prepares the system to be safely powered off and restarted without losing configuration data.

- Reboot: Shut down and restart gracefully.
- Restart the console: Restart the communications, database, and HTTP server processes. This is typically used during troubleshooting.

Tip For information on shutting down/restarting a 7000/8000 series device, including restarting the Snort process, see Shut Down or Restart a 7000/8000 Series Device, on page 947. For virtual devices, refer to the documentation for your virtual platform. For VMware in particular, custom power options are part of VMware Tools.

Shut Down or Restart the FMC

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Choose **Process**.

Step 3 Do one of the following:

Shut down	Click **Run Command** next to **Shutdown Management Center**.
Reboot	Click **Run Command** next to **Reboot Management Center**. **Note** Rebooting logs you out, and the system runs a database check that can take up to an hour to complete.
Restart the console	Click **Run Command** next to **Restart Management Center Console**. **Note** Restarting may cause deleted hosts to reappear in the network map.

Related Topics

Snort® Restart Scenarios, on page 348

Remote Storage Management

On Firepower Management Centers, you can use the following for local or remote storage for backups and reports:

- Network File System (NFS)
- Server Message Block (SMB)/Common Internet File System (CIFS)
- Secure Shell (SSH)

You cannot send backups to one remote system and reports to another, but you can choose to send either to a remote system and store the other on the Firepower Management Center.

Tip After configuring and selecting remote storage, you can switch back to local storage **only** if you **have not** increased the connection database limit.

Configuring Local Storage

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Choose **Remote Storage Device**.

Step 3 Choose **Local (No Remote Storage)** from the **Storage Type** drop-down list.

Step 4 Click **Save**.

Configuring NFS for Remote Storage

Before you begin

- Ensure that your external remote storage system is functional and accessible from your FMC.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **Remote Storage Device**.

Step 3 Choose **NFS** from the **Storage Type** drop-down list.

Step 4 Add the connection information:

- Enter the IPv4 address or hostname of the storage system in the **Host** field.
- Enter the path to your storage area in the **Directory** field.

Step 5 Optionally, check the **Use Advanced Options** check box and enter any required command line options; see Remote Storage Management Advanced Options, on page 893.

Step 6 Under **System Usage**:

- Choose **Use for Backups** to store backups on the designated host.
- Choose **Use for Reports** to store reports on the designated host.
- Enter **Disk Space Threshold** for backup to remote storage. Default is 90%.

Step 7 To test the settings, click **Test**.

Step 8 Click **Save**.

Configuring SMB for Remote Storage

Before you begin

Ensure that your external remote storage system is functional and accessible from your FMC:

- The system recognizes top-level SMB shares, not full file paths. You must use Windows to share the exact directory you want to use.
- Make sure the Windows user you will use to access the SMB share from the FMC has ownership of and read/change access to the share location.

- To ensure security, you should install SMB 2.0 or greater.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **Remote Storage Device**.

Step 3 Choose **SMB** from the **Storage Type** drop-down list.

Step 4 Add the connection information:

- Enter the IPv4 address or hostname of the storage system in the **Host** field.
- Enter the share of your storage area in the **Share** field.
- Optionally, enter the domain name for the remote storage system in the **Domain** field.
- Enter the user name for the storage system in the **Username** field and the password for that user in the **Password** field.

Step 5 Optionally, check the **Use Advanced Options** check box and enter any required command line options; see Remote Storage Management Advanced Options, on page 893.

Step 6 Under **System Usage**:

- Choose **Use for Backups** to store backups on the designated host.
- Choose **Use for Reports** to store reports on the designated host.

Step 7 To test the settings, click **Test**.

Step 8 Click **Save**.

Configuring SSH for Remote Storage

Before you begin

- Ensure that your external remote storage system is functional and accessible from your Firepower Management Center.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **Remote Storage Device**.

Step 3 Choose **SSH** from the **Storage Type** drop-down list.

Step 4 Add the connection information:

- Enter the IP address or host name of the storage system in the **Host** field.
- Enter the path to your storage area in the **Directory** field.

- Enter the storage system's user name in the **Username** field and the password for that user in the **Password** field. To specify a network domain as part of the connection user name, precede the user name with the domain followed by a forward slash (/).
- To use SSH keys, copy the content of the **SSH Public Key** field and place it in your authorized_keys file.

Step 5 Optionally, check the **Use Advanced Options** check box and enter any required command line options; see Remote Storage Management Advanced Options, on page 893.

Step 6 Under System Usage:

- Choose **Use for Backups** to store backups on the designated host.
- Choose **Use for Reports** to store reports on the designated host.

Step 7 If you want to test the settings, you must click **Test**.

Step 8 Click **Save**.

Remote Storage Management Advanced Options

If you select the Network File System (NFS) protocol, Server Message Block (SMB) protocol, or `SSH` to use secure file transfer protocol (SFTP) to store your reports and backups, you can select the **Use Advanced Options** check box to use one of the mount binary options as documented in an NFS, SMB, or SSH mount main page.

If you select SMB or NFS storage type, you can specify the version number of the remote storage in the **Command Line Option** field using the following format:

`vers=`*`version`*

where `version` is the version number of SMB or NFS remote storage you want to use. For example, to select NFSv4, enter `vers=4.0`.

If SMB encryption is enabled for a file server, only SMB version 3.0 clients are allowed to access the file server. To access encrypted SMB file server from the FMC, type the following in the **Command Line Option** field:

`vers=`*`3.0`*

where you select encrypted SMBv3 to copy or save backup files from the FMC to the encrypted SMB file server.

Change Reconciliation

To monitor the changes that users make and ensure that they follow your organization's preferred standard, you can configure the system to send, via email, a detailed report of changes made over the past 24 hours. Whenever a user saves changes to the system configuration, a snapshot is taken of the changes. The change reconciliation report combines information from these snapshots to present a clear summary of recent system changes.

The following sample graphic displays a User section of an example change reconciliation report and lists both the previous value for each configuration and the value after changes. When users make multiple changes

to the same configuration, the report lists summaries of each distinct change in chronological order, beginning with the most recent.

You can view changes made during the previous 24 hours.

Configuring Change Reconciliation

Before you begin

- Configure an email server to receive emailed reports of changes made to the system over a 24 hour period; see Configuring a Mail Relay Host and Notification Address, on page 906 for more information.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **Change Reconciliation**.

Step 3 Check the **Enable** check box.

Step 4 Choose the time of day you want the system to send out the change reconciliation report from the **Time to Run** drop-down lists.

Step 5 Enter email addresses in the **Email to** field.

Tip Once you have added email addresses, click **Resend Last Report** to send recipients another copy of the most recent change reconciliation report.

Step 6 If you want to include policy changes, check the **Include Policy Configuration** check box.

Step 7 If you want to include all changes over the past 24 hours, check the **Show Full Change History** check box.

Step 8 Click **Save**.

Related Topics

Using the Audit Log to Examine Changes, on page 2419

Change Reconciliation Options

The **Include Policy Configuration** option controls whether the system includes records of policy changes in the change reconciliation report. This includes changes to access control, intrusion, system, health, and network discovery policies. If you do not select this option, the report will not show changes to any policies. This option is available on Firepower Management Centers only.

The **Show Full Change History** option controls whether the system includes records of all changes over the past 24 hours in the change reconciliation report. If you do not select this option, the report includes only a consolidated view of changes for each category.

Note The change reconciliation report does not include changes to Firepower Threat Defense interfaces and routing settings.

Policy Change Comments

You can configure the Firepower System to track several policy-related changes using the comment functionality when users modify access control, intrusion, or network analysis policies.

With policy change comments enabled, administrators can quickly assess why critical policies in a deployment were modified. Optionally, you can have changes to intrusion and network analysis policies written to the audit log.

Configuring Comments to Track Policy Changes

You can configure the Firepower System to prompt users for comments when they modify an access control policy, intrusion policy, or network analysis policy. You can use comments to track users' reasons for policy changes. If you enable comments on policy changes, you can make the comment optional or mandatory. The system prompts the user for a comment when each new change to a policy is saved.

Procedure

Step 1 Choose **System** > **Configuration**.

The system configuration options appear in the left navigation panel.

Step 2 Configure the policy comment preferences for any of the following:

- Click **Access Control Preferences** for comment preferences for access control policies.
- Click **Intrusion Policy Preferences** for comment preferences for intrusion policies.
- Click **Network Analysis Policy Preferences** for comment preferences for network analysis policies.

Step 3 You have the following choices for each policy type:

- **Disabled**—Disables change comments.
- **Optional**—Gives users the option to describe their changes in a comment.
- **Required**—Requires users to describe their changes in a comment before saving.

Step 4 Optionally for intrusion or network analysis policy comments:

- Check **Write changes in Intrusion Policy to audit log** to write all intrusion policy changes to the audit log.
- Check **Write changes in Network Analysis Policy to audit log** to write all network analysis policy changes to the audit log.

Step 5 Click **Save**.

Access List

You can limit access to the FMC by IP address and port. By default, the following ports are enabled for any IP address:

- 443 (HTTPS) for web interface access.
- 22 (SSH) for CLI access.

You can also add access to poll for SNMP information over port 161. Because SNMP is disabled by default, you must first enable SNMP before you can add SNMP access rules. For more information, see Configure SNMP Polling, on page 908.

Caution By default, access is not restricted. To operate in a more secure environment, consider adding access for specific IP addresses and then deleting the default **any** option.

Configure an Access List

This access list does not control external database access. See Enabling External Access to the Database, on page 878.

Caution If you delete access for the IP address that you are currently using to connect to the FMC, and there is no entry for "`IP=any port=443`", you will lose access when you save.

To configure access lists for Classic devices, use device platform settings. See Configure Access Lists for Classic Devices, on page 933.

Before you begin

By default, the access list includes rules for HTTPS and SSH. To add SNMP rules to the access list, you must first enable SNMP. For more information, see Configure SNMP Polling, on page 908.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 (Optional) Click **SNMP** to configure SNMP if you want to add SNMP rules to the access list. By default, SNMP is disabled; see Configure SNMP Polling, on page 908.

Step 3 Click **Access List**.

Step 4 To add access for one or more IP addresses, click **Add Rules**.

Step 5 In the **IP Address** field, enter an IP address or address range, or `any`.

Step 6 Choose **SSH**, **HTTPS**, **SNMP**, or a combination of these options to specify which ports you want to enable for these IP addresses.

Step 7 Click **Add**.

Step 8 Click **Save**.

Related Topics

Firepower System IP Address Conventions, on page 15

Audit Logs

The Firepower Management Center records user activity in read-only audit logs. You can review audit log data in several ways:

- Use the web interface: Auditing the System, on page 2415.

  Audit logs are presented in a standard event view where you can view, sort, and filter audit log messages based on any item in the audit view. You can easily delete and report on audit information and you can view detailed reports of the changes that users make.

- Stream audit log messages to the syslog: Stream Audit Logs to Syslog, on page 897..
- Stream audit log messages to an HTTP server: Stream Audit Logs to an HTTP Server, on page 898.

Streaming audit log data to an external server allows you to conserve space on the FMC. Note that sending audit information to an external URL may affect system performance.

Optionally, you can secure the channel for audit log streaming, enable TLS and mutual authentication using TLS certificates; see Audit Log Certificate, on page 900.

Classic devices also maintain audit logs. To stream audit logs from a Classic devices, see Stream Audit Logs from Classic Devices, on page 934.

Stream Audit Logs to Syslog

When this feature is enabled, audit log records appear in the syslog in the following format :

```
Date Time Host [Tag] Sender: User_Name@User_IP, Subsystem, Action
```

Where the local date, time, and originating hostname precede the bracketed optional tag, and the sending device name precedes the audit log message.

For example, if you specify a tag of `FMC-AUDIT-LOG` for audit log messages from your management center, a sample audit log message from your FMC could appear as follows:

```
Mar 01 14:45:24 localhost [FMC-AUDIT-LOG] Dev-MC7000: admin@10.1.1.2, Operations > Monitoring,
Page View
```

If you specify a severity and facility, these values do not appear in syslog messages; instead, they tell the system that receives the syslog messages how to categorize them.

To stream audit logs from Classic devices, use device platform settings: Stream Audit Logs from Classic Devices, on page 934.

Before you begin

Make sure the FMC can communicate with the syslog server. When you save your configuration, the system uses ICMP/ARP and TCP SYN packets to verify that the syslog server is reachable. Then, the system uses port 514/UDP to stream audit logs. If you secure the channel (optional, see Audit Log Certificate, on page 900), you must manually configure port 1470 for TCP.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **Audit Log**.

Step 3 Choose **Enabled** from the **Send Audit Log to Syslog** drop-down menu.

Step 4 The following fields are applicable only for audit logs sent to syslog:

Option	Description
Host	The IP address or the fully qualified name of the syslog server to which you will send audit logs.
Facility	The subsystem that creates the message. Choose a facility described in Syslog Alert Facilities, on page 2079. For example, choose AUDIT.
Severity	The severity of the message. Choose a severity described in Syslog Severity Levels, on page 2080.
Tag	An optional tag to include in audit log syslog messages. Best practice: Enter a value in this field to easily differentiate audit log messages from other, similar syslog messages such as health alerts. For example, if you want all audit log records sent to the syslog to be labeled with `FMC-AUDIT-LOG`, enter `FMC-AUDIT-LOG` in the field.

Step 5 (Optional) To test whether the IP address of the syslog server is valid, click **Test Syslog Server**.

The system sends the following packets to verify whether the syslog server is reachable:

a. ICMP echo request

b. TCP SYN on 443 and 80 ports

c. ICMP time stamp query

d. TCP SYN on random ports

Note If the FMC and syslog server are in the same subnet, ARP is used instead of ICMP.

The system displays the result for the server.

Step 6 Click **Save**.

Stream Audit Logs to an HTTP Server

When this feature is enabled, the appliance sends audit log records to an HTTP server in the following format:

```
Date Time Host [Tag] Sender: User_Name@User_IP, Subsystem, Action
```

Where the local date, time, and originating hostname precede the bracketed optional tag, and the sending appliance or device name precedes the audit log message.

For example, if you specify a tag of `FROMMC`, a sample audit log message could appear as follows:

```
Mar 01 14:45:24 localhost [FROMMC] Dev-MC7000: admin@10.1.1.2, Operations > Monitoring,
Page View
```

To stream audit logs from Classic devices, use device platform settings: Stream Audit Logs from Classic Devices, on page 934.

Before you begin

Make sure the device can communicate with the HTTP server. Optionally, secure the channel; see Audit Log Certificate, on page 900.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **Audit Log**.

Step 3 Optionally, in the **Tag** field, enter the tag name that you want to appear with the message. For example, if you want all audit log records to be preceded with `FROMMC`, enter `FROMMC` in the field.

Step 4 Choose **Enabled** from the **Send Audit Log to HTTP Server** drop-down list.

Step 5 In the **URL to Post Audit** field, designate the URL where you want to send the audit information. Enter a URL that corresponds to a Listener program that expects the HTTP POST variables as listed:

- `subsystem`
- `actor`
- `event_type`
- `message`
- `action_source_ip`
- `action_destination_ip`
- `result`
- `time`
- `tag` (if defined; see Step 3)

Caution To allow encrypted posts, use an HTTPS URL. Sending audit information to an external URL may affect system performance.

Step 6 Click **Save**.

Audit Log Certificate

You can use Transport Layer Security (TLS) certificates to secure communications between Firepower appliances and a trusted audit log server.

Client Certificates (Required)

For *each appliance* (client certificates are unique), you must generate a certificate signing request (CSR), submit it to a Certificate Authority (CA) for signing, then import the signed certificate onto the appliance.

You cannot use the FMC to import audit log certificates onto its managed devices. These certificates are unique to each appliance, and you must log into *each appliance* to import them locally:

- For the FMC, use the local system configuration: Obtain a Signed Audit Log Client Certificate for the FMC, on page 901 and Import an Audit Log Client Certificate into the FMC, on page 902.
- For 7000/8000 series devices, use the local system configuration: Obtain a Signed Client Certificate for Secure Audit Log Streaming on a 7000/8000 Series Device, on page 942.
- For ASA FirePOWER and NGIPSv, generate a CSR with a tool like OpenSSL, then use the CLI to import the signed certificate: **configure audit_cert import**.

Server Certificates (Optional)

For additional security, we recommend you require mutual authentication between Firepower appliances and the audit log server. To accomplish this, load one or more certificate revocation lists (CRLs). You cannot stream audit logs to servers with revoked certificates listed in those CRLs.

Firepower supports CRLs encoded in Distinguished Encoding Rules (DER) format. Note that these are the same CRLs that the system uses to validate HTTPS client certificates for the FMC web interface.

To require valid audit log server certificates, use the FMC web interface:

- For the FMC, use the local system configuration: Require Valid Audit Log Server Certificates, on page 903.
- For Classic devices (including 7000/8000 series), use device platform settings: Require Valid Audit Log Server Certificates for Classic Devices, on page 935.

Securely Stream Audit Logs

If you stream the audit log to a trusted HTTP server or syslog server, you can use Transport Layer Security (TLS) certificates to secure the channel between the FMC and the server. You must generate a unique client certificate for each appliance you want to audit.

To securely stream audit logs to Classic devices, see Stream Audit Logs from Classic Devices, on page 934.

Before you begin

See ramifications of requiring client and server certificates at Audit Log Certificate, on page 900.

Procedure

Step 1 Obtain and install a signed client certificate on the FMC:

a) Obtain a Signed Audit Log Client Certificate for the FMC, on page 901:

Generate a Certificate Signing Request (CSR) from the FMC based on your system information and the identification information you supply.

Submit the CSR to a recognized, trusted certificate authority (CA) to request a signed client certificate.

If you will require mutual authentication between the FMC and the audit log server, the client certificate must be signed by the same CA that signed the server certificate to be used for the connection.

b) After you receive the signed certificate from the certificate authority, import it into the FMC. See Import an Audit Log Client Certificate into the FMC, on page 902.

Step 2 Configure the communication channel with the server to use Transport Layer Security (TLS) and enable mutual authentication.

See Require Valid Audit Log Server Certificates, on page 903.

Step 3 Configure audit log streaming if you have not yet done so.

See Stream Audit Logs to Syslog, on page 897 or Stream Audit Logs to an HTTP Server, on page 898.

Obtain a Signed Audit Log Client Certificate for the FMC

To obtain a certificate for a managed Classic device, see Obtain a Signed Client Certificate for Secure Audit Log Streaming on a 7000/8000 Series Device, on page 942.



Important The **Audit Log Certificate** page is not available on a standby Firepower Management Center in a high availability setup. You cannot perform this task from a standby Firepower Management Center.

The system generates certificate request keys in Base-64 encoded PEM format.

Before you begin

Keep the following in mind:

- You must generate a certificate signing request (CSR) from the device or appliance on which you will install the certificate. (For example, you cannot generate a certificate signing request for Device B from Appliance A.) You must generate a unique certificate signing request from each device and appliance.
- To ensure security, use a globally recognized and trusted Certificate Authority (CA) to sign your certificate.
- If you will require mutual authentication between the appliance and the audit log server, the same Certificate Authority must sign both the client certificate and the server certificate.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **Audit Log Certificate**.

Step 3 Click **Generate New CSR**.

Step 4 Enter a country code in the **Country Name (two-letter code)** field.

Step 5 Enter a state or province postal abbreviation in the **State or Province** field.

Step 6 Enter a **Locality or City**.

Step 7 Enter an **Organization** name.

Step 8 Enter an **Organizational Unit (Department)** name.

Step 9 Enter the fully qualified domain name of the server for which you want to request a certificate in the **Common Name** field.

Note If the common name and the DNS hostname do not match, audit log streaming will fail.

Step 10 Click **Generate**.

Step 11 Open a new blank file with a text editor.

Step 12 Copy the entire block of text in the certificate request, including the `BEGIN CERTIFICATE REQUEST` and `END CERTIFICATE REQUEST` lines, and paste it into a blank text file.

Step 13 Save the file as *`clientname`*`.csr`, where *`clientname`* is the name of the appliance where you plan to use the certificate.

Step 14 Click **Close**.

What to do next

- Submit the certificate signing request to the certificate authority that you selected using the guidelines in the "Before You Begin" section of this procedure.
- When you receive the signed certificate, import it to the appliance; see Import an Audit Log Client Certificate into the FMC, on page 902.

Import an Audit Log Client Certificate into the FMC

In an FMC high availability setup, you *must* use the active peer.

For ASA FirePOWER and NGIPSv, use the CLI to import a signed certificate: **configure audit_cert import**. For 7000/8000 series devices, use the system configuration (**System** > **Configuration**) on the device's local web interface: Obtain a Signed Client Certificate for Secure Audit Log Streaming on a 7000/8000 Series Device, on page 942.

Before you begin

- Obtain a Signed Audit Log Client Certificate for the FMC, on page 901.
- Make sure you are importing the signed certificate for the correct appliance. Each certificate is unique to a specific appliance or device.

- If the signing authority that generated the certificate requires you to trust an intermediate CA, be prepared to provide the necessary certificate chain (or certificate path). The CA that signed the client certificate must be the same CA that signed any intermediate certificates in the certificate chain.

Procedure

Step 1 On the FMC, choose **System** > **Configuration**.

Step 2 Click **Audit Log Certificate**.

Step 3 Click **Import Audit Client Certificate**.

Step 4 Open the client certificate in a text editor, copy the entire block of text, including the `BEGIN CERTIFICATE` and `END CERTIFICATE` lines. Paste this text into the **Client Certificate** field.

Step 5 To upload a private key, open the private key file and copy the entire block of text, including the `BEGIN RSA PRIVATE KEY` and `END RSA PRIVATE KEY` lines. Paste this text into the **Private Key** field.

Step 6 Open any required intermediate certificates, copy the entire block of text for each, and paste it into the **Certificate Chain** field.

Step 7 Click **Save**.

Require Valid Audit Log Server Certificates

The system supports validating audit log server certificates using imported CRLs in Distinguished Encoding Rules (DER) format.

Note If you choose to verify certificates using CRLs, the system uses the same CRLs to validate both audit log server certificates and certificates used to secure the HTTP connection between an appliance and a web browser.

Important You cannot perform this procedure on the standby Firepower Management Center in a high availablity pair.

Before you begin

- Understand the ramifications of requiring mutual authentication and of using certificate revocation lists (CRLs) to ensure that certificates are still valid. See Audit Log Certificate, on page 900.
- Obtain and import the client certificate following the steps in Securely Stream Audit Logs, on page 900 and the topics referenced in that procedure.

Procedure

Step 1 On the FMC, choose **System** > **Configuration**.

Step 2 Click **Audit Log Certificate**.

Step 3 To use Transport Layer Security to securely stream the audit log to an external server, choose **Enable TLS**.

Step 4 If you want to accept server certificates without verification (not recommended):

a) Deselect **Enable Mutual Authentication**.

b) Click **Save** and skip the remainder of this procedure.

Step 5 To verify the certificate of the audit log server, choose **Enable Mutual Authentication**.

Step 6 (If you enabled mutual authentication) To automatically recognize certificates that are no longer valid:

a) Select **Enable Fetching of CRL**.

Note Enabling fetching of the CRL creates a scheduled task to regularly update the CRL or CRLs.

b) Enter a valid URL to an existing CRL file and click **Add CRL**.

Repeat to add up to 25 CRLs.

c) Click **Refresh CRL** to load the current CRL or CRLs from the specified URL or URLs.

Step 7 Verify that you have a valid server certificate generated by the same certificate authority that created the client certificate.

Step 8 Click **Save**.

What to do next

(Optional) Set the frequency of CRL updates. See Configuring Certificate Revocation List Downloads, on page 182.

View the Audit Log Client Certificate on the FMC

You can view the audit log client certificate only for the appliance that you are logged in to. In FMC high availability pairs, you can view the certificate only on the active peer.

To view audit log certificates on Classic devices, use the system configuration on a 7000/8000 series device's web interface, or the CLI: **show audit_cert**.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **Audit Log Certificate**.

Dashboard Settings

Dashboards provide you with at-a-glance views of current system status through the use of widgets: small, self-contained components that provide insight into different aspects of the Firepower System. The Firepower System is delivered with several predefined dashboard widgets.

You can configure the Firepower Management Center so that Custom Analysis widgets are enabled on the dashboard.

Related Topics

About Dashboards, on page 255

Enabling Custom Analysis Widgets for Dashboards

Use Custom Analysis dashboard widgets to create a visual representation of events based on a flexible, user-configurable query.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **Dashboard**.

Step 3 Check the **Enable Custom Analysis Widgets** check box to allow users to add Custom Analysis widgets to dashboards.

Step 4 Click **Save**.

DNS Cache

You can configure the system to resolve IP addresses automatically on the event view pages. You can also configure basic properties for DNS caching performed by the appliance. Configuring DNS caching allows you to identify IP addresses you previously resolved without performing additional lookups. This can reduce the amount of traffic on your network and speed the display of event pages when IP address resolution is enabled.

Configuring DNS Cache Properties

DNS resolution caching is a system-wide setting that allows the caching of previously resolved DNS lookups.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Choose **DNS Cache**.

Step 3 From the **DNS Resolution Caching** drop-down list, choose one of the following:

- **Enabled**—Enable caching.
- **Disabled**—Disable caching.

Step 4 In the **DNS Cache Timeout (in minutes)** field, enter the number of minutes a DNS entry remains cached in memory before it is removed for inactivity.

The default setting is 300 minutes (five hours).

Step 5 Click **Save**.

Related Topics

Configuring Event View Settings, on page 33

Email Notifications

Configure a mail host if you plan to:

- Email event-based reports
- Email status reports for scheduled tasks
- Email change reconciliation reports
- Email data-pruning notifications
- Use email for discovery event, impact flag, correlation event alerting, intrusion event alerting, and health event alerting

When you configure email notification, you can select an encryption method for the communication between the system and mail relay host, and can supply authentication credentials for the mail server if needed. After configuring, you can test the connection.

Configuring a Mail Relay Host and Notification Address

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **Email Notification**.

Step 3 In the **Mail Relay Host** field, enter the hostname or IP address of the mail server you want to use. The mail host you enter **must** allow access from the appliance.

Step 4 In the **Port Number** field, enter the port number to use on the email server.

Typical ports include:

- 25, when using no encryption
- 465, when using SSLv3
- 587, when using TLS

Step 5 Choose an **Encryption Method**:

- **TLS**—Encrypt communications using Transport Layer Security.
- **SSLv3**—Encrypt communications using Secure Socket Layers.
- **None**—Allow unencrypted communication.

Note Certificate validation is not required for encrypted communication between the appliance and mail server.

Step 6 In the **From Address** field, enter the valid email address you want to use as the source email address for messages sent by the appliance.

Step 7 Optionally, to supply a user name and password when connecting to the mail server, choose **Use Authentication**. Enter a user name in the **Username** field. Enter a password in the **Password** field.

Step 8 To send a test email using the configured mail server, click **Test Mail Server Settings**.

A message appears next to the button indicating the success or failure of the test.

Step 9 Click **Save**.

Language Selection

You can use the Language page to specify a different language for the web interface.

Set the Language for the Web Interface

The language you specify here is used for the web interface for every user. You can choose from:

- English
- Chinese (simplified)
- Japanese
- Korean

To set the language for 7000/8000 series devices, use device platform settings: Set the Language for the 7000/8000 Series Web Interface, on page 938.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **Language**.

Step 3 Choose the language you want to use.

Step 4 Click **Save**.

Login Banners

You can use the Login Banner page to specify session, login, or custom message banners for a security appliance or shared policy.

You can use ASCII characters and carriage returns to create a custom login banner. The system does not preserve tab spacing. If your login banner is too large or causes errors, Telnet or SSH sessions can fail when the system attempts to display the banner.

Customize the Login Banner

To customize login banners for Classic devices, use device platform settings. See Customize the Login Banner for Classic Devices , on page 938.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Choose **Login Banner**.

Step 3 In the **Custom Login Banner** field, enter the login banner text you want to use.

Step 4 Click **Save**.

SNMP Polling

You can enable Simple Network Management Protocol (SNMP) polling. This feature supports use of versions 1, 2, and 3 of the SNMP protocol. This feature allows access to the standard management information base (MIB), which includes system details such as contact, administrative, location, service information, IP addressing and routing information, and transmission protocol usage statistics.

Note When selecting SNMP versions for the SNMP protocol, note that SNMPv2 only supports read-only communities and SNMPv3 only supports read-only users. SNMPv3 also supports encryption with AES128.

Enabling SNMP polling does not cause the system to send SNMP traps; it only makes the information in the MIBs available for polling by your network management system.

Configure SNMP Polling

To configure SNMP polling on Classic managed devices, use the device platform settings. See Configure SNMP Polling on Classic Devices, on page 940.

Before you begin

Add SNMP access for each computer you plan to use to poll the system. See Configure an Access List, on page 896.

Note The SNMP MIB contains information that could be used to attack your deployment. We recommend that you restrict your access list for SNMP access to the specific hosts that will be used to poll for the MIB. We also recommend you use SNMPv3 and use strong passwords for network management access.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **SNMP**.

Step 3 From the **SNMP Version** drop-down list, choose the SNMP version you want to use:

- **Version 1** or **Version 2**: Enter a read-only SNMP community name in the **Community String** field, then skip to the end of the procedure.

  Note Do not include special characters (< > / % # & ? ', etc.) in the SNMP community string name.

- **Version 3**: Click **Add User** to display the user definition page. SNMPv3 only supports read-only users and encryption with AES128.

Step 4 Enter a **Username**.

Step 5 Choose the protocol you want to use for authentication from the **Authentication Protocol** drop-down list.

Step 6 Enter the password required for authentication with the SNMP server in the **Authentication Password** field.

Step 7 Re-enter the authentication password in the **Verify Password** field.

Step 8 Choose the privacy protocol you want to use from the **Privacy Protocol** list, or choose **None** to not use a privacy protocol.

Step 9 Enter the SNMP privacy key required by the SNMP server in the **Privacy Password** field.

Step 10 Re-enter the privacy password in the **Verify Password** field.

Step 11 Click **Add**.

Step 12 Click **Save**.

Security Certifications Compliance

Your organization might be required to use only equipment and software complying with security standards established by the U.S. Department of Defense and global certification organizations. The following security certifications standards are supported:

- Common Criteria (CC): a global standard established by the international Common Criteria Recognition Arrangement, defining properties for security products
- Unified Capabilities Approved Products List (UCAPL): a list of products meeting security requirements established by the U.S. Defense Information Systems Agency (DISA)



Note The U.S. Government has changed the name of the Unified Capabilities Approved Products List (UCAPL) to the Department of Defense Information Network Approved Products List (DODIN APL). References to UCAPL in this documentation and the FMC web interface can be interpreted as references to DODIN APL.

- Federal Information Processing Standards (FIPS) 140: a requirements specification for encryption modules

You can enable security certifications compliance in CC mode or UCAPL mode. Enabling security certifications compliance does not guarantee strict compliance with all requirements of the security mode selected. For more information on hardening procedures, refer to the guidelines for this product provided by the certifying entity.

Caution After you enable this setting, you cannot disable it. If you need to take an appliance out of CC or UCAPL mode, you must reimage.

Security Certifications Compliance Characteristics

The following table describes behavior changes when you enable CC or UCAPL mode. (Restrictions on login accounts refers to command line or shell access, not web interface access.)

System Change	CC Mode	UCAPL Mode
FIPS compliance is enabled.	Yes	Yes
The system does not allow remote storage for backups or reports.	Yes	Yes
The system starts an additional system audit daemon.	No	Yes
The system boot loader is secured.	No	Yes
The system applies additional security to login accounts.	No	Yes
The system disables the reboot key sequence Ctrl-Alt-Del.	No	Yes
The system enforces a maximum of ten simultaneous login sessions.	No	Yes
The system supports exporting event data using eStreamer only for Version 6.2.0.3 or subsequent 6.2.0.x patches.	Yes	Yes
The system applies more stringent safeguards for login accounts: • Passwords must be at least fifteen alphanumeric characters of mixed case and must include at least one numeric character. • Passwords cannot be a word that appears in a dictionary or include consecutive repeating characters. • The system locks out a user after three failed login attempts in a row. In this case, the password must be reset by an administrator. • The system stores password history.	No	Yes

Security Certifications Compliance Recommendations

Cisco recommends that you observe the following best practices when using a system with security certifications compliance enabled:

- To enable security certifications compliance in your deployment, enable it first on the Firepower Management Center, then enable it in the same mode on all managed devices.

Caution The Firepower Management Center will not receive event data from a managed device unless both are operating in the same security certifications compliance mode.

- For all users, enable password strength checking and set the minimum password length to the value required by the certifying agency.
- If you are using Firepower Management Centers in a high-availability configuration, configure them both to use the same security certifications compliance mode.
- Do not configure the system to use any of the following features:
  - Email reports, alerts, or data pruning notifications.
  - Nmap Scan, Cisco IOS Null Route, Set Attribute Value, or ISE EPS remediations.
  - Remote storage for backups or reports.
  - Third-party client access to the system database.
  - External notifications or alerts transmitted via email, SNMP trap, or syslog.
  - Audit log messages transmitted to an HTTP server or to a syslog server without using SSL certificates to secure the channel between the appliance and the server.
- You may configure the system to export event data to an external client using eStreamer only for Version 6.2.0.3 and subsequent 6.2.0.x patches.
- Do not enable external authentication using LDAP or RADIUS in deployments using CC mode.
- Do not enable CACs in deployments using CC mode.
- Disable access to the Firepower Management Center and managed devices via the Firepower REST API in deployments using CC or UCAPL mode.
- Enable CACs in deployments using UCAPL mode.

Note The Firepower System does not support CC or UCAPL mode for classic devices in stacks or high availability pairs.

Enable Security Certifications Compliance

This configuration applies to either a Firepower Management Center or a Classic managed device:

- For the Firepower Management Center, this configuration is part of the system configuration.
- For a managed device, you apply this configuration from the FMC as part of a platform settings policy.

In either case, the configuration does not take effect until you save your system configuration changes or deploy the shared platform settings policy.

Caution After you enable this setting, you cannot disable it. If you need to take the appliance out of CC or UCAPL mode, you must reimage.

Before you begin

- We recommend you register all devices that you plan to be part of your deployment to the FMC before enabling security certifications compliance on any appliances.
- You must be an Admin user to perform this task.

Procedure

Step 1 Depending on whether you are configuring an FMC or a managed device:

- FMC: Choose **System** > **Configuration**.
- Classic device: Choose **Devices** > **Platform Settings** and create or edit a Firepower policy.

Step 2 Click **UCAPL/CC Compliance**.

Note Appliances reboot when you enable UCAPL or CC compliance. The FMC reboots when you save the system configuration; managed devices reboot when you deploy configuration changes.

Step 3 To *permanently* enable security certifications compliance on the appliance, you have two choices:

- To enable security certifications compliance in Common Criteria mode, choose **CC** from the drop-down list.
- To enable security certifications compliance in Unified Capabilities Approved Products List mode, choose **UCAPL** from the drop-down list.

Step 4 Click **Save**.

What to do next

- If you have not already, apply Control and Protection licenses to all Classic devices in your deployment.
- If your appliances were updated from versions earlier than Version 5.2.0, enabling security certifications compliance regenerates appliance certificates. After you enable security certifications compliance in the same mode across your deployment, reregister managed devices to the FMC.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Time and Time Synchronization

Synchronizing the system time on your Firepower Management Center (FMC) and its managed devices is essential to successful operation of your Firepower System. We recommend that you specify NTP servers during FMC initial configuration, but you can use the information in this section to establish or change time sychronization settings after intial configuration is complete.

Use a Network Time Protocol (NTP) server to synchronize system time on the FMC and all devices.

Caution Unintended consequences can occur when time is not synchronized between the FMC and managed devices.

To synchronize time on FMC and managed devices, see:

- Recommended: Synchronize Time on the FMC with an NTP Server, on page 913

  This topic provides instructions for configuring your FMC to synchronize with an NTP server or servers and includes links to instructions on configuring managed devices to synchronize with the same NTP server or servers.

- Otherwise: Synchronize Time Without Access to a Network NTP Server, on page 914

  This topic provides instructions for setting the time on your FMC, configuring your FMC to serve as an NTP server, and links to instructions on configuring managed devices to synchronize with the FMC NTP server.

Synchronize Time on the FMC with an NTP Server

Time synchronization among all of the components of your system is critically important.

The best way to ensure proper time synchronization between FMC and all managed devices is to use an NTP server on your network.

The FMC supports NTPv4.

You must have Admin or Network Admin privileges to do this procedure.

Before you begin

Note the following:

- If your FMC and managed devices cannot access a network NTP server, do not use this procedure. Instead, see Synchronize Time Without Access to a Network NTP Server, on page 914.
- Do not specify an untrusted NTP server.
- Connections to NTP servers do not use configured proxy settings.
- Firepower 4100 Series devices and Firepower 9300 devices cannot use this procedure to set the system time. Instead, configure those devices to use the same NTP server(s) that you configure using this procedure. For instructions, see the documentation for your hardware model.

Caution If the FMC is rebooted and your DHCP server sets an NTP server record different than the one you specify here, the DHCP-provided NTP server will be used instead. To avoid this situation, configure your DHCP server to use the same NTP server.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **Time Synchronization**.

Step 3 If **Serve Time via NTP** is **Enabled**, choose **Disabled** to disable the FMC as an NTP server.

Step 4 For the **Set My Clock** option, choose **Via NTP from** and enter the hostname or IP address of an NTP server.

If your organization has corroborative NTP servers, enter multiple NTP servers as a comma-separated list.

Step 5 Click **Save**.

What to do next

Set managed devices to synchronize with the same NTP server or servers:

- Configure device platform settings: Configure NTP Time Synchronization for Threat Defense, on page 976 and Synchronize Time on Classic Devices with an NTP Server, on page 939.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Synchronize Time Without Access to a Network NTP Server

If your devices cannot directly reach the network NTP server, or your organization does not have a network NTP server, a physical-hardware FMC can serve as an NTP server.

Important

- Do not use this procedure unless you have no other NTP server. Instead, use the procedure in Synchronize Time on the FMC with an NTP Server, on page 913.
- Do not use a virtual FMC as an NTP server.

To change the time manually **after** configuring the FMC as an NTP server, you must disable the NTP option, change the time manually, and then re-enable the NTP option.

Procedure

Step 1 Manually set the system time on the FMC:

a) Choose **System** > **Configuration**.

b) Click **Time Synchronization**.

c) If **Serve Time via NTP** is **Enabled**, choose **Disabled**.
d) Click **Save**.
e) For **Set My Clock**, choose **Manually in Local Configuration**.
f) Click **Save**.
g) In the navigation panel at the left side of the screen, click **Time**.
h) Use the **Set Time** drop-down lists to set the time.
i) If the time zone displayed is not UTC, click it and set the time zone to **UTC**.
j) Click **Save**.
k) Click **Done**.
l) Click **Apply**.

Step 2 Set the FMC to serve as an NTP server:

a) In the navigation panel at the left side of the screen, click **Time Synchronization**.
b) For **Serve Time via NTP**, choose **Enabled**.
c) Click **Save**.

Step 3 Set managed devices to synchronize with the FMC NTP server:

a) In the Time Synchronization settings for the platform settings policy assigned to your managed devices, set the clock to synchronize **Via NTP from Management Center**.
b) Deploy the change to managed devices.

For instructions:

- For FTD devices, see Configure NTP Time Synchronization for Threat Defense, on page 976.
- For all other devices, see Synchronize Time on Classic Devices with an NTP Server, on page 939.

About Changing Time Synchronization Settings

- Your Firepower Management Center and its managed devices are heavily dependent on accurate time. The system clock is a system facility that maintains the time of the Firepower System. The system clock is set to Universal Coordinated Time (UTC), which is the primary time standard by which the world regulates clocks and time.

  DO NOT ATTEMPT TO CHANGE THE SYSTEM TIME. Changing the system time zone from UTC is NOT supported, and doing so will require you to reimage the device to recover from an unsupported state.

- If you configure the FMC to serve time using NTP, and then later disable it, the NTP service on managed devices still attempts to synchronize time with the FMC. You must update and redeploy any applicable platform settings policies to establish a new time source.

- To change the time manually **after** configuring the Firepower Management Center as an NTP server, you must disable the NTP option, change the time manually, and then re-enable the NTP option.

View Current System Time, Source, and NTP Server Connection Status

Time settings are displayed on most pages in local time using the time zone you set on the Time Zone page in User Preferences (the default is America/New York), but are stored on the appliance using UTC time.

Restriction The Time Zone function (in User Preferences) assumes that the default system clock is set to UTC time. DO NOT ATTEMPT TO CHANGE THE SYSTEM TIME. Be advised that changing the system time from UTC is NOT supported, and doing so will require you to reimage the device to recover from an unsupported state.

Note To view time and time source information on your 7000- and 8000-Series hardware device, see View System Time for 7000/8000 Series Devices, on page 947.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **Time**.

The current time is displayed using the time zone specified for your account in User Preferences.

If your appliance uses an NTP server: For information about the table entries, see NTP Server Status, on page 916.

NTP Server Status

If you are synchronizing time from an NTP server, you can view connection status on the **Time** page (choose **System > Configuration**).

Table 72: NTP Status

Column	Description
NTP Server	The IP address or name of the configured NTP server.
Status	The status of the NTP server time synchronization: • **Being Used** indicates that the appliance is synchronized with the NTP server. • **Available** indicates that the NTP server is available for use, but time is not yet synchronized. • **Not Available** indicates that the NTP server is in your configuration, but the NTP daemon is unable to use it. • **Pending** indicates that the NTP server is new or the NTP daemon was recently restarted. Over time, its value should change to **Being Used**, **Available**, or **Not Available**. • **Unknown** indicates that the status of the NTP server is unknown.

Column	Description
Offset	The number of milliseconds of difference between the time on the appliance and the configured NTP server. Negative values indicate that the appliance is behind the NTP server, and positive values indicate that it is ahead.
Last Update	The number of seconds that have elapsed since the time was last synchronized with the NTP server. The NTP daemon automatically adjusts the synchronization times based on a number of conditions. For example, if you see larger update times such as 300 seconds, that indicates that the time is relatively stable and the NTP daemon has determined that it does not need to use a lower update increment.

Session Timeouts

Unattended login sessions may be security risks. You can configure the amount of idle time before a user's login session times out due to inactivity.

Note that you can exempt specific web interface users from timeout, for scenarios where you plan to passively, securely monitor the system for long periods of time. Users with the Administrator role, whose complete access to menu options poses an extra risk if compromised, cannot be made exempt from session timeouts.

Configure Session Timeouts

To configure session timeouts for Classic devices, use device platform settings. See Configure Session Timeouts for Classic Devices, on page 940.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **Shell Timeout**.

Step 3 Configure session timeouts:

- Web interface (FMC only): Configure the **Browser Session Timeout (Minutes)**. The default value is `60`; the maximum value is `1440` (24 hours).

  To exempt users from this session timeout, see User Account Login Options, on page 64.

- CLI: Configure the **Shell Timeout (Minutes)** field. The default value is `0`; the maximum value is `1440` (24 hours).

Step 4 Click **Save**.

Vulnerability Mapping

The Firepower System automatically maps vulnerabilities to a host IP address for any application protocol traffic received or sent from that address, when the server has an application ID in the discovery event database and the packet header for the traffic includes a vendor and version.

For any servers which do not include vendor or version information in their packets, you can configure whether the system associates vulnerabilities with server traffic for these vendor and versionless servers.

For example, a host serves SMTP traffic that does not have a vendor or version in the header. If you enable the SMTP server on the Vulnerability Mapping page of a system configuration, then save that configuration to the Firepower Management Center managing the device that detects the traffic, all vulnerabilities associated with SMTP servers are added to the host profile for the host.

Although detectors collect server information and add it to host profiles, the application protocol detectors will not be used for vulnerability mapping, because you cannot specify a vendor or version for a custom application protocol detector and cannot select the server for vulnerability mapping.

Mapping Vulnerabilities for Servers

This procedure requires any Smart License or the Protection classic license.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Choose **Vulnerability Mapping**.

Step 3 You have the following choices:

- To prevent vulnerabilities for a server from being mapped to hosts that receive application protocol traffic without vendor or version information, clear the check box for that server.
- To cause vulnerabilities for a server to be mapped to hosts that receive application protocol traffic without vendor or version information, check the check box for that server.

Tip You can check or clear all check boxes at once using the check box next to **Enabled**.

Step 4 Click **Save**.

Remote Console Access Management

You can use a Linux system console for remote access on supported systems via either the VGA port (which is the default) or the serial port on the physical appliance. Use the Console Configuration page to choose the option most suitable to the physical layout of your organization's Firepower deployment.

On supported physical-hardware-based Firepower systems, you can use Lights-Out Management (LOM) on a Serial Over LAN (SOL) connection on the default (`eth0`) management interface to remotely monitor or manage the system without logging into the management interface of the system. You can perform limited tasks, such as viewing the chassis serial number or monitoring such conditions as fan speed and temperature,

using a command line interface on an out-of-band management connection. For information about the cable connection to support LOM, see the *Firepower Management Center Getting Started Guide* for your hardware model.

You must enable LOM for both the system and the user you want to manage the system. After you enable the system and the user, you use a third-party Intelligent Platform Management Interface (IPMI) utility to access and manage your system.

Configuring Remote Console Settings on the System

You must be an Admin user to perform this procedure.

Before you begin

- Disable Spanning Tree Protocol (STP) on any third-party switching equipment connected to the device's management interface.
- If you plan to enable Lights-Out Management see the *Getting Started Guide* for your appliance for information about installing and using an Intelligent Platform Management Interface (IPMI) utility.

Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **Console Configuration**.

Step 3 Click **Save**.

Step 4 The system displays the following warning: "You will have to reboot your system for these changes to take effect." Click **OK** to reboot now or **Cancel** to reboot later.

What to do next

- If you configured serial access, be sure the rear-panel serial port is connected to a local computer, terminal server, or other device that can support remote serial access over ethernet as described in the *Getting Started Guide* for your FMC model.
- If you configured Lights-Out Management, enable a Lights-Out Management user; see Lights-Out Management User Access Configuration, on page 919.

Lights-Out Management User Access Configuration

You must explicitly grant Lights-Out Management permissions to users who will use the feature. LOM users also have the following restrictions:

- You must assign the Administrator role to the user.
- The username may have up to 16 alphanumeric characters. Hyphens and longer user names are not supported for LOM users.

- A user's LOM password is the same as that user's system password. Cisco recommends that you use a complex, non-dictionary-based password of the maximum supported length for your appliance and change it every three months.
- If LOM is enabled on a Firepower 7110, 7115, 7120, or 7125 device, the password may have up to 16 alphanumeric characters.
- Physical Firepower Management Centers and 8000 Series devices can have up to 13 LOM users. 7000 Series devices can have up to eight LOM users.

Note that if you deactivate, then reactivate, a user with LOM while a that user is logged in, or restore a user from a backup during that user's login session, that user may need to log back into the web interface to regain access to `impitool` commands.

Enabling Lights-Out Management User Access

You must be an Admin user to perform this procedure.

You configure LOM and LOM users on a per-system basis using each system's local web interface. You cannot use the Firepower Management Center to configure LOM on a managed device. Similarly, because users are managed independently per appliance, enabling or creating a LOM-enabled user on the Firepower Management Center does not transfer that capability to users on managed devices.

Procedure

Step 1 Choose **System** > **Users** > **Users**.

Step 2 To grant LOM user access to an existing user, click **Edit** ( ) next to a user name in the list.

Step 3 Under **User Configuration**, enable the Administrator role.

Step 4 Check the **Allow Lights-Out Management Access** check box.

Step 5 Click **Save**.

Serial Over LAN Connection Configuration

You use a third-party IPMI utility on your computer to create a Serial Over LAN connection to the appliance. If your computer uses a Linux-like or Mac environment, use IPMItool; for Windows environments, you can use IPMIutil or IPMItool, depending on your Windows version.

Note Cisco recommends using IPMItool version 1.8.12 or greater.

Linux

IPMItool is standard with many distributions and is ready to use.

Mac

You must install IPMItool on a Mac. First, confirm that your Mac has Apple's XCode Developer tools installed, making sure that the optional components for command line development are installed (UNIX Development and System Tools in newer versions, or Command Line Support in older versions). Then you can install macports and the IPMItool. Use your favorite search engine for more information or try these sites:

```
https://developer.apple.com/technologies/tools/
http://www.macports.org/
http://github.com/ipmitool/ipmitool/
```

Windows

For Windows Versions 10 and greater with Windows Subsystem for Linux (WSL) enabled, as well as some older versions of Windows Server, you can use IPMItool. Otherwise, you must compile IPMIutil on your Windows system; you can use IPMIutil itself to compile. Use your favorite search engine for more information or try this site:

```
http://ipmiutil.sourceforge.net/man.html#ipmiutil
```

Understanding IPMI Utility Commands

Commands used for IPMI utilities are composed of segments as in the following example for IPMItool on Mac:

```
ipmitool -I lanplus -H IP_address -U user_name command
```

where:

- `ipmitool` invokes the utility.
- `-I lanplus` specifies to use an encrypted IPMI v2.0 RMCP+ LAN Interface for the session.
- `-H IP_address` indicates the IP address you have configured for Lights-Out Management on the appliance you want to access.
- `-U user_name` is the name of an authorized remote session user.
- *`command`* is the name of the command you want to use.



Note Cisco recommends using IPMItool version 1.8.12 or greater.

The same command for IMPIutil on Windows looks like this:

```
ipmiutil command -V 4 -J 3 -N IP_address -Uuser_name
```

This command connects you to the command line on the appliance where you can log in as if you were physically present at the appliance. You may be prompted to enter a password.

Configuring Serial Over LAN with IPMItool

You must be an Admin user with LOM access to perform this procedure.

Procedure

Using IPMItool, enter the following command, and a password if prompted:

```
ipmitool -I lanplus -H IP_address -U user_name sol activate
```

Configuring Serial Over LAN with IPMIutil

You must be an Admin user with LOM access to perform this procedure.

Procedure

Using IPMIutil, enter the following command, and a password if prompted:

```
ipmiutil -J 3 -N IP_address -U username sol -a
```

Lights-Out Management Overview

Lights-Out Management (LOM) provides the ability to perform a limited set of actions over an SOL connection on the default (`eth0`) management interface without the need to log into the system. You use the command to create a SOL connection followed by one of the LOM commands. After the command is completed, the connection ends. Note that not all power control commands are valid on 70xx Family devices.

Note The baseboard management controller (BMC) for a Firepower 71xx, Firepower 82xx, or a Firepower 83xx device is only accessible via 1 Gbps link speeds when the host is powered on. When the device is powered down, the BMC can only establish Ethernet link at 10 and 100 Mbps. Therefore if LOM is being used to remotely power the device, connect the device to the network using 10 and 100 Mbps link speeds only.

Caution In rare cases, if your computer is on a different subnet than the system's management interface and the system is configured for DHCP, attempting to access LOM features can fail. If this occurs, you can either disable and then re-enable LOM on the system, or use a computer on the same subnet as the system to ping its management interface. You should then be able to use LOM.

Caution Cisco is aware of a vulnerability inherent in the Intelligent Platform Management Interface (IPMI) standard (CVE-2013-4786). Enabling Lights-Out Management (LOM) on an system exposes this vulnerability. To mitigate this vulnerability, deploy your systems on a secure management network accessible only to trusted users and use a complex, non-dictionary-based password of the maximum supported length for your system and change it every three months. To prevent exposure to this vulnerability, do not enable LOM.

If all attempts to access your system have failed, you can use LOM to restart your system remotely. Note that if a system is restarted while the SOL connection is active, the LOM session may disconnect or time out.

Caution Do **not** restart your system unless it does not respond to any other attempts to restart. Remotely restarting does not gracefully reboot the system and you may lose data.

Table 73: Lights-Out Management Commands

IPMItool	IPMIutil	Description
(not applicable)	`-V 4`	Enables admin privileges for the IPMI session
`-I lanplus`	`-J 3`	Enables encryption for the IPMI session
`-H` *`hostname/IP address`*	`-N` *`nodename/IP address`*	Indicates the LOM IP address or hostname for the
`-U`	`-U`	Indicates the username of an authorized LOM acc
`sol activate`	`sol -a`	Starts the SOL session
`sol deactivate`	`sol -d`	Ends the SOL session
`chassis power cycle`	`power -c`	Restarts the appliance (not valid on 70xx Family d
`chassis power on`	`power -u`	Powers up the appliance
`chassis power off`	`power -d`	Powers down the appliance (not valid on 70xx Fan
`sdr`	`sensor`	Displays appliance information, such as fan speeds temperatures

For example, to display a list of appliance information, the IPMItool command is:

```
ipmitool -I lanplus -H IP_address -U user_name sdr
```

Note Cisco recommends using IPMItool version 1.8.12 or greater.

The same command with the IPMIutil utility is:

```
ipmiutil sensor -V 4 -J 3 -N IP_address -U user_name
```

Configuring Lights-Out Management with IPMItool

You must be an Admin user with LOM access to perform this procedure.

Procedure

Enter the following command for IPMItool and a password if prompted:

```
ipmitool -I lanplus -H IP_address -U user_name command
```

Configuring Lights-Out Management with IPMIutil

You must be an Admin user with LOM access to perform this procedure.

Procedure

Enter the following command for IPMIutil and a password if prompted:

```
ipmiutil -J 3 -N IP_address -U username command
```

REST API Preferences

The Firepower REST API provides a lightweight interface for third-party applications to view and manage appliance configuration using a REST client and standard HTTP methods. For more information on the Firepower REST API, see the *Firepower REST API Quick Start Guide*.

By default, the Firepower Management Center allows requests from applications using the REST API. You can configure the Firepower Management Center to block this access.

Enabling REST API Access

Note In deployments using the FMC high availability, this feature is available only in the active FMC.

Procedure

Step 1 Choose **System** > **Configuration**

Step 2 Click **REST API Preferences**.

Step 3 To enable or disable REST API access to the FMC, check or uncheck the **Enable REST API** check box.

Step 4 Click **Save**.

Step 5 Access the REST API Explorer at:

```
https://<management_center_IP_or_name>:<https_port>/api/api-explorer
```

VMware Tools and Virtual Systems

VMware Tools is a suite of performance-enhancing utilities intended for virtual machines. These utilities allow you to make full use of the convenient features of VMware products. Firepower virtual appliances running on VMware support the following plugins:

- guestInfo
- powerOps
- timeSync
- vmbackup

You can also enable VMware Tools on all supported versions of ESXi. For a list of supported versions, see the Cisco Firepower NGIPSv Quick Start Guide for VMware. For information on the full functionality of VMware Tools, see the VMware website (http://www.vmware.com/).

Enabling VMware Tools on the Firepower Management Center for VMware

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Firepower Management Center	Global only	Admin

Because NGIPSv does not have a web interface, you must use the CLI to enable VMware Tools on that platform; see the Cisco Firepower NGIPSv Quick Start Guide for VMware.

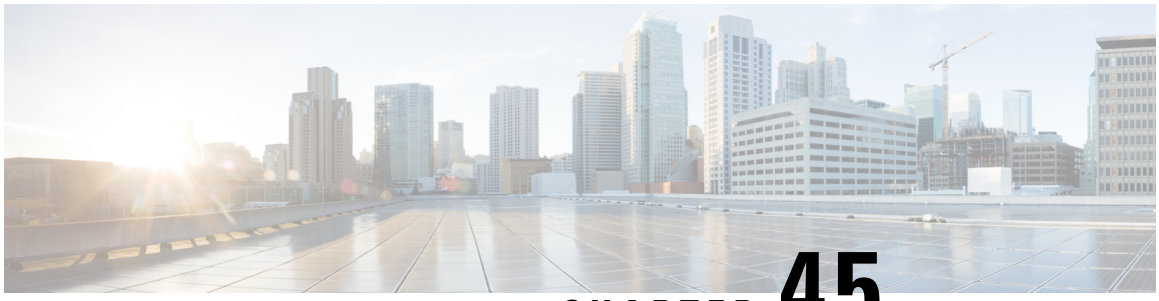
Procedure

Step 1 Choose **System** > **Configuration**.

Step 2 Click **VMware Tools**.

Step 3 Click **Enable VMware Tools**.

Step 4 Click **Save**.



CHAPTER 45

Platform Settings Policies

The following topics explain platform settings policies and how to deploy them to managed devices:

- Introduction to Platform Settings, on page 927
- Requirements and Prerequisites for Platform Settings Policies, on page 928
- Managing Platform Settings Policies, on page 928
- Create a Platform Settings Policy, on page 929
- Setting Target Devices for a Platform Settings Policy, on page 929

Introduction to Platform Settings

A platform settings policy is a shared set of features or parameters that define the aspects of a managed device that are likely to be similar to other managed devices in your deployment, such as time settings and external authentication.

A shared policy makes it possible to configure multiple managed devices at once, which provides consistency in your deployment and streamlines your management efforts. Any changes to a platform settings policy affects all the managed devices where you applied the policy. Even if you want different settings per device, you must create a shared policy and apply it to the desired device.

For example, your organization's security policies may require that your appliances have a "No Unauthorized Use" message when a user logs in. With platform settings, you can set the login banner once in a platform settings policy.

You can also benefit from having multiple platform settings policies on a Firepower Management Center. For example, if you have different mail relay hosts that you use under different circumstances or if you want to test different access lists, you can create several platform settings policies and switch between them, rather than editing a single policy.

Related Topics

Configure Platform Settings for Classic Devices, on page 933

System Configuration Settings, on page 869

Requirements and Prerequisites for Platform Settings Policies

Model Support

Any, but you must create the correct type of policy for the target devices:

- **Firepower Settings** to create a shared policy for Classic managed devices: ASA FirePOWER, NGIPSv, 7000 & 8000 Series.
- **Threat Defense Settings** to create a shared policy for Firepower Threat Defense managed devices.

Supported Domains

Any

User Roles

Admin

Access Admin

Network Admin

Managing Platform Settings Policies

Use the Platform Settings page (**Devices** > **Platform Settings**) to manage platform settings policies. This page indicates the type of device for each policy. The Status column shows the device targets for the policy.

Procedure

Step 1 Choose **Devices** > **Platform Settings**.

Step 2 Manage your platform settings policies:

- Create — To create a new platform settings policy, click **New Policy**; see Create a Platform Settings Policy, on page 929.
- Copy — To copy a platform settings policy, click **Copy** ( ).
- Edit — To modify the settings in an existing platform settings policy, click **Edit** ( ).
- Delete — To delete a policy that is not in use, click **Delete** ( ), then confirm your choice.

Caution You should not delete a policy that is the last deployed policy on any of its target devices, even if it is out of date. Before you delete the policy completely, it is good practice to deploy a different policy to those targets.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Create a Platform Settings Policy

Platform settings for Firepower Threat Defense devices differ from platform settings for Classic devices. When you create a new platform settings policy you must choose a type: *Firepower* (for Classic managed devices) or *Threat Defense* (for FTD devices).

Procedure

Step 1 Choose **Devices** > **Platform Settings**.

Step 2 Click **New Policy**.

Step 3 Choose a device type from the drop-down list:

- **Firepower Settings** to create a shared policy for Classic managed devices.
- **Threat Defense Settings** to create a shared policy for Firepower Threat Defense managed devices.

Step 4 Enter a **Name** for the new policy and optionally, a **Description**.

Step 5 Optionally, choose the **Available Devices** where you want to apply the policy and click **Add to Policy** (or drag and drop) to add the selected devices. You can enter a search string in the **Search** field to narrow the list of devices.

Step 6 Click **Save**.
The system creates the policy and opens it for editing.

Step 7 Configure the platform settings based on the device platform type:

- For Firepower Settings, see Platform Settings for Classic Devices, on page 931.
- For Threat Defense Settings, see Platform Settings for Firepower Threat Defense, on page 949.

Step 8 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Setting Target Devices for a Platform Settings Policy

You can add targeted devices at the same time you create a new policy, or you can change them later.

Procedure

Step 1 Choose **Devices** > **Platform Settings**.

Step 2 Click **Edit** ( ) next to the platform settings policy that you want to edit.

Step 3 Click **Policy Assignment**.

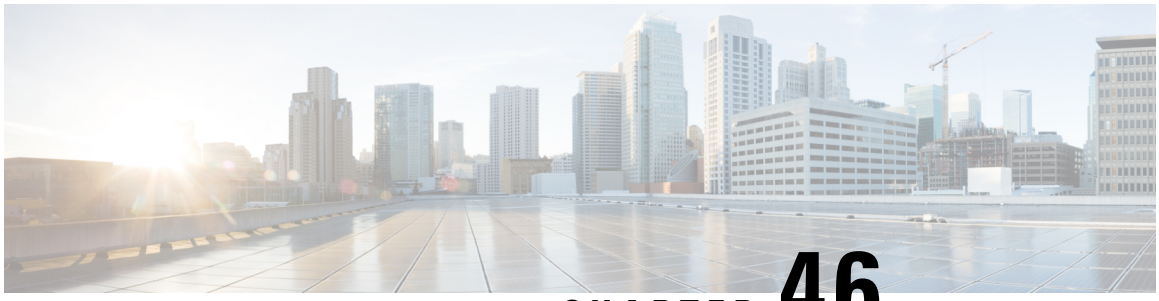
Step 4 Do any of the following:

- To assign a device, stack, high-availability pair, or device group to the policy, select it in the **Available Devices** list and click **Add to Policy**. You can also drag and drop.
- To remove a device assignment, click **Delete** ( ) next to a device, stack, high-availability pair, or device group in the **Selected Devices** list.

Step 5 Click **OK**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.



CHAPTER 46

Platform Settings for Classic Devices

The following topics explain Firepower platform settings and how to configure them on Classic devices:

- About Platform Settings for Classic Devices, on page 931
- Requirements for Platform Settings for Classic Devices, on page 932
- Configure Platform Settings for Classic Devices, on page 933
- Local System Configuration for 7000/8000 Series Devices, on page 941

About Platform Settings for Classic Devices

Platform settings for managed devices are policy-based so that you can apply the same configuration to multiple devices. Use a *Firepower* platform settings policy with Classic devices:

- 7000/8000 series devices
- ASA FirePOWER modules
- NGIPSv

Note that for the FMC, many of these settings are handled in the *system configuration;* see System Configuration, on page 867.

Table 74: Firepower Platform Settings for Classic Devices

Platform Setting	Description	See
Access List	Control which computers can access the system on specific ports.	Configure Access Lists for Classic Devices, on page 933
Audit Log	Configure the system to send an audit log to an external host.	Stream Audit Logs from Classic Devices, on page 934
Audit Log Certificate	As part of audit log secure streaming, require mutual authentication between Classic devices and the audit log server.	Require Valid Audit Log Server Certificates for Classic Devices, on page 935

Platform Setting	Description	See
External Authentication	Set the default user role for any 7000/8000 series device user who is authenticated by an external RADIUS, LDAP or Microsoft Active Directory repository.	Enable External Authentication to 7000/8000 Series Devices, on page 936
Language	Specify a different language for the web interface on a 7000/8000 series device.	Set the Language for the 7000/8000 Series Web Interface, on page 938
Login Banner	Create a custom login banner that appears when users log in.	Customize the Login Banner for Classic Devices , on page 938
Shell Timeout	Configure the amount of idle time, in minutes, before a user's login session times out due to inactivity.	Configure Session Timeouts for Classic Devices, on page 940
SNMP	Enable Simple Network Management Protocol (SNMP) polling.	Configure SNMP Polling on Classic Devices, on page 940
Time Synchronization	Manage time synchronization on the system.	Synchronize Time on Classic Devices with an NTP Server, on page 939
UCAPL/CC Compliance	Enable compliance with specific requirements set out by the United States Department of Defense.	Enable Security Certifications Compliance, on page 911

Requirements for Platform Settings for Classic Devices

License Requirements

None.

Model Requirements

You can apply a Firepower platform settings policy to any Classic device.

Some platform settings apply only to 7000/8000 series devices because those devcies have a web interface: external authentication settings, display language, session timeouts, and so on. Applying these settings to ASA FirePOWER or NGIPSv has no effect.

You can also log into the local web interface on 7000/8000 series devices for non-policy based system configurations. See Local System Configuration for 7000/8000 Series Devices, on page 941.

Domain Requirements

None.

You can apply a Firepower platform setting policy at any Domain level.

Configure Platform Settings for Classic Devices

Platform settings for managed devices are policy-based so that you can apply the same configuration to multiple devices. Use a Firepower platform settings policy with Classic devices.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit a Firepower policy.

See About Platform Settings for Classic Devices, on page 931 and Create a Platform Settings Policy, on page 929.

Step 2 Choose the **Available Devices** where you want to deploy the policy by clicking **Policy Assignment**.

Step 3 Click **Add to Policy** (or drag and drop) to add the selected devices.

Step 4 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configure Access Lists for Classic Devices

By default, access to Firepower devices is not restricted. Port 22 (SSH) is open for CLI access. For 7000/8000 series devices, port 443 (HTTPS) is also open for web interface access.

To operate in a more secure environment, consider adding access for specific IP addresses. You can also add access to poll for SNMP information over port 161.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit a Firepower policy.

Step 2 Click **Access List**.

Step 3 To add access for one or more IP addresses, click **Add Rules**.

Step 4 In the **IP Address** field, enter an IP address or address range, or `any`.

Step 5 Choose **SSH**, **HTTPS**, **SNMP**, or a combination of these options to specify which ports you want to enable for these IP addresses.

Step 6 Click **Add**.

Step 7 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Stream Audit Logs from Classic Devices

Firepower appliances generate records (or *audit logs*) of user interactions. You can stream these audit logs to a syslog or HTTP server. Note that sending audit information to an external URL may affect system performance.



Tip On 7000/8000 series devices, you can also review audit logs on the device's web interface: Auditing the System, on page 2415.

Optionally, you can use Transport Layer Security (TLS) certificates to secure communications between Firepower devices and a trusted audit log server. For *each device* (client certificates are unique), you must generate a certificate signing request (CSR), submit it to a Certificate Authority (CA) for signing, then import the signed certificate onto the device. You cannot use the FMC to import audit log certificates onto its managed devices. These certificates are unique to each device, and you must log into *each device* to import them.

To ensure security, use a globally recognized and trusted CA. The same CA must sign:

- Both the client certificate and the server certificate, if you plan to require mutual authentication between the device and the audit log server.
- Any intermediate certificates in the certificate chain. If the signing CA requires you to trust an intermediate CA, you must provide the necessary certificate chain (or certificate path).

Audit logs have the following format:

```
timestamp host [tag] appliance_name: username@ip_address, subsystem, action
```

For example:

```
Mar 01 14:45:24 localhost [FIREPOWER] MyFirepowerAppliance: admin@10.1.1.2, System >
Configuration, Page View
```

Note that the tag is optional and user-configurable. Syslog events also have an optional facility and severity..

Before you begin

Make sure your devices can communicate with the server or servers where you plan to stream audit logs. For syslog streaming, the system uses port 7/UDP to verify that the syslog server is reachable when you save the configuration. Then, the system uses port 514/UDP to stream audit logs. If you secure the channel, you must manually configure port 1470 for TCP.

Procedure

Step 1 (Optional) Set up secure communications with the audit log server.

For ASA FirePOWER and NGIPSv, you can generate a CSR with a tool like OpenSSL, then use the CLI to import the signed certificate: **configure audit_cert import**. For 7000/8000 series devices, use the system configuration (**System** > **Configuration**) on the device's web interface: Obtain a Signed Client Certificate for Secure Audit Log Streaming on a 7000/8000 Series Device, on page 942.

To verify that the certificate imported correctly, use the 7000/8000 series device's web interface, or the CLI: **show audit_cert**.

Step 2 On the FMC, choose **Devices** > **Platform Settings** and create or edit a Firepower policy.

Step 3 Click **Audit Log** to configure audit log streaming.

Syslog streaming:

a) Set **Send Audit Log to Syslog** to **Enabled**.
b) Provide **Host** information for the syslog server: IP address or fully qualified name.
c) Choose a **Facility** (Syslog Alert Facilities, on page 2079) and **Severity** (Syslog Severity Levels, on page 2080).

HTTP streaming:

a) Set **Send Audit Log to HTTP Server** to **Enabled**.
b) Provide a **URL to Post Audit** where you want to send audit logs. HTTPS is supported.

The URL must correspond to a Listener program that expects the following HTTP POST variables: `subsystem`, `actor`, `event_type`, `message`, `action_source_ip`, `action_destination_ip`, `result`, `time`, `tag` (if provided).

Step 4 (Optional) Enter a **Tag** in include in each message. For example, you might want to tag Firepower audit logs with **`FIREPOWER`**.

Step 5 Click **Save**.
If you configured syslog streaming, the system verifies that the syslog server is reachable.

What to do next

- (Optional) If you configured secure communications, we recommend you also require mutual authentication between the device and the audit log server: Require Valid Audit Log Server Certificates for Classic Devices, on page 935.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Require Valid Audit Log Server Certificates for Classic Devices

For additional security, we recommend you require mutual authentication between Firepower appliances and the audit log server. To accomplish this, load one or more certificate revocation lists (CRLs). You cannot stream audit logs to servers with revoked certificates listed in those CRLs.

Firepower supports CRLs encoded in Distinguished Encoding Rules (DER) format. Note that these are the same CRLs that the system uses to validate HTTPS client certificates for the FMC web interface.

Before you begin

Obtain and import a signed client certificate onto each device.

- For ASA FirePOWER and NGIPSv, you can generate a CSR with a tool like OpenSSL, then use the CLI to import the signed certificate: **configure audit_cert import**.
- For 7000/8000 series devices, use the system configuration (**System** > **Configuration**) on the device's web interface: Obtain a Signed Client Certificate for Secure Audit Log Streaming on a 7000/8000 Series Device, on page 942.

Use a globally recognized and trusted CA. The same CA must sign the client certificates you imported *and* the server certificate you will require with this procedure.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit a Firepower policy.

Step 2 Click **Audit Log Certificate**.

Step 3 Select **Enable TLS**, then **Enable Mutual Authentication**.

We recommend you enable mutual authentication. If you do not, the device will accept server certificates without verification.

Step 4 Select **Enable Fetching of CRL**, provide the URL to a CRL file, and click **Add CRL**.

You can add up to 25 CRLs. When you deploy, the system will schedule CRL updates. To set the update frequency, see Configuring Certificate Revocation List Downloads, on page 182.

Step 5 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Enable External Authentication to 7000/8000 Series Devices

Use device platform settings to allow users of 7000/8000 series devices to authenticate to an LDAP or RADIUS server, rather than using the local database.

Before you begin

Configure external authentication objects. See External Authentication, on page 71.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit a Firepower policy.

Step 2 Click **External Authentication**.

Step 3 From the **Status** drop-down list, choose **Enabled**.

Step 4 From the **Default User Role** drop-down list, choose user roles to define the default permissions you want to grant to externally authenticated users.

Step 5 If you want to use the external server to authenticate CLI or shell access accounts, choose **Enabled** from the **Shell Authentication** drop-down list.

Step 6 If you want to enable CAC authentication and authorization, choose an available CAC authentication object from the **CAC Authentication** drop-down list.

For more information, see CAC Authentication, on page 74.

Step 7 Check the check boxes next to the each external authentication object that you want to use. If you enable more than 1 object, then users are checked against servers in the order specified. See the next step to reorder servers.

If you enable shell authentication, you must enable an external authentication object that includes a **Shell Access Filter**. CLI/shell access users can only authenticate against the server whose authentication object is highest in the list.

If you need both CLI and CAC authentication, you must use separate authentication objects for each purpose.

Step 8 (Optional) Use the up and down arrows to change the order in which authentication servers are accessed when an authentication request occurs.

Step 9 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

About External Authentication for 7000/8000 Series Devices

If you create an authentication object referencing an external authentication server, you can enable external authentication to let users logging into the managed device authenticate to that server, rather than using the local database.

When you enable external authentication, the system verifies the user credentials against users on an LDAP or RADIUS server. In addition, if a user has local, internal authentication enabled and the user credentials are not found in the internal database, the system then checks the external server for a set of matching credentials. If a user has the same username on multiple systems, all passwords across all servers work. Note, however, that if authentication fails on the available external authentication servers, the system does not revert to checking the local database.

When you enable external authentication, you can set the default user role for any user whose account is externally authenticated. You can select multiple roles, as long as those roles can be combined. For example, if you enable external authentication that retrieves only users in the Network Security group in your company, you may set the default user role to include the Security Analyst role so users can access collected event data without any additional user configuration on your part. However, if your external authentication retrieves records for other personnel in addition to the security group, you would probably want to leave the default role unselected.

If no access role is selected, users can log in but cannot access any functionality. After a user attempts to log in, their account is listed on the user management page (**System** > **Users**), where you can edit the account settings to grant additional permissions.

Tip If you configure the system to use one user role and apply the policy, then later modify the configuration to use different default user roles, any user accounts created before the modification retain the first user role until you modify the accounts, or delete and recreate them.

If you want to specify the set of users who can authenticate against the LDAP server for CLI/shell access or for CAC authentication and authorization, you must create separate authentication objects for each and enable the objects separately.

If a user with internal authentication attempts to log in, the system first checks if that user is in the local user database. If the user exists, the system then checks the username and password against the local database. If a match is found, the user logs in successfully. If the login fails, however, and external authentication is enabled, the system checks the user against each external authentication server in the authentication order

shown in the configuration. If the username and password match results from an external server, the system changes the user to an external user with the default privileges for that authentication object.

If an external user attempts to log in, the system checks the username and password against the external authentication server. If a match is found, the user logs in successfully. If the login fails, the user login attempt is rejected. External users cannot authenticate against the user list in the local database. If the user is a new external user, an external user account is created in the local database with the default privileges from the external authentication object.

Set the Language for the 7000/8000 Series Web Interface

The language you specify here is used for the web interface for every user who logs in. You can choose from:

- English
- Chinese (simplified)
- Japanese
- Korean

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit a Firepower policy.

Step 2 Click **Language**.

Step 3 Choose the language you want to use.

Step 4 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Customize the Login Banner for Classic Devices

You can customize the CLI login banner for Classic devices. For 7000/8000 series devices, the login banner also appears in the web interface. Note that if the banner is too large or causes errors, CLI sessions can fail when the system attempts to display the banner.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit a Firepower policy.

Step 2 Choose **Login Banner**.

Step 3 In the **Custom Login Banner** field, enter the login banner text you want to use.

The system will not preserve tab spacing.

Step 4 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Synchronize Time on Classic Devices with an NTP Server

Synchronizing the system time on your FMC and all its managed devices is essential to successful operations. If your deployment includes the Firepower Threat Defense devices, see Configure NTP Time Synchronization for Threat Defense, on page 976.

The device supports NTPv4.

Caution Unintended consequences can occur when time is not synchronized between the FMC and managed devices.

After you deploy, it may take a few minutes for managed devices to synchronize with the configured NTP servers.

Before you begin

Make sure the device can communicate with the NTP server or servers you plan to use. You can either:

- (Recommended.) Use the same NTP servers as the FMC: Synchronize Time on the FMC with an NTP Server, on page 913.

  If you choose this option, the device gets its time directly from the configured NTP server. If the device's configured NTP servers are not reachable for any reason, it synchronizes its time with the FMC.

- If your device cannot reach an NTP server or your organization does not have one, you must use the **Via NTP from Management Center** option discussed in the following proecedure.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit a Firepower policy.

Step 2 Click **Time Synchronization**.

Step 3 Specify how time is synchronized:

- **Via NTP from**: If your Firepower Management Center is using NTP servers on the network, select this option and enter the fully-qualified DNS name (such as ntp.example.com), or IPv4 or IPv6 address, of the same NTP servers you specified in **System > Configuration > Time Synchronization**. If the NTP servers are not reachable, the Firepower Management Center acts as an NTP server.
- **Via NTP from Management Center**: (Default). The managed device gets time from the NTP servers you configured for the Firepower Management Center (except for authenticated NTP servers) and synchronizes time with those servers directly. However, if any of the following are true, the managed device synchronizes time from the Firepower Management Center:
  - The Firepower Management Center's NTP servers are not reachable by the device.

- The Firepower Management Center has no unauthenticated servers.

Step 4 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configure Session Timeouts for Classic Devices

Unattended login sessions may be security risks. You can configure the amount of idle time before a user's login session times out due to inactivity. The maximum value is 24 hours, or 1440 minutes.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit a Firepower policy.

Step 2 Click **Shell Timeout**.

Step 3 Configure session timeouts:

- Web interface (7000/8000 series only): Enter a **Browser Session Timeout (Minutes)**.

  You can exempt specific web interface users from timeout, for scenarios where you plan to passively, securely monitor the system for long periods of time. For more information, see User Account Login Options, on page 64.

- CLI: Enter a **Shell Timeout (Minutes)**.

Step 4 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configure SNMP Polling on Classic Devices

Simple Network Management Protocol (SNMP) polling allows access to the standard management information base (MIB) on Firepower devices, which includes system details such as contact, administrative, location, service information, IP addressing and routing information, and transmission protocol usage statistics. Additional MIBs for 7000/8000 series devices include statistics on traffic passing through physical interfaces, logical interfaces, virtual interfaces, ARP, NDP, virtual bridges, and virtual routers. Note that enabling SNMP polling does not cause the system to send SNMP traps; it only makes the information in the MIBs available for polling by your network management system.

The system supports SNMPv1, v2, and v3. SNMPv2 only supports read-only communities and SNMPv3 only supports read-only users. SNMPv3 also supports encryption with AES128.

Before you begin

Add SNMP access for each computer you plan to use to poll the system. See Configure Access Lists for Classic Devices, on page 933.

Note The SNMP MIB contains information that could be used to attack your deployment. We recommend that you restrict your access list for SNMP access to the specific hosts that will be used to poll for the MIB. We also recommend you use SNMPv3 and use strong passwords for network management access.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit a Firepower policy.

Step 2 Click **SNMP**.

Step 3 From the **SNMP Version** drop-down list, choose the SNMP version you want to use:

- **Version 1** or **Version 2**: Enter a read-only SNMP community name in the **Community String** field, then skip to the end of the procedure.

  Note Do not include special characters (< > / % # & ? ', etc.) in the SNMP community string name.

- **Version 3**: Click **Add User** to display the user definition page. SNMPv3 only supports read-only users and encryption with AES128.

Step 4 Enter a **Username**.

Step 5 Choose the protocol you want to use for authentication from the **Authentication Protocol** drop-down list.

Step 6 Enter the password required for authentication with the SNMP server in the **Authentication Password** field.

Step 7 Re-enter the authentication password in the **Verify Password** field.

Step 8 Choose the privacy protocol you want to use from the **Privacy Protocol** list, or choose **None** to not use a privacy protocol.

Step 9 Enter the SNMP privacy key required by the SNMP server in the **Privacy Password** field.

Step 10 Re-enter the privacy password in the **Verify Password** field.

Step 11 Click **Add**.

Step 12 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Local System Configuration for 7000/8000 Series Devices

You can log into the local web interface on 7000/8000 series devices for non-policy based system configurations. Many of these configurations parallel FMC system configurations, and are documented in the FMC system configuration chapter: System Configuration, on page 867.

Table 75: Local System Configurations for 7000/8000 Series Devices

System Configuration	Description	See
Audit Log Certificate	As part of audit log secure streaming, obtain and import a signed client certificate for 7000/8000 series devices.	Obtain a Signed Client Certificate for Secure Audit Log Streaming on a 7000/8000 Series Device, on page 942
Change Reconciliation	Send a detailed report of changes to the system over the last 24 hours.	Change Reconciliation, on page 893
Console Configuration	Configure console access via VGA or serial port, or via Lights-Out Management (LOM).	Remote Console Access Management, on page 918
HTTPS Certificate	Request an HTTPS server certificate, if needed, from a trusted authority and upload certificates to the system.	HTTPS Certificates, on page 871
Information	View current information about the appliance and edit the display name.	Appliance Information, on page 870
Management Interfaces	Change options such as the IP address, hostname, and proxy settings of the appliance.	Configure Management Interfaces on a 7000/8000 Series Device, on page 944
Process	Shut down, reboot, or restart Firepower processes.	Shut Down or Restart a 7000/8000 Series Device, on page 947
Prohibit Packet Transfer	Disable sending packet data from 7000/8000 series devices to the FMC in a low-bandwidth deployment.	Prohibit Packet Transfer to FMC, on page 944
Time	View the current time settings.	View System Time for 7000/8000 Series Devices, on page 947

Obtain a Signed Client Certificate for Secure Audit Log Streaming on a 7000/8000 Series Device

Optionally, you can use Transport Layer Security (TLS) certificates to secure communications between Firepower devices and a trusted audit log server. For *each device* (client certificates are unique), you must generate a certificate signing request (CSR), submit it to a Certificate Authority (CA) for signing, then import the signed certificate onto the device. You cannot use the FMC to import audit log certificates onto its managed devices. These certificates are unique to each device, and you must log into *each device* to import them.

To ensure security, use a globally recognized and trusted CA. The same CA must sign:

- Both the client certificate and the server certificate, if you plan to require mutual authentication between the device and the audit log server.

- Any intermediate certificates in the certificate chain. If the signing CA requires you to trust an intermediate CA, you must provide the necessary certificate chain (or certificate path).

The system generates certificate request keys in Base-64 encoded PEM format.

Procedure

Step 1 Log into the device's web interface and choose **System** > **Configuration**.

Step 2 Click **Audit Log Certificate**.

Step 3 Generate a CSR.

a) Click **Generate New CSR**.
b) Fill out the required location and organizational information.
c) Enter the fully qualified domain name of the server for which you want to request a certificate in the **Common Name** field. If the common name and the DNS hostname do not match, audit log streaming will fail.
d) Click **Generate**.

Step 4 Create a text file for the CSR.

a) Copy and paste the entire block of text in the certificate request, including the `BEGIN CERTIFICATE REQUEST` and `END CERTIFICATE REQUEST` lines.
b) Save the file as *`clientname`*`.csr`, where *`clientname`* is the name of the device where you plan to use the certificate.

Step 5 Submit the CSR to the CA and wait to receive the signed certificate.

Step 6 Import the signed certificate onto the device.

If you left the page, browse back to **System** > **Configuration** > **Audit Log Certificate**, then click **Import Audit Client Certificate**. Copy and paste the following:

- **Client Certificate**: All of the text in the signed certificate, including the `BEGIN CERTIFICATE` and `END CERTIFICATE` lines.
- **Private Key**: All of the text in the private key file, including the `BEGIN RSA PRIVATE KEY` and `END RSA PRIVATE KEY` lines.
- **Certificate Chain**: All of the text in each required intermediate certificate.

Make sure you import the correct certificate. Client certificates are unique.

Step 7 Click **Save**.

What to do next

If you have not already, use the device platform settings on the FMC to configure audit log streaming: Stream Audit Logs from Classic Devices, on page 934.

Prohibit Packet Transfer to FMC

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Any	7000 & 8000 Series	N/A	Admin

You may want to disable sending packet data from 7000 or 8000 Series devices to the Firepower Management Center in a low-bandwidth deployment if you are not concerned about the specific content of the packet that triggered an intrusion policy violation.

Procedure

Step 1 In the local web interface of your 7000 or 8000 Series device, choose **System** > **Configuration**.

Step 2 Click **Information**.

Step 3 Select **Prohibit Packet Transfer to the Management Center**.

Step 4 Click **Save**.

Configure Management Interfaces on a 7000/8000 Series Device

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Any	7000 & 8000 Series	Global only	Admin

Modify the management interface settings on the managed device using the web interface. You can optionally enable an event interface if your model supports it. For more information on management interfaces, see About Device Management Interfaces, on page 221.

Caution Be careful when making changes to the management interface; if you cannot re-connect because of a configuration error, you will need to access the device console port and reconfigure the settings at the CLI.

Procedure

Step 1 Choose **System** > **Configuration**, and then choose **Management Interfaces**.

Step 2 In the **Interfaces** area, click **Edit** next to the interface that you want to configure.

All available interfaces are listed in this section. You cannot add more interfaces.

You can configure the following options on each management interface:

- **Enabled**—Enable the management interface. Do **not** disable the default eth0 management interface. Some processes require the eth0 interface.
- **Channels**—(8000 series only) Configure an event-only interface. You can enable the eth1 management interface on your 8000 series device to act as an event interface. To do so, uncheck the **Management**

Traffic check box, and leave the **Event Traffic** check box checked. For the eth0 management interface, leave both check boxes checked.

The Firepower Management Center event-only interface cannot accept management channel traffic, so you should simply disable the management channel on the device event interface.

You can optionally disable **Event Traffic** for the management interface. In either case, the device will try to send events on the event-only interface, and if that interface is down, it will send events on the management interface even if you disable the event channel.

You cannot disable both event and management channels on an interface.

- **Mode**—Specify a link mode. Note that any changes you make to auto-negotiation are ignored for GigabitEthernet interfaces.
- **MTU**—Set the maximum transmission unit (MTU). The default is 1500. The range within which you can set the MTU can vary depending on the model and interface type.

  Because the system automatically trims 18 bytes from the configured MTU value, any value below 1298 does not comply with the minimum IPv6 MTU setting of 1280, and any value below 594 does not comply with the minimum IPv4 MTU setting of 576. For example, the system automatically trims a configured value of 576 to 558.
- **MDI/MDIX**—Set the **Auto-MDIX** setting.
- **IPv4 Configuration**—Set the IPv4 IP address. Choose:
  - **Static**—Manually enter the **IPv4 Management IP** address and **IPv4 Netmask**.
  - **DHCP**—Set the interface to use DHCP (eth0 only).
  - **Disabled**—Disable IPv4. Do **not** disable both IPv4 and IPv6.
- **IPv6 Configuration**—Set the IPv6 IP address. Choose:
  - **Static**—Manually enter the **IPv6 Management IP** address and **IPv6 Prefix Length**.
  - **DHCP**—Set the interface to use DHCPv6 (eth0 only).
  - **Router Assigned**—Enable stateless autoconfiguration.
  - **Disabled**—Disable IPv6. Do **not** disable both IPv4 and IPv6.

Step 3 In the **Routes** area, edit a static route by clicking **Edit** (), or add a route by clicking **Add** (). View the route statistics by clicking **View** ().

Note You need to add a static route for the event-only interface if the Firepower Management Center is on a remote network; otherwise, all traffic will match the default route through the management interface. For the default route, you can change only the gateway IP address. The default route always uses the eth0 interface. For information about routing, see Network Routes on Device Management Interfaces, on page 222.

You can configure the following settings for a static route:

- **Destination**—Set the destination address of the network to which you want to create a route.
- **Netmask** or **Prefix Length**—Set the netmask (IPv4) or prefix length (IPv6) for the network.

- **Interface**—Set the egress management interface.
- **Gateway**—Set the gateway IP address.

Step 4 In the **Shared Settings** area, set network parameters shared by all interfaces.

Note If you selected **DHCP** for the eth0 interface, you cannot manually specify some shared settings derived from the DHCP server.

You can configure the following shared settings:

- **Hostname**—Set the device hostname. If you change the hostname, reboot the device if you want the new hostname reflected in syslog messages. Syslog messages do not reflect a new hostname until after a reboot.
- **Domains**—Set the search domain(s) for the device, separated by commas. These domains are added to hostnames when you do not specify a fully-qualified domain name in a command, for example, **ping system**. The domains are used only on the management interface, or for commands that go through the management interface.
- **Primary DNS Server**, **Secondary DNS Server**, **Tertiary DNS Server**—Set the DNS servers to be used in order of preference.
- **Remote Management Port**—Set the remote management port for communication with the FMC. The FMC and managed devices communicate using a two-way, SSL-encrypted communication channel, which by default is on port 8305.

  Note Cisco **strongly** recommends that you keep the default settings for the remote management port, but if the management port conflicts with other communications on your network, you can choose a different port. If you change the management port, you must change it for **all** devices in your deployment that need to communicate with each other.

Step 5 In the **LCD Panel** area, check the **Allow reconfiguration of network settings** check box to enable changing network settings using the device's LCD panel.

You can use the LCD panel to edit the IP address for the device. Confirm that any changes you make are reflected on the managing Firepower Management Center. In some cases, you may need to update the data manually on the Firepower Management Center as well.

Caution Allowing reconfiguration using the LCD panel can present a security risk. You need only physical access, not authentication, to configure network settings using the LCD panel. The web interface warns you that enabling this option is a potential security issue.

Step 6 In the **Proxy** area, configure HTTP proxy settings.

The device is configured to directly-connect to the internet on ports TCP/443 (HTTPS) and TCP/80 (HTTP). You can use a proxy server, to which you can authenticate via HTTP Digest.

Note Proxies that use NT LAN Manager (NTLM) authentication are not supported.

a) Check the **Enabled** check box.
b) In the **HTTP Proxy** field, enter the IP address or fully-qualified domain name of your proxy server.
c) In the **Port** field, enter a port number.
d) Supply authentication credentials by choosing **Use Proxy Authentication**, and then provide a **User Name** and **Password**.

Step 7 Click **Save**.

Step 8 If you changed the management IP address, it might affect communication between the FMC and the managed device.

Changing the IP address will not affect the current connection. However, if the device or FMC reloads, then the connection needs to be reestablished. You need at least one of the devices (FMC or managed device) to have the correct IP address of the peer. For example, if you specified a NAT ID (instead of an IP address) for the FMC during device setup, then the device IP address that you defined on the FMC when you added the device will be wrong, and the FMC will not be able to reestablish communications. In this case, you must change the management IP address of the device in the FMC; see Update the Hostname or IP Address in FMC, on page 234.

Shut Down or Restart a 7000/8000 Series Device

Procedure

Step 1 On the device's web interface, choose **System** > **Configuration**.

Step 2 Choose **Process**.

Step 3 Do one of the following:

Shut down	Click **Run Command** next to **Shutdown Appliance**. **Caution** Do **not** shut off Firepower appliances using the power button; it may cause a loss of data. Using the web interface (or CLI) prepares the system to be safely powered off and restarted without losing configuration data.
Reboot	Click **Run Command** next to **Reboot Appliance**. **Note** Rebooting logs you out, and the system runs a database check that can take up to an hour to complete.
Restart the console	Click **Run Command** next to **Restart Appliance Console**.
Restart the Snort process	Click **Run Command** next to **Restart Snort**. **Caution** Restarting the Snort process temporarily interrupts traffic inspection. Whether traffic drops during this interruption or passes without inspection depends on how the device is configured. See Snort® Restart Traffic Behavior, on page 350 for more information.

View System Time for 7000/8000 Series Devices

Time settings are displayed on most pages in local time using the time zone you set on the Time Zone page in User Preferences, but are stored on the appliance using UTC time. In addition, the current time appears in

UTC at the top of the Time Synchronization page (local time is displayed in the Manual clock setting option, if enabled).



Restriction The Time Zone function (in User Preferences) assumes that the default system clock is set to UTC time. ***Do not change this.*** Changing the system time from UTC is ***not*** supported, and you will have to reimage the device.

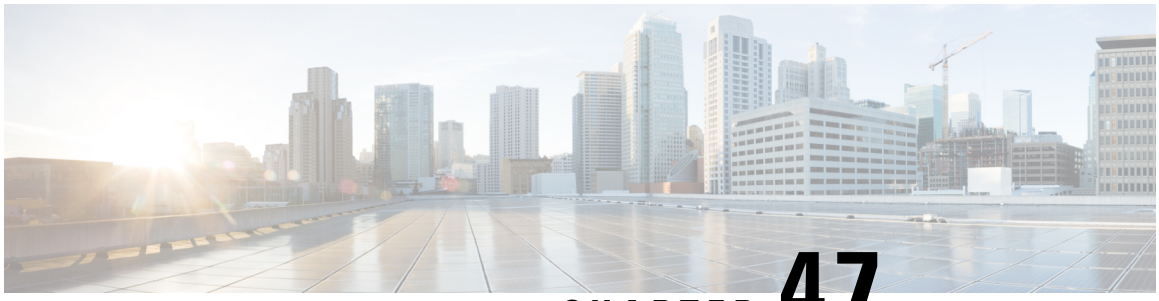
Use this procedure to verify system time information on 7000 and 8000 Series devices.

Procedure

Step 1 Log into the device's web interface and choose **System** > **Configuration**.

Step 2 Click **Time**.

If you are using NTP, see NTP Server Status, on page 916.



CHAPTER 47

Platform Settings for Firepower Threat Defense

Platform settings for Firepower Threat Defense devices configure a range of unrelated features whose values you might want to share among several devices. Even if you want different settings per device, you must create a shared policy and apply it to the desired device.

- Configure ARP Inspection, on page 949
- Configure Banners, on page 950
- Configure Fragment Handling, on page 951
- Configure HTTP, on page 952
- Configure ICMP Access Rules, on page 953
- Configure Secure Shell, on page 954
- Configure SMTP, on page 956
- Configure SNMP for Threat Defense, on page 956
- About Configuring Syslog, on page 962
- Configure Global Timeouts, on page 974
- Configure NTP Time Synchronization for Threat Defense, on page 976
- History for Firepower Threat Defense Platform Settings, on page 977

Configure ARP Inspection

By default, all ARP packets are allowed between bridge group members. You can control the flow of ARP packets by enabling ARP inspection.

ARP inspection prevents malicious users from impersonating other hosts or routers (known as ARP spoofing). ARP spoofing can enable a “man-in-the-middle” attack. For example, a host sends an ARP request to the gateway router; the gateway router responds with the gateway router MAC address. The attacker, however, sends another ARP response to the host with the attacker MAC address instead of the router MAC address. The attacker can now intercept all the host traffic before forwarding it on to the router.

ARP inspection ensures that an attacker cannot send an ARP response with the attacker MAC address, so long as the correct MAC address and the associated IP address are in the static ARP table.

When you enable ARP inspection, the FTD device compares the MAC address, IP address, and source interface in all ARP packets to static entries in the ARP table, and takes the following actions:

- If the IP address, MAC address, and source interface match an ARP entry, the packet is passed through.
- If there is a mismatch between the MAC address, the IP address, or the interface, then the FTD device drops the packet.

- If the ARP packet does not match any entries in the static ARP table, then you can set the FTD device to either forward the packet out all interfaces (flood), or to drop the packet.

Note The dedicated interface never floods packets even if this parameter is set to flood.

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.

Step 2 Select **ARP Inspection**.

Step 3 Add entries to the ARP inspection table.

a) Click **Add** to create a new entry, or click **Edit** if the entry already exists.

b) Select the desired options.

- **Inspect Enabled**—To perform ARP inspection on the selected interfaces and zones.
- **Flood Enabled**—Whether to flood ARP requests that do not match static ARP entries out all interfaces other than the originating interface or the dedicated management interface. This is the default behavior.

  If you do not elect to flood ARP requests, then only those requests that exactly match static ARP entries are allowed.
- **Security Zones**—Add the zones that contain the interfaces on which to perform the selected actions. The zones must be switched zones. For interfaces not in a zone, you can type the interface name into the field below the Selected Security Zone list and click **Add**. These rules will be applied to a device only if the device includes the selected interfaces or zones.

c) Click **OK**.

Step 4 Add static ARP entries according to Add a Static ARP Entry, on page 647.

Step 5 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure Banners

You can configure messages to show users when they connect to the device command line interface (CLI).

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.

Step 2 Select **Banner**.

Step 3 Configure the banner.

Following are some tips and requirements for banners.

- Only ASCII characters are allowed. You can use line returns (press Enter), but you cannot use tabs.
- You can dynamically add the hostname or domain name of the device by including the variables **$(hostname)** or **$(domain)**.
- Although there is no absolute length restriction on banners, Telnet or SSH sessions will close if there is not enough system memory available to process the banner messages.
- From a security perspective, it is important that your banner discourage unauthorized access. Do not use the words "welcome" or "please," as they appear to invite intruders in. The following banner sets the correct tone for unauthorized access:

```
You have logged in to a secure device.
If you are not authorized to access this device,
log out immediately or risk criminal charges.
```

Step 4 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure Fragment Handling

By default, the Firepower Threat Defense device allows up to 24 fragments per IP packet, and up to 200 fragments awaiting reassembly. You might need to let fragments on your network if you have an application that routinely fragments packets, such as NFS over UDP. However, if you do not have an application that fragments traffic, we recommend that you do not allow fragments by setting **Chain** to 1. Fragmented packets are often used as Denial of Service (DoS) attacks.



Note These settings establish the defaults for devices assigned this policy. You can override these settings for specific interfaces on a device by selecting **Override Default Fragment Setting** in the interface configuration. When you edit an interface, you can find the option on **Advanced** > **Security Configuration**. Select **Devices** > **Device Management**, edit a Firepower Threat Defense device, and select **Interfaces** to edit interface properties..

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.

Step 2 Select **Fragment Settings**.

Step 3 Configure the following options. Click **Reset to Defaults** if you want to use the default settings.

- **Size (Block)**—The maximum number of packet fragments from all connections collectively that can be waiting for reassembly. The default is 200 fragments.

- **Chain (Fragment)**—The maximum number of packets into which a full IP packet can be fragmented. The default is 24 packets. Set this option to 1 to disallow fragments.
- **Timeout (Sec)**—The maximum number of seconds to wait for an entire fragmented packet to arrive. The default is 5 seconds. If all fragments are not received within this time, all fragments are discarded.

Step 4 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure HTTP

If you want to allow HTTPS connections to one or more interfaces on the Firepower Threat Defense device, configure HTTPS settings. You can use HTTPS to download packet captures for troubleshooting.

Before you begin

- When you manage the Firepower Threat Defense using the FMC, HTTPS access to the Firepower Threat Defense is only for viewing packet capture files. The Firepower Threat Defense does not have a web interface for configuration in this management mode.
- HTTPS local users can only be configured at the CLI using the **configure user add** command. By default, there is an **admin** user for which you configured the password during initial setup. AAA external authentication is not supported.
- The physical management interface is shared between the Diagnostic logical interface and the Management logical interface; this configuration applies only to the Diagnostic logical interface, if used, or to other data interfaces. The Management logical interface is separate from the other interfaces on the device. It is used to set up and register the device to the FMC. It has a separate IP address and static routing.
- To use HTTPS, you do not need an access rule allowing the host IP address. You only need to configure HTTPS access according to this section.
- You can only use HTTPS to a reachable interface; if your HTTPS host is located on the outside interface, you can only initiate a management connection directly to the outside interface.
- The device allows a maximum of 5 concurrent HTTPS connections.
- You need network objects that define the hosts or networks you will allow to make HTTPS connections to the device. Select **Objects** > **Object Management** to configure objects.

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.
Step 2 Select **HTTP**.
Step 3 Enable the HTTPS server by clicking **Enable HTTP server**.
Step 4 (Optional) Change the HTTPS port. The default is 443.
Step 5 Identify the interfaces and IP addresses that allow HTTPS connections.

Use this table to limit which interfaces will accept HTTPS connections, and the IP addresses of the clients who are allowed to make those connections. You can use network addresses rather than individual IP addresses.

a) Click **Add** to add a new rule, or click **Edit** to edit an existing rule.
b) Configure the rule properties:

- **IP Address**—The network object that identifies the hosts or networks you are allowing to make HTTPS connections. Choose an object from the drop-down menu, or add a new network object by clicking +.
- **Security Zones**—Add the zones that contain the interfaces to which you will allow HTTPS connections. For interfaces not in a zone, you can type the interface name into the field below the Selected Security Zone list and click **Add**. These rules will be applied to a device only if the device includes the selected interfaces or zones.

c) Click **OK**.

Step 6 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure ICMP Access Rules

By default, you can send ICMP packets to any interface using either IPv4 or IPv6, with these exceptions:

- The FTD does not respond to ICMP echo requests directed to a broadcast address.
- The FTD only responds to ICMP traffic sent to the interface that traffic comes in on; you cannot send ICMP traffic through an interface to a far interface.

To protect the device from attacks, you can use ICMP rules to limit ICMP access to interfaces to particular hosts, networks, or ICMP types. ICMP rules function like access rules, where the rules are ordered, and the first rule that matches a packet defines the action.

If you configure any ICMP rule for an interface, an implicit deny ICMP rule is added to the end of the ICMP rule list, changing the default behavior. Thus, if you want to simply deny a few message types, you must include a permit any rule at the end of the ICMP rule list to allow the remaining message types.

We recommend that you always grant permission for the ICMP unreachable message type (type 3). Denying ICMP unreachable messages disables ICMP path MTU discovery, which can halt IPsec and PPTP traffic. Additionally ICMP packets in IPv6 are used in the IPv6 neighbor discovery process.

Before you begin

Ensure that the objects needed in the rules already exist. Select **Objects** > **Object Management** to configure objects. You need network objects that define the desired hosts or networks, and port objects that define the ICMP message types you want to control.

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.

Step 2 Select **ICMP**.

Step 3 Configure ICMP rules.

a) Click **Add** to add a new rule, or click **Edit** to edit an existing rule.

b) Configure the rule properties:

- **Action**—Whether to permit (allow) or deny (drop) matching traffic.
- **ICMP Service**—The port object that identifies the ICMP message type.
- **Network**—The network object that identifies the hosts or networks whose access you are controlling.
- **Security Zones**—Add the zones that contain the interfaces that you are protecting. For interfaces not in a zone, you can type the interface name into the field below the Selected Security Zone list and click **Add**. These rules will be applied to a device only if the device includes the selected interfaces or zones.

c) Click **OK**.

Step 4 (Optional.) Set rate limits on ICMPv4 Unreachable messages.

- **Rate Limit**—Sets the rate limit of unreachable messages, between 1 and 100 messages per second. The default is 1 message per second.
- **Burst Size**—Sets the burst rate, between 1 and 10. The system sends this number of replies, but subsequent replies are not sent until the rate limit is reached.

Step 5 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure Secure Shell

This section describes how to enable SSH connections to one or more *data* interfaces on the Firepower Threat Defense. SSH is not supported to the Diagnostic logical interface.



Note SSH is enabled by default on the Management interface; however, this screen does not affect Management SSH access.

The Management interface is separate from the other interfaces on the device. It is used to set up and register the device to the FMC. SSH for data interfaces shares the internal user list with SSH for the Management interface. Other settings are configured separately: for data interfaces, enable SSH and access lists using this screen; SSH traffic for data interfaces uses the regular routing configuration, and not any static routes configured at setup or at the CLI.

For the Management interface, to configure an SSH access list, see the **configure ssh-access-list** command in the Firepower Threat Defense Command Reference. To configure a static route, see the **configure network static-routes** command. By default, you configure the default route through the Management interface at initial setup.

To use SSH, you do not also need an access rule allowing the host IP address. You only need to configure SSH access according to this section.

You can only SSH to a reachable interface; if your SSH host is located on the outside interface, you can only initiate a management connection directly to the outside interface.

The device allows a maximum of 5 concurrent SSH connections.

Before you begin

- SSH local users can only be configured at the CLI using the **configure user add** command; see Creating Local User Accounts for the FTD CLI, on page 67. By default, there is an **admin** user for which you configured the password during initial setup. AAA external authentication is not supported.
- You need network objects that define the hosts or networks you will allow to make SSH connections to the device. Select **Objects** > **Object Management** to configure objects.

Note You cannot use the system-provided **any** network object. Instead, use **any-ipv4** or **any-ipv6**.

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.

Step 2 Select **Secure Shell**.

Step 3 Identify the interfaces and IP addresses that allow SSH connections.

Use this table to limit which interfaces will accept SSH connections, and the IP addresses of the clients who are allowed to make those connections. You can use network addresses rather than individual IP addresses.

a) Click **Add** to add a new rule, or click **Edit** to edit an existing rule.

b) Configure the rule properties:

- **IP Address**—The network object that identifies the hosts or networks you are allowing to make SSH connections. Choose an object from the drop-down menu, or add a new network object by clicking +.
- **Security Zones**—Add the zones that contain the interfaces to which you will allow SSH connections. For interfaces not in a zone, you can type the interface name into the field below the Selected Security Zone list and click **Add**. These rules will be applied to a device only if the device includes the selected interfaces or zones.

c) Click **OK**.

Step 4 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure SMTP

You must identity an SMTP server if you configure email alerts in the Syslog settings. The source email address you configure for Syslog must be a valid account on the SMTP servers.

Before you begin

Ensure that the network objects that define the host address of the primary and secondary SMTP servers exist. Select **Objects** > **Object Management** to define the objects. Alternatively, you can create the objects while editing the policy.

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.

Step 2 Click **SMTP Server**.

Step 3 Select the network objects that identify the **Primary Server IP Address** and optionally, the **Secondary Server IP Address**.

Step 4 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure SNMP for Threat Defense

Simple Network Management Protocol (SNMP) defines a standard way for network management stations running on PCs or workstations to monitor the health and status of many types of devices, including switches, routers, and security appliances. You can use the SNMP page to configure a firewall device for monitoring by SNMP management stations.

The Simple Network Management Protocol (SNMP) enables monitoring of network devices from a central location. Cisco security appliances support network monitoring using SNMP versions 1, 2c, and 3, as well as traps and SNMP read access; SNMP write access is not supported.

SNMPv3 supports read-only users and encryption with DES, 3DES, AES256, AES192, and AES128.

Note To create an alert to an external SNMP server, access **Policies** > **Action** > **Alerts**

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.

Step 2 Select **SNMP**.

Step 3 Enable SNMP and configure basic options.

- **Enable SNMP Servers**—Whether to provide SNMP information to the configured SNMP hosts. You can deselect this option to disable SNMP monitoring while retaining the configuration information.
- **Read Community String**, **Confirm**—Enter the password used by a SNMP management station when sending requests to the Firepower Threat Defense device. The SNMP community string is a shared secret among the SNMP management stations and the network nodes being managed. The security device uses the password to determine if the incoming SNMP request is valid. The password is a case-sensitive alphanumeric string of up to 32 characters; spaces and special characters are not permitted.
- **System Administrator Name**—Enter the name of the device administrator or other contact person. This string is case-sensitive and can be up to 127 characters. Spaces are accepted, but multiple spaces are shortened to a single space.
- **Location**—Enter the location of this security device (for example, Building 42,Sector 54). This string is case-sensitive and can be up to 127 characters. Spaces are accepted, but multiple spaces are shortened to a single space.
- **Port**—Enter the UDP port on which incoming requests will be accepted. The default is 161.

Step 4 (SNMPv3 only.) Add SNMPv3 Users, on page 957.

Step 5 Add SNMP Hosts, on page 959.

Step 6 Configure SNMP Traps, on page 960.

Step 7 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Add SNMPv3 Users

Note You create users for SNMPv3 only. These steps are not applicable for SNMPv1 or SNMPv2c.

Note that SNMPv3 only supports read-only users.

SNMP users have a specified username, an authentication password, an encryption password, and authentication and encryption algorithms to use. The authentication algorithm options are MD5 and SHA. The encryption algorithm options are DES, 3DES, AES256, AES192, and AES128.

Note When using SNMPv3 with clustering or High Availability, if you add a new cluster unit after the initial cluster formation or you replace a High Availability unit, then SNMPv3 users are not replicated to the new unit. You must remove the users, re-add them, and then redeploy your configuration to force the users to replicate to the new unit.

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.

Step 2 Click **SNMP** > **Users**.

Step 3 Click **Add**.

Step 4 Select the security level for the user from the **Security Level** drop-down list.

- **Auth**—Authentication but No Privacy, which means that messages are authenticated.
- **No Auth**—No Authentication and No Privacy, which means that no security is applied to messages.
- **Priv**—Authentication and Privacy, which means that messages are authenticated and encrypted.

Step 5 Enter the name of the SNMP user in the **Username** field. Usernames must be 32 characters or less.

Step 6 Select the type of password, you want to use in the **Encryption Password Type** drop-down list.

- **Clear text**—The Firepower Threat Defense device will still encrypt the password when deploying to the device.
- **Encrypted**—The Firepower Threat Defense device will directly deploy the encrypted password.

Step 7 In the **Auth Algorithm Type** drop-down list, select the type of authentication you want to use: MD5, SHA.

Step 8 In the **Authentication Password** field, enter the password to use for authentication. If you selected Encrypted as the Encrypt Password Type, the password must be formatted as xx:xx:xx..., where xx are hexadecimal values.

Note The length of the password will depend on the authentication algorithm selected. For all passwords, the length must be 256 characters or less.

If you selected Clear Text as the Encrypt Password Type, repeat the password in the **Confirm** field.

Step 9 In the **Encryption Type** drop-down list, select the type of encryption you want to use: AES128, AES192, AES256, 3DES, or DES.

Note To use AES or 3DES encryption, you must have the appropriate license installed on the device.

Step 10 Enter the password to use for encryption in the **Encryption Password** field. If you selected Encrypted as the Encrypt Password Type, the password must be formatted as xx:xx:xx..., where xx are hexadecimal values. For encrypted passwords, the length of the password depends on the encryption type selected. The password sizes are as follows (where each xx is one octal):

- AES 128 requires 16 octals
- AES 192 requires 24 octals
- AES 256 requires 32 octals
- 3DES requires 32 octals
- DES can be any size

Note For all passwords, the length must be 256 characters or less.

If you selected Clear Text as the Encrypt Password Type, repeat the password in the **Confirm** field.

Step 11 Click **OK**.

Step 12 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Add SNMP Hosts

Use Host to add or edit entries in the SNMP Hosts table on the SNMP page. These entries represent SNMP management stations allowed to access the Firepower Threat Defense device.

Before you begin

Ensure that the network objects that define the SNMP management stations exist. Select **Device** > **Object Management** to configure network objects.

Note Only IPv4 addresses are supported.

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.

Step 2 Click **SNMP** > **Hosts**.

Step 3 Click **Add**.

Step 4 In the **IP Address** field, either enter a valid IPv6 or IPv4 host or select the network object that defines the SNMP management station's host address.

Step 5 Select the appropriate SNMP version from the **SNMP version** drop-down list.

Step 6 (SNMPv3 only.) Select the username of the SNMP user that you configured from the **User Name** drop-down list.

Note You can associate up to 23 SNMP users per SNMP host.

Step 7 (SNMPv1, 2c only.) In the **Read Community String** field, enter the community string that you have already configured, for read access to the device. Re-enter the string to confirm it.

Note This string is required, only if the string used with this SNMP station is different from the one already defined in the **Enable SNMP Server** section.

Step 8 Select the type of communication between the device and the SNMP management station. You can select both types.

- **Poll**—The management station periodically requests information from the device.
- **Trap**—The device sends trap events to the management station as they occur.

Step 9 In the **Port** field, enter a UDP port number for the SNMP host. The default value is 162. The valid range is 1 to 65535.

Step 10 Click **Add** to enter or select the interface on which this SNMP management station contacts the device.

Step 11 In the **Zones/Interfaces** list, add the zones that contain the interfaces through which the device communicates with the management station. For interfaces not in a zone, you can type the interface name into the field below the **Selected Zone/Interface** list and click **Add**. The host will be configured on a device only if the device includes the selected interfaces or zones.

Step 12 Click **OK**.

Step 13 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure SNMP Traps

Use SNMP Traps to configure SNMP traps (event notifications) for the Firepower Threat Defense device. Traps are different from browsing; they are unsolicited "comments" from the Firepower Threat Defense device to the management station for certain events, such as linkup, linkdown, and syslog event generated. An SNMP object ID (OID) for the device appears in SNMP event traps sent from the device.

Some traps are not applicable to certain hardware models. These traps will be ignored if you apply the policy to one of these models. For example, not all models have field-replaceable units, so the **Field Replaceable Unit Insert/Delete** trap will not be configured on those models.

SNMP traps are defined in either standard or enterprise-specific MIBs. Standard traps are created by the IETF and documented in various RFCs. SNMP traps are compiled into the Firepower Threat Defense software.

If needed, you can download RFCs, standard MIBs, and standard traps from the following location:

http://www.ietf.org/

Browse the complete list of Cisco MIBs, traps, and OIDs from the following location:

SNMP Object Navigator

In addition, download Cisco OIDs by FTP from the following location:

ftp://ftp.cisco.com/pub/mibs/oid/oid.tar.gz

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.

Step 2 Click **SNMP** > **SNMP Traps** to configure SNMP traps (event notifications) for the Firepower Threat Defense device.

Step 3 Select the appropriate Enable Traps options. You can select either or both options.

a) Check **Enable All SNMP Traps** to quickly select all traps in the subsequent four sections.

b) Check **Enable All Syslog Traps** to enable transmission of trap-related syslog messages.

Note SNMP traps are of higher priority than other notification messages from the Firepower Threat Defense as they are expected to be near real-time. When you enable all SNMP or syslog traps, it is possible for the SNMP process to consume excess resources in the agent and in the network, causing the system to hang. If you notice system delays, unfinished requests, or timeouts, you can selectively enable SNMP and syslog traps. You can also limit the rate at which syslog messages are generated by severity level or message ID. For example, all syslog message IDs that begin with the digits 212 are associated with the SNMP class; see Limit the Rate of Syslog Message Generation, on page 971.

Step 4 The event-notification traps in the **Standard** section are enabled by default for an existing policy:

- **Authentication** – Unauthorized SNMP access. This authentication failure occurs for packets with an incorrect community string.
- **Link Up** – One of the device's communication links has become available (it has "come up"), as indicated in the notification.
- **Link Down** – One of the device's communication links has failed, as indicated in the notification.
- **Cold Start** – The device is reinitializing itself such that its configuration or the protocol entity implementation may be altered.
- **Warm Start** – The device is reinitializing itself such that its configuration and the protocol entity implementation is unaltered.

Step 5 Select the desired event-notification traps in the **Entity MIB** section:

- **Field Replaceable Unit Insert** – A Field Replaceable Unit (FRU) has been inserted, as indicated. (FRUs include assemblies such as power supplies, fans, processor modules, interface modules, etc.)
- **Field Replaceable Unit Delete** – A Field Replaceable Unit (FRU) has been removed, as indicated in the notification
- **Configuration Change** – There has been a hardware change, as indicated in the notification

Step 6 Select the desired event-notification traps in the **Resource** section:

- **Connection Limit Reached** – This trap indicates that a connection attempt was rejected because the configured connections limit has been reached.

Step 7 Select the desired event-notification traps in the **Other** section:

- **NAT Packet Discard** – This notification is generated when IP packets are discarded by the NAT function. Available Network Address Translation addresses or ports have fallen below configured threshold.

Step 8 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

About Configuring Syslog

You can enable system logging (syslog) for Firepower Threat Defense devices. Logging information can help you identify and isolate network or device configuration problems. The following topics explain logging and how to configure it.

About Syslog

System logging is a method of collecting messages from devices to a server running a syslog daemon. Logging to a central syslog server helps in aggregation of logs and alerts. Cisco devices can send their log messages to a UNIX-style syslog service. A syslog service accepts messages and stores them in files, or prints them according to a simple configuration file. This form of logging provides protected long-term storage for logs. Logs are useful both in routine troubleshooting and in incident handling.

Table 76: System Logs for Firepower Threat Defense

Logs Related To	Details	Configure In
Device and system health, network configuration	This syslog configuration generates messages for features running on the data plane, that is, features that are defined in the CLI configuration that you can view with the **show running-config** command. This includes features such as routing, VPN, data interfaces, DHCP server, NAT, and so forth. Data plane syslog messages are numbered, and they are the same as those generated by devices running ASA software. However, Firepower Threat Defense does not necessarily generate every message type that is available for ASA Software. For information on these messages, see *Cisco Firepower Threat Defense Syslog Messages* at https://www.cisco.com/c/en/us/td/docs/security/firepower/Syslogs/b_fptd_syslog_guide.html. This configuration is explained in the following topics.	**Platform Settings**
Policies, rules, and events	This syslog configuration generates alerts for access control rules, intrusion rules, and other advanced services as described in Configurations Supporting Alert Responses, on page 2076. These messages are not numbered. For information on configuring this type of syslog, see Creating a Syslog Alert Response, on page 2078.	**Alert Responses**

You can configure more than one syslog server, and control the messages and events sent to each server. You can also configure different destinations, such as console, email, internal buffer, and so forth.

Severity Levels

The following table lists the syslog message severity levels.

Table 77: Syslog Message Severity Levels

Level Number	Severity Level	Description
0	**emergencies**	System is unusable.

Level Number	Severity Level	Description
1	**alert**	Immediate action is needed.
2	**critical**	Critical conditions.
3	**error**	Error conditions.
4	**warning**	Warning conditions.
5	**notification**	Normal but significant conditions.
6	**informational**	Informational messages only.
7	**debugging**	Debugging messages only. Log at this level only temporarily, when debugging issues. This log level can potentially generate so many messages that system performance can be affected.



Note ASA and FTD do not generate syslog messages with a severity level of zero (emergencies).

Syslog Message Filtering

You can filter generated syslog messages so that only certain syslog messages are sent to a particular output destination. For example, you could configure the FTD device to send all syslog messages to one output destination and to send a subset of those syslog messages to a different output destination.

Specifically, you can direct syslog messages to an output destination according to the following criteria:

- Syslog message ID number
- Syslog message severity level
- Syslog message class (equivalent to a functional area)

You customize these criteria by creating a message list that you can specify when you set the output destination. Alternatively, you can configure the FTD device to send a particular message class to each type of output destination independently of the message list.

Syslog Message Classes

You can use syslog message classes in two ways:

- Specify an output location for an entire category of syslog messages.
- Create a message list that specifies the message class.

The syslog message class provides a method of categorizing syslog messages by type, equivalent to a feature or function of the device. For example, the rip class denotes RIP routing.

All syslog messages in a particular class share the same initial three digits in their syslog message ID numbers. For example, all syslog message IDs that begin with the digits 611 are associated with the vpnc (VPN client) class. Syslog messages associated with the VPN client feature range from 611101 to 611323.

In addition, most of the ISAKMP syslog messages have a common set of prepended objects to help identify the tunnel. These objects precede the descriptive text of a syslog message when available. If the object is not known at the time that the syslog message is generated, the specific heading = value combination does not appear.

The objects are prefixed as follows:

Group = *groupname*, Username = *user*, IP = *IP_address*

Where the group is the tunnel-group, the username is the username from the local database or AAA server, and the IP address is the public IP address of the remote access client or Layer 2 peer.

The following table lists the message classes and the range of message IDs in each class.

Table 78: Syslog Message Classes and Associated Message ID Numbers

Class	Definition	Syslog Message ID Numbers
auth	User Authentication	109, 113
—	Access Lists	106
—	Application Firewall	415
—	Botnet Traffic Filtering	338
bridge	Transparent Firewall	110, 220
ca	PKI Certification Authority	717
citrix	Citrix Client	723
—	Clustering	747
—	Card Management	323
config	Command Interface	111, 112, 208, 308
csd	Secure Desktop	724
cts	Cisco TrustSec	776
dap	Dynamic Access Policies	734
eap, eapoudp	EAP or EAPoUDP for Network Admission Control	333, 334
eigrp	EIGRP Routing	336
email	E-mail Proxy	719
—	Environment Monitoring	735
ha	Failover	101, 102, 103, 104, 105, 210, 311, 709

Class	Definition	Syslog Message ID Numbers
—	Identity-based Firewall	746
ids	Intrusion Detection System	400, 733
—	IKEv2 Toolkit	750, 751, 752
ip	IP Stack	209, 215, 313, 317, 408
ipaa	IP Address Assignment	735
ips	Intrusion Protection System	400, 401, 420
—	IPv6	325
—	Licensing	444
mdm-proxy	MDM Proxy	802
nac	Network Admission Control	731, 732
nacpolicy	NAC Policy	731
nacsettings	NAC Settings to apply NAC Policy	732
—	NAT and PAT	305
—	Network Access Point	713
np	Network Processor	319
—	NP SSL	725
ospf	OSPF Routing	318, 409, 503, 613
—	Password Encryption	742
—	Phone Proxy	337
rip	RIP Routing	107, 312
rm	Resource Manager	321
—	Smart Call Home	120
session	User Session	106, 108, 201, 202, 204, 302, 303, 304, 305, 314, 405, 406, 407, 500, 502, 607, 608, 609, 616, 620, 703, 710
snmp	SNMP	212
—	ScanSafe	775
ssl	SSL Stack	725
svc	SSL VPN Client	722

Class	Definition	Syslog Message ID Numbers
sys	System	199, 211, 214, 216, 306, 307, 315, 414, 604, 605, 606, 610, 612, 614, 615,701, 711, 741
—	Threat Detection	733
tag-switching	Service Tag Switching	779
vm	VLAN Mapping	730
vpdn	PPTP and L2TP Sessions	213, 403, 603
vpn	IKE and IPsec	316, 320, 402, 404, 501, 602, 702, 713, 714, 715
vpnc	VPN Client	611
vpnfo	VPN Failover	720
vpnlb	VPN Load Balancing	718
—	VXLAN	778
webfo	WebVPN Failover	721
webvpn	WebVPN and AnyConnect Client	716

Guidelines for Logging

This section includes guidelines and limitations that you should review before configuring logging.

IPv6 Guidelines

- IPv6 is supported. Syslogs can be sent using TCP or UDP.
- Ensure that the interface configured for sending syslogs is enabled, IPv6 capable, and the syslog server is reachable through the designated interface.
- Secure logging over IPv6 is not supported.

Additional Guidelines

- The syslog server must run a server program called syslogd. Windows provides a syslog server as part of its operating system.
- To view logs generated by the FTD device, you must specify a logging output destination. If you enable logging without specifying a logging output destination, the FTD device generates messages but does not save them to a location from which you can view them. You must specify each different logging output destination separately.
- It is not possible to have two different lists or classes being assigned to different syslog servers or same locations.

- You can configure up to 16 syslog servers.
- The syslog server should be reachable through the FTD device. You should configure the device to deny ICMP unreachable messages on the interface through which the syslog server is reachable and to send syslogs to the same server. Make sure that you have enabled logging for all severity levels. To prevent the syslog server from crashing, suppress the generation of syslogs 313001, 313004, and 313005.
- The number of UDP connections for syslog is directly related to the number of CPUs on the hardware platform and the number of syslog servers you configure. At any point in time, there can be as many UDP syslog connections as there are CPUs times the number of configured syslog servers. This is the expected behavior. Note that the global UDP connection idle timeout applies to these sessions, and the default is 2 minutes. You can adjust that setting if you want to close these session more quickly, but the timeout applies to all UDP connections, not just syslog.
- When the FTD device sends syslogs via TCP, the connection takes about one minute to initiate after the syslogd service restarts.

Configure Syslog Logging for FTD Devices

To configure syslog settings, perform the following steps:

Before you begin

See requirements in Guidelines for Logging, on page 966.

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.

Step 2 Click **Syslog** from the table of contents.

Step 3 Click **Logging Setup** to enable logging, specify FTP Server settings, and specify Flash usage. For more information, see Enable Logging and Configure Basic Settings, on page 968

Step 4 Click **Logging Destinations** to enable logging to specific destinations and to specify filtering on message severity level, event class, or on a custom event list. For more information, see Enable Logging Destinations, on page 969

You must enable a logging destination to see messages at that destination.

Step 5 Click **E-mail Setup** to specify the e-mail address that is used as the source address for syslog messages that are sent as e-mail messages. For more information, see Send Syslog Messages to an E-mail Address, on page 970

Step 6 Click **Events List** to define a custom event list that includes an event class, a severity level, and an event ID. For more information, see Create a Custom Event List, on page 970

Step 7 Click **Rate Limit** to specify the volume of messages being sent to all configured destinations and define the message severity level to which you want to assign rate limits. For more information, see Limit the Rate of Syslog Message Generation, on page 971

Step 8 Click **Syslog Settings** to specify the logging facility, enable the inclusion of a time stamp, and enable other settings to set up a server as a syslog destination. For more information, see Configure Syslog Settings, on page 972

Step 9 Click **Syslog Servers** to specify the IP address, protocol used, format, and security zone for the syslog server that is designated as a logging destination. For more information, see Configure a Syslog Server, on page 973

Enable Logging and Configure Basic Settings

You must enable logging for the system to generate syslog messages for data plane events.

You can also set up archiving on flash or an FTP server as a storage location when the local buffer becomes full. You can manipulate logging data after it is saved. For example, you could specify actions to be executed when certain types of syslog messages are logged, extract data from the log and save the records to another file for reporting, or track statistics using a site-specific script.

The following procedure explains some of the basic syslog settings.

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.

Step 2 Select **Syslog** > **Logging Setup**.

Step 3 Enable logging and configure basic logging settings.

- **Enable Logging**—Turns on data plane system logging for the Firepower Threat Defense device.
- **Enable Logging on the Failover Standby Unit**—Turns on logging for the standby for the Firepower Threat Defense device, if available.
- **Send syslogs in EMBLEM format**—Enables EMBLEM format logging for every logging destination. If you enable EMBLEM, you must use the UDP protocol to publish syslog messages; EMBLEM is not compatible with TCP.

  Note Syslog messages in RFC5424 format, typically displays the priority value (PRI). However, in FMC, if you want to display the PRI value in the syslog messages of the managed FTD, ensure to enable the EMBLEM format. For more information on PRI, see RFC5424.

- **Send debug messages as syslogs**—Redirects all the debug trace output to the syslog. The syslog message does not appear in the console if this option is enabled. Therefore, to see debug messages, you must enable logging at the console and configure it as the destination for the debug syslog message number and logging level. The syslog message number used is 71101. Default logging level for this syslog is debug.
- **Memory Size of Internal Buffer**—Specify the size of the internal buffer to which syslog messages are saved if the logging buffer is enabled. When the buffer fills up, it is overwritten. The default is 4096 bytes. The range is 4096 to 52428800.

Step 4 (Optional) Configure an FTP server if you want to save log buffer contents to the server before the buffer is overwritten. Specify the FTP Server information.

- **FTP Server Buffer Wrap**— To save the buffer contents to the FTP server before it is overwritten, check this box and enter the necessary destination information in the following fields. To remove the FTP configuration, deselect this option.
- **IP Address**—Select the host network object that contains the IP address of the FTP server.
- **User Name**—Enter the user name to use when connecting to the FTP server.
- **Path**—Enter the path, relative to the FTP root, where the buffer contents should be saved.

- **Password/ Confirm**—Enter and confirm the password used to authenticate the user name to the FTP server.

Step 5 (Optional) Specify Flash size if you want to save log buffer contents to flash before the buffer is overwritten.

- **Flash**—To save the buffer contents to the flash memory before it is overwritten, check this box.
- **Maximum flash to be used by logging (KB)**—Specify the maximum space to be used in the flash memory for logging(in KB). The range is 4-8044176 kilobytes.
- **Minimum free space to be preserved (KB)**—Specifies the minimum free space to be preserved in flash memory (in KB). The range is 0-8044176 kilobytes.

Step 6 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Enable Logging Destinations

You must enable a logging destination to see messages at that destination. When enabling a destination, you must also specify the message filter for the destination.

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.

Step 2 Select **Syslog** > **Logging Destinations**.

Step 3 Click **Add** to enable a destination and apply a logging filter, or edit an existing destination.

Step 4 In the **Logging Destinations** dialog box, select a destination and configure the filter to use for a destination:

a) Choose the destination you are enabling in the **Logging Destination** drop-down list. You can create one filter per destination: Console, E-Mail, Internal buffer, SNMP trap, SSH Sessions, and Syslog servers.

Note Console and SSH session logging works in the diagnostic CLI only. Enter **system support diagnostic-cli**.

b) In **Event Class**, choose the filter that will apply to all classes not listed in the table.

You can configure these filters:

- **Filter on severity** —Select the severity level. Messages at this level or higher are sent to the destination
- **Use Event List** —Select the event list that defines the filter. You create these lists on the **Event Lists** page.
- **Disable Logging** —Prevents messages from being sent to this destination.

c) If you want to create filters per event class, click **Add** to create a new filter, or edit an existing filter, and select the event class and severity level to limit messages in that class. Click **OK** to save the filter.

For an explanation of the event classes, see Syslog Message Classes, on page 963.

d) Click **OK** .

Step 5 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Send Syslog Messages to an E-mail Address

You can set up a list of recipients for syslog messages to be sent as e-mails.

Before you begin

- Configure an SMTP server on the SMTP Server platform settings page
- Enable Logging and Configure Basic Settings, on page 968
- Enable Logging Destinations

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.

Step 2 Select **Syslog** > **Email Setup**.

Step 3 Specify the e-mail address that is used as the source address for syslog messages that are sent as e-mail messages.

Step 4 Click **Add** to enter a new e-mail address recipient of the specified syslog messages.

Step 5 Choose the severity level of the syslog messages that are sent to the recipient from the drop-down list.

The syslog message severity filter used for the destination e-mail address causes messages of the specified severity level and higher to be sent. For information on the levels, see Severity Levels, on page 962.

Step 6 Click **OK**.

Step 7 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Create a Custom Event List

An event list is a custom filter you can apply to a logging destination to control which messages are sent to the destination. Normally, you filter messages for a destination based on severity only, but you can use an event list to fine-tune which messages are sent based on a combination of event class, severity, and message identifier (ID).

Creating a custom event list is a two-step process. You create a custom list in the **Event Lists**, and then use the event list to define the logging filter for the various types of destination, in the **Logging Destinations**.

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.

Step 2 Select **Syslog** > **Events List**.

Step 3 Configure an event list.

a) Click **Add** to add a new list, or edit an existing list.

b) Enter a name for the event list in the **Name** field. Spaces are not allowed.

c) To identify messages based on severity or event class, select the **Severity/Event Class** tab and add or edit entries.

For information on the available classes see Syslog Message Classes, on page 963.

For information on the levels, see Severity Levels, on page 962.

Certain event classes are not applicable for the device in transparent mode. If such options are configured then they will be bypassed and not deployed.

d) To identify messages specifically by message ID, select the **Message ID** and add or edit the IDs.

You can enter a range of IDs using a hyphen, for example, 100000-200000. IDs are six digits. For information on how the initial three digits map to features, see Syslog Message Classes, on page 963.

For specific message numbers, see Cisco ASA Series Syslog Messages.

e) Click **OK** to save the event list.

Step 4 Click **Logging Destinations** and add or edit the destination that should use the filter.

See Enable Logging Destinations, on page 969.

Step 5 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Limit the Rate of Syslog Message Generation

You can limit the rate at which syslog messages are generated by severity level or message ID. You can specify individual limits for each logging level and each Syslog message ID. If the settings conflict, the Syslog message ID limits take precedence.

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.

Step 2 Select **Syslog** > **Rate Limit**.

Step 3 To limit message generation by severity level, click **Logging Level** > **Add** and configure the following options:

- **Logging Level**—The severity level you are rate limiting. For information on the levels, see Severity Levels, on page 962.
- **Number of messages**—The maximum number of messages of the specified type allowed in the specified time period.

- **Interval**—The number of seconds before the rate limit counter resets.

Step 4 Click **OK**.

Step 5 To limit message generation by syslog message ID, click **Syslog Level** > **Add** and configure the following options:

- **Syslog ID**—The syslog message ID you are rate limiting. For specific message numbers, see Cisco ASA Series Syslog Messages.
- **Number of messages**—The maximum number of messages of the specified type allowed in the specified time period.
- **Interval**—The number of seconds before the rate limit counter resets.

Step 6 Click **OK**.

Step 7 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure Syslog Settings

You can configure general syslog settings to set the facility code to be included in syslog messages that are sent to syslog servers, specify whether a timestamp is included in each message, specify the device ID to include in messages, view and modify the severity levels for messages, and disable the generation of specific messages.

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.

Step 2 Select **Syslog** > **Syslog Settings**.

Step 3 Select a system log facility for syslog servers to use as a basis to file messages in the **Facility** drop-down list.

The default is LOCAL4(20), which is what most UNIX systems expect. However, because your network devices share available facilities, you might need to change this value for system logs.

Step 4 Select the **Enable timestamp on each syslog message** check box to include the date and time a message was generated in the syslog message.

Timestamp timezone is always UTC.

Step 5 If you want to add a device identifier to syslog messages (which is placed at the beginning of the message), check the **Enable Syslog Device ID** check box and then select the type of ID.

- **Interface**—To use the IP address of the selected interface, regardless of the interface through which the appliance sends the message. Select the security zone that identifies the interface. The zone must map to a single interface.
- **User Defined ID**—To use a text string (up to 16 characters) of your choice.
- **Host Name**—To use the hostname of the device.

Step 6 Use the Syslog Message table to alter the default settings for specific syslog messages. You need to configure rules in this table only if you want to change the default settings. You can change the severity assigned to a message, or you can disable the generation of a message.

By default, Netflow is enabled and the entries are shown in the table.

a) To suppress syslog messages that are redundant because of Netflow, select **Netflow Equivalent Syslogs**.

This adds the messages to the table as suppressed messages.

Note If any of these syslog equivalents are already in the table, your existing rules are not overwritten.

b) To add a rule, click **Add**.

c) You select the message number whose configuration you want to change, from the **Syslog ID** drop down list and then select the new severity level from the **Logging Level** drop down list, or select **Suppressed** to disable the generation of the message. Typically, you would not change the severity level and disable the message, but you can make changes to both fields if desired.

d) Click **OK** to add the rule to the table.

Step 7 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configure a Syslog Server

To configure a syslog server to handle messages generated from your system, perform the following steps.

Before you begin

- See requirements in Guidelines for Logging, on page 966.
- Make sure your devices can reach your syslog collector on the network.

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.

Step 2 Select **Syslog** > **Syslog Server**.

Step 3 Check the **Allow user traffic to pass when TCP syslog server is down** check box, to allow traffic if any syslog server that is using the TCP protocol is down.

Step 4 Enter a size of the queue for storing syslog messages on the security appliance when syslog server is busy in the **Message queue size (messages)** field. The minimum is 1 message. The default is 512. Specify 0 to allow an unlimited number of messages to be queued (subject to available block memory).

Step 5 Click **Add** to add a new syslog server.

a) In the **IP Address** drop-down list, select a network host object that contains the IP address of the syslog server.

b) Choose the protocol (either TCP or UDP) and enter the port number for communications between the Firepower Threat Defense device and the syslog server.

UDP is faster and uses less resources on the device than TCP.

The default ports are 514 for UDP, 1470 for TCP. Valid non-default port values for either protocol are 1025 through 65535.

c) Check the **Log messages in Cisco EMBLEM format (UDP only)** check box to specify whether to log messages in Cisco EMBLEM format (available only if UDP is selected as the protocol).

Note Syslog messages in RFC5424 format, typically displays the priority value (PRI). However, in FMC, only when you enable logging in Cisco EMBLEM format, the PRI value in the syslog messages of the managed FTD is displayed. For more information on PRI, see RFC5424.

d) Check the **Enable Secure Syslog** check box to encrypt the connection between the device and server using SSL/TLS over TCP.

e) Add the zones that contain the interfaces used to communicate with the syslog server. For interfaces not in a zone, you can type the interface name into the field below the **Selected Zones/Interface** list and click **Add**. These rules will be applied to a device only if the device includes the selected interfaces or zones.

Note If the syslog server is on the network attached to the physical Management interface, you must type the name of that interface into the **Interface Name** field below the **Selected Security Zones** list and click **Add**. You must also configure this name (if not already configured), and an IP address, for the Diagnostic interface (edit the device from the Device Management page and select the Interfaces tab). For more information about the management/diagnostic interface, see Diagnostic Interface, on page 609.

f) Click **OK**.

Step 6 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configure Global Timeouts

You can set the global idle timeout durations for the connection and translation slots of various protocols. If the slot has not been used for the idle time specified, the resource is returned to the free pool.

You can also set a time out for console sessions with the device.

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.

Step 2 Select **Timeouts**.

Step 3 Configure the timeouts you want to change.

For any given setting, select **Custom** to define your own value, **Default** to return to the system default value. In most cases, the maximum timeout is 1193 hours.

You can disable some timeouts by selecting **Disable**.

- **Console Timeout**—The idle time until a connection to the console is closed, range is 0 to 60 minutes. The default is 0, which means the session does not time out. If you change the value, existing console sessions use the old timeout value. The new value applies to new connections only.
- **Translation Slot (xlate)**—The idle time until a NAT translation slot is freed. This duration must be at least 1 minute. The default is 3 hours.
- **Connection (Conn)**—The idle time until a connection slot is freed. This duration must be at least 5 minutes. The default is 1 hour.
- **Half-Closed**—The idle time until a TCP half-closed connection closes. A connection is considered half-closed if both the FIN and FIN-ACK have been seen. If only the FIN has been seen, the regular connection timeout applies. The minimum is 30 seconds. The default is 10 minutes.
- **UDP**—The idle time until a UDP connection closes. This duration must be at least 1 minute. The default is 2 minutes.
- **ICMP**—The idle time after which general ICMP states are closed. The default (and minimum) is 2 seconds.
- **RPC/Sun RPC**—The idle time until a SunRPC slot is freed. This duration must be at least 1 minute. The default is 10 minutes.
- **H.225**—The idle time until an H.225 signaling connection closes. The default is 1 hour. To close a connection immediately after all calls are cleared, a timeout of 1 second (0:0:1) is recommended.
- **H.323**—The idle time after which H.245 (TCP) and H.323 (UDP) media connections close. The default (and minimum) is 5 minutes. Because the same connection flag is set on both H.245 and H.323 media connections, the H.245 (TCP) connection shares the idle timeout with the H.323 (RTP and RTCP) media connection.
- **SIP**—The idle time until a SIP signaling port connection closes. This duration must be at least 5 minutes. The default is 30 minutes.
- **SIP Media**—The idle time until a SIP media port connection closes. This duration must be at least 1 minute. The default is 2 minutes. The SIP media timer is used for SIP RTP/RTCP with SIP UDP media packets, instead of the UDP inactivity timeout.
- **SIP Disconnect**—The idle time after which SIP session is deleted if the 200 OK is not received for a CANCEL or a BYE message, between 0:0:1 and 0:10:0. The default is 2 minutes (0:2:0).
- **SIP Invite**—The idle time after which pinholes for PROVISIONAL responses and media xlates will be closed, between 0:1:0 and 00:30:0. The default is 3 minutes (0:3:0).

- **SIP Provisional Media**—The timeout value for SIP provisional media connections, between 1 and 30 minutes. The default is 2 minutes.
- **Floating Connection**—When multiple routes exist to a network with different metrics, the system uses the one with the best metric at the time of connection creation. If a better route becomes available, then this timeout lets connections be closed so a connection can be reestablished to use the better route. The default is 0 (the connection never times out). To make it possible to use better routes, set the timeout to a value between 0:0:30 and 1193:0:0.
- **Xlate PAT**—The idle time until a PAT translation slot is freed, between 0:0:30 and 0:5:0. The default is 30 seconds. You may want to increase the timeout if upstream routers reject new connections using a freed PAT port because the previous connection might still be open on the upstream device.
- **TCP Proxy Reassembly**—The idle timeout after which buffered packets waiting for reassembly are dropped, between 0:0:10 and 1193:0:0. The default is 1 minute (0:1:0).
- **ARP Timeout**—(Transparent mode only.) The number of seconds between ARP table rebuilds, from 60 to 4294967. The default is 14,400 seconds (4 hours).

Step 4 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure NTP Time Synchronization for Threat Defense

Use a Network Time Protocol (NTP) server to synchronize the clock settings on your devices. We recommend you configure all Firepower Threat Defenses managed by an FMC to use the same NTP server as the FMC. The Firepower Threat Defense gets its time directly from the configured NTP server. If the Firepower Threat Defense's configured NTP servers are not reachable for any reason, it synchronizes its time with the FMC.

The device supports NTPv4.



Note If you are deploying Firepower Threat Defense on the Firepower 4100/9300 chassis, you must configure NTP on the Firepower 4100/9300 chassis so that Smart Licensing will work properly and to ensure proper timestamps on device registrations. You should use the same NTP server for the Firepower 4100/9300 chassis and the FMC.

Before you begin

- If your organization has one or more NTP servers that your FTD can reach, use the same NTP server or servers for your devices that you have configured for Time Synchronization on the **System > Configuration** page on your FMC. Copy the specified value.
- If your device cannot reach an NTP server or your organization does not have one, you must use the **Via NTP from Defense Center** option as discussed in the following procedure.

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit the Firepower Threat Defense policy.

Step 2 Select **Time Synchronization**.

Step 3 Configure one of the following clock options:

- **Via NTP from Defense Center**—(Default). The managed device gets time from the NTP servers you configured for the FMC (except for authenticated NTP servers) and synchronizes time with those servers directly. However, if any of the following are true, the managed device synchronizes time from the FMC:
  - The FMC's NTP servers are not reachable by the device.
  - The FMC has no unauthenticated servers.
- **Via NTP from**—If your FMC is using NTP servers on the network, select this option and enter the fully-qualified DNS name (such as ntp.example.com), or IPv4 or IPv6 address, of the same NTP servers you specified in **System > Configuration > Time Synchronization**. If the NTP servers are not reachable, the FMC acts as an NTP server.

Step 4 Click **Save**.

What to do next

- Make sure the policy is assigned to your devices. See Setting Target Devices for a Platform Settings Policy, on page 929.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.
- If your system includes Classic devices, set up time synchronization for those devices. See Synchronize Time on Classic Devices with an NTP Server, on page 939.

History for Firepower Threat Defense Platform Settings

Feature	Version	Details
External Authentication for SSH and HTML removed	6.1.0	Due to changes to support converged management access, only local users are supported for SSH and HTML to data interfaces. Also, you can no longer SSH to the logical Diagnostic interface; instead you can SSH to the logical Management interface (which shares the same physical port). Previously, only external authentication was supported for SSH and HTML access to Diagnostic and data interfaces, while only local users were supported to the Management interface. New/Modified screen: **Devices** > **Platform Settings** > **External Authentication** Supported platforms: Firepower Threat Defense

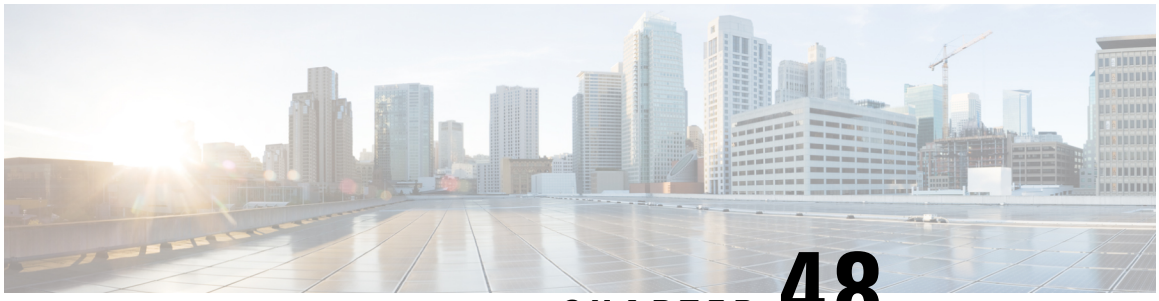
Feature	Version	Details
Firepower Threat Defense support	6.0.1	This feature was introduced. New/Modified screen: **Devices** > **Platform Settings** Supported platforms: Firepower Threat Defense



PART XIII

Network Address Translation (NAT)

- NAT Policy Management, on page 981
- NAT for 7000 and 8000 Series Devices, on page 987
- Network Address Translation (NAT) for Firepower Threat Defense, on page 1005



CHAPTER 48

NAT Policy Management

The following topics describe how to manage NAT policies for your Firepower System:

- Requirements and Prerequisites for NAT Policies, on page 981
- Managing NAT Policies, on page 981
- Creating NAT Policies, on page 982
- Configuring NAT Policies, on page 983
- Configuring NAT Policy Targets, on page 984
- Copying NAT Policies, on page 985

Requirements and Prerequisites for NAT Policies

Model Support

Any, but you must select the correct type of policy for the device model:

- **Firepower NAT** for 7000 & 8000 Series Series devices.
- **Threat Defense NAT** for FTD devices.

Supported Domains

Any

User Roles

Admin

Access Admin

Network Admin

Managing NAT Policies

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Administrators in ancestor domains can target NAT policies to devices in descendant domains, which descendant domains can use or replace with customized local policies. If a NAT policy targets devices in different descendant domains, administrators in the descendant domains can view information about target devices belonging to their domain only.

Procedure

Step 1 Choose **Devices** > **NAT** .

Step 2 Manage your NAT policies:

- Copy — Click **Copy** ( ) next to the policy you want to copy; see Copying NAT Policies, on page 985.
- Create — Click **New Policy**; see Creating NAT Policies, on page 982.
- Delete — Click **Delete** ( ) next to the policy you want to delete, then click **OK**. When prompted whether to continue, you are also informed if another user has unsaved changes in the policy.

  Caution After you have deployed a NAT policy to a managed device, you cannot delete the policy from the device. Instead, you must deploy a NAT policy with no rules to remove the NAT rules already present on the managed device. You also cannot delete a policy that is the last deployed policy on any of its target devices, even if it is out of date. Before you can delete the policy completely, you must deploy a different policy to those targets.

- Deploy—Click **Deploy**; see Deploy Configuration Changes, on page 346.
- Edit — Click **Edit** ( ); see Configuring NAT Policies, on page 983.
- Report—Click **Report** ( ); see Generating Current Policy Reports, on page 356.

Creating NAT Policies

When you create a new NAT policy you must, at minimum, give it a unique name. Although you are not required to identify policy targets at policy creation time, you must perform this step before you can deploy the policy. If you apply a NAT policy with no rules to a device, the system removes all NAT rules from that device.

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Administrators in ancestor domains can target NAT policies to devices in descendant domains, which descendant domains can use or replace with customized local policies. If a NAT policy targets devices in different descendant domains, administrators in the descendant domains can view information about target devices belonging to their domain only.

Procedure

Step 1 Choose **Devices** > **NAT** .

Step 2 From the **New Policy** drop-down list, choose one of the following:

- **Firepower NAT** for 7000 & 8000 Series devices.
- **Threat Defense NAT** for Firepower Threat Defense devices.

Step 3 Enter a unique **Name**.

In a multidomain deployment, policy names must be unique within the domain hierarchy. The system may identify a conflict with the name of a policy you cannot view in your current domain.

Step 4 Optionally, enter a **Description**.

Step 5 Choose the devices where you want to deploy the policy:

- Choose a device in the **Available Devices** list, and click **Add to Policy**.
- Click and drag a device from the **Available Devices** list to the **Selected Devices** list.
- Remove a device from the **Selected Devices** list by clicking **Delete** ( ) next to the device.

Step 6 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configuring NAT Policies

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Administrators in ancestor domains can target NAT policies to devices in descendant domains, which descendant domains can use or replace with customized local policies. If a NAT policy targets devices in different descendant domains, administrators in the descendant domains can view information about target devices belonging to their domain only.

If you change the type of an interface to a type that is not valid for use with a NAT policy that targets a device with that interface, the policy labels the interface as deleted. Click **Save** in the NAT policy to automatically remove the interface from the policy.

Note Rule attributes differ by NAT policy type. When adding or editing rules, click ? in the dialog box for more information, or see the relevant chapter: Network Address Translation (NAT) for Firepower Threat Defense, on page 1005 or NAT for 7000 and 8000 Series Devices, on page 987.

Procedure

Step 1 Choose **Devices** > **NAT** .

Step 2 Click **Edit** ( ) next to the NAT policy you want to modify.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Configure your NAT policies:

- To modify the policy name or description, click the **Name** or **Description** field, delete any characters as needed, then enter the new name or description. In a multidomain deployment, policy names must be unique within the domain hierarchy. The system may identify a conflict with the name of a policy you cannot view in your current domain.
- To manage policy targets, see Configuring NAT Policy Targets, on page 984.
- To save your policy changes, click **Save**.
- To add a rule to a policy, click **Add Rule**.
- To edit an existing rule, click **Edit** ( ) next to the rule.
- To delete a rule, click **Delete** ( ) next to the rule, then click **OK**.
- To enable or disable an existing rule, right-click a rule, choose **State**, and choose **Disable** or **Enable**.
- (Firepower NAT only) To display the configuration page for a specific rule attribute, click the name or value in the column for the condition on the row for the rule. For example, click the name or value in the **Source Networks** column to display the Source Network page for the selected rule.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configuring NAT Policy Targets

You can identify the managed devices you want to target with your policy while creating or editing a policy. You can search a list of available devices, 7000 or 8000 Series stacks, and high-availability pairs, and add them to a list of selected devices.

You cannot target stacked devices running different versions of the Firepower System (for example, if an upgrade on one of the devices fails).

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Administrators in ancestor domains can target NAT policies to devices in descendant domains, which descendant domains can use or replace with customized local policies. If a NAT policy targets devices in different descendant domains, administrators in the descendant domains can view information about target devices belonging to their domain only.

Procedure

Step 1 Choose **Devices** > **NAT** .

Step 2 Click **Edit** ( ) next to the NAT policy you want to modify.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Policy Assignments**.

Step 4 Do any of the following:

- To assign a device, stack, high-availability pair, or device group to the policy, select it in the **Available Devices** list and click **Add to Policy**. You can also drag and drop.
- To remove a device assignment, click **Delete** ( ) next to a device, stack, high-availability pair, or device group in the **Selected Devices** list.

Step 5 Click **OK**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Copying NAT Policies

You can make a copy of a NAT policy. The copy includes all policy rules and configurations.

In a multidomain deployment, you can copy policies from current and ancestor domains.

Procedure

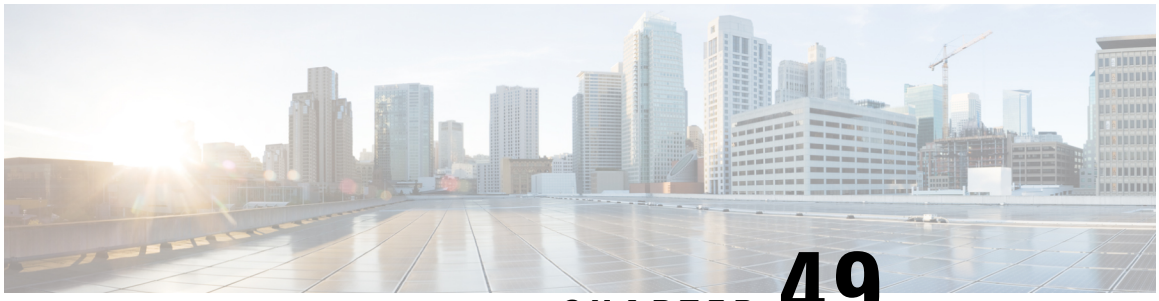
Step 1 Choose **Devices** > **NAT** .

Step 2 Click **Copy** ( ) next to the NAT policy you want to copy.

Step 3 Enter a unique **Name** for the policy.

In a multidomain deployment, policy names must be unique within the domain hierarchy. The system may identify a conflict with the name of a policy you cannot view in your current domain.

Step 4 Click **OK**.



CHAPTER 49

NAT for 7000 and 8000 Series Devices

The following topics describe how to configure NAT for 7000 and 8000 Series devices:

- NAT Policy Configuration, on page 987
- Rule Organization in a NAT Policy, on page 988
- Organizing NAT Rules, on page 989
- NAT Policy Rules Options, on page 990

NAT Policy Configuration

You can configure NAT policies in different ways to manage specific network needs. You can:

- *Expose an internal server to an external network.*

  In this configuration, you define a static translation from an external IP address to an internal IP address so the system can access an internal server from outside the network. Traffic sent to the server targets the external IP address or IP address and port, and is translated into the internal IP address or IP address and port. Return traffic from the server is translated back to the external address.

- *Allow an internal host/server to connect to an external application.*

  In this configuration, you define a static translation from an internal address to an external address. This definition allows the internal host or server to initiate a connection to an external application that is expecting the internal host or server to have a specific IP address and port. Therefore, the system cannot dynamically allocate the address of the internal host or server.

- *Hide private network addresses from an external network.*

  You can obscure your internal network addresses using either of the following configurations:

  - If you have a sufficient number of external IP addresses to satisfy your internal network needs, you can use a block of IP addresses. In this configuration, you create a dynamic translation that automatically converts the source IP address of any outgoing traffic to an unused IP address from your externally facing IP addresses.

  - If you have an insufficient number of external IP addresses to satisfy your internal network needs, you can use a limited block of IP addresses and port translation. In this configuration, you create a dynamic translation that automatically converts the source IP address and port of outgoing traffic to an unused IP address and port from your externally facing IP addresses.

Caution In 7000 or 8000 Series device high-availability pairs, only select an individual peer interface for a static NAT rule on a paired device if all networks affected by the NAT translations are private. Do **not** use configurations for static NAT rules affecting traffic between public and private networks.

NAT Policies Configuration Guidelines

To configure a NAT policy, you must give the policy a unique name and identify the devices, or *targets*, where you want to deploy the policy. You can also add, edit, delete, enable, and disable NAT rules. After you create or modify a NAT policy, you can deploy the policy to all or some targeted devices.

You can deploy NAT policies to a 7000 or 8000 Series device high-availability pair, including paired stacks, as you would a standalone device. However, you can define static NAT rules for interfaces on individual paired devices or the entire high-availability pair and use the interfaces in source zones. For dynamic rules, you can use only the interfaces on the whole high-availability pair in source or destination zones.

Caution In 7000 or 8000 Series device high-availability pairs, only select an individual peer interface for a static NAT rule on a paired device if all networks affected by the NAT translations are private. Do **not** use this configuration for static NAT rules affecting traffic between public and private networks.

If you configure dynamic NAT on a device high-availability pair without HA link interfaces established, both paired devices independently allocate dynamic NAT entries, and the system cannot synchronize the entries between devices.

You can deploy NAT policies to a device stack as you would a standalone device. If you establish a device stack from devices that were included in a NAT policy and had rules associated with interfaces from the secondary device that was a member of the stack, the interfaces from the secondary device remain in the NAT policy. You can save and deploy policies with the interfaces, but the rules do not provide any translation.

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain. Administrators in ancestor domains can target NAT policies to devices in descendant domains, which descendant domains can use or replace with customized local policies

Rule Organization in a NAT Policy

The Edit page for the NAT policy lists static NAT rules and dynamic NAT rules separately. The system sorts static rules alphabetically by name, and you cannot change the display order. You cannot create static rules with identical matching values. The system inspects static translations for a match before it inspects any dynamic translations.

Dynamic rules are processed in numerical order. The numeric position of each dynamic rule appears on the left side of the page next to the rule. You can move or insert dynamic rules and otherwise change the rule order. For example, if you move dynamic rule 10 under dynamic rule 3, rule 10 becomes rule 4 and all subsequent numbers increment accordingly.

A dynamic rule's position is important because the system compares packets to dynamic rules in the rules' numeric order on the policy Edit page. When a packet meets all the conditions of a dynamic rule, the system applies the conditions of that rule to the packet and ignores all subsequent rules for that packet.

You can specify a dynamic rule's numeric position when you add or edit a dynamic rule. You can also highlight a dynamic rule before adding a new dynamic rule to insert the new rule below the rule you highlighted.

You can select one or more dynamic rules by clicking a blank space in the row for the rule. You can drag and drop selected dynamic rules into a new location, thereby changing the position of the rules you moved and all subsequent rules.

You can cut or copy selected rules and paste them above or below an existing rule. You can only paste static rules in the Static Translations list and only dynamic rules in the Dynamic Translations list. You can also delete selected rules and insert new rules into any location in the list of existing rules.

You can display explanatory warnings to identify rules that will never match because they are preempted by preceding rules.

If you have access control policies in your deployment, the system does not translate traffic until it has passed through access control.

Organizing NAT Rules

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Any	Admin/Network Admin

Procedure

Step 1 Choose **Devices** > **NAT** .

Step 2 Click the edit icon ( ) next to the NAT policy you want to modify.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Organize your NAT rules:

- To choose a rule, click a blank area in the row for a rule.
- To clear rule selections, click the reload icon ( ) on the lower right side of the page. To clear individual rules, click a blank area in a rule's row while holding the Ctrl key.
- To cut or copy selected rules, right-click a blank area in the row for a selected rule, then select **Cut** or **Copy**.
- To paste rules you have cut or copied into the rule list, right-click a blank area in the row for a rule where you want to paste selected rules, then select **Paste above** or **Paste below**.
- To move selected rules, drag and drop selected rules beneath a new location, indicated by a horizontal blue line that appears above your pointer as you drag.
- To delete a rule, click the delete icon ( ) next to the rule, then click **OK**.
- To show warnings, click **Show Warnings**.

NAT Rule Warnings and Errors

The conditions of a NAT rule may preempt a subsequent rule from matching traffic. Any type of rule condition can preempt a subsequent rule.

A rule also preempts an identical subsequent rule where all configured conditions are the same. A subsequent rule would not be preempted if any condition were different.

If you create a rule that causes the NAT policy to fail upon deploy, an error icon ( ) appears next to the rule. An error occurs if there is a conflict in the static rules, or if you edit a network object used in the policy that now makes the policy invalid. For example, an error occurs if you change a network object to use only IPv6 addresses and the rule that uses that object no longer has any valid networks where at least one network is required. Error icons appear automatically; you do not have to click **Show Warnings**.

Showing and Hiding NAT Rule Warnings

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Any	Admin/Network Admin

Procedure

Step 1 Choose **Devices** > **NAT** .

Step 2 Click the edit icon ( ) next to the NAT policy you want to modify.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 To show warnings, click **Show Warnings**.

The page updates with an warning icon ( ) next to each preempted rule.

Step 4 To display the warning for a rule, hover your pointer over the warning icon ( ) next to a rule.
A message indicates which rule preempts the rule.

Step 5 To clear warnings, click **Hide Warnings**.
The page refreshes and the warnings disappear.

NAT Policy Rules Options

A NAT rule is simply a set of configurations and conditions that:

- qualifies network traffic
- specifies how the traffic that matches those qualifications is translated

You create and edit NAT rules from within an existing NAT policy. Each rule belongs to only one policy.

The web interface for adding or editing a rule is similar. You specify the rule name, state, type, and position (if dynamic) at the top of the page. You build conditions using the tabs on the left side of the page; each condition type has its own tab.

The following list summarizes the configurable components of a NAT rule.

Name

Give each rule a unique name. For static NAT rules, use a maximum of 22 characters. For dynamic NAT rules, use a maximum of 30 characters. You can use printable characters, including spaces and special characters, with the exception of the colon (:).

Rule State

By default, rules are enabled. If you disable a rule, the system does not use it to evaluate network traffic for translation. When viewing the list of rules in a NAT policy, disabled rules are grayed out, although you can still modify them.

Type

A rule's type determines how the system handles traffic that matches the rule's conditions. When you create and edit NAT rules, the configurable components vary according to rule type.

Position (Dynamic Rules Only)

Dynamic rules in a NAT policy are numbered, starting at 1. The system matches traffic to NAT rules in top-down order by ascending rule number.

When you add a rule to a policy, you specify its position by placing it **above** or **below** a specific rule, using rule numbers as a reference point. When editing an existing rule, you can **Move** the rule in a similar fashion.

Conditions

Rule conditions identify the specific traffic you want to translate. Conditions can match traffic by any combination of multiple attributes, including security zone, network, and transport protocol port.

Related Topics

Creating and Editing NAT Rules, on page 991

Creating and Editing NAT Rules

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Any	Admin/Network Admin

In a multidomain deployment, the system displays policies and rules created in the current domain, which you can edit. It also displays policies and rules created in ancestor domains, which you cannot edit. To view and edit rules created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Devices** > **NAT** .

Step 2 Click the edit icon ( ) next to the NAT policy where you want to add a rule.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Add a new rule or edit an existing rule:

- To add a new rule, click **Add Rule**.
- To edit an existing rule, click the edit icon ( ) next to the rule you want to edit.

Step 4 Enter a unique rule **Name**.

Step 5 Configure the following rule components:

- Specify whether the rule is **Enabled**.
- Specify a rule **Type**.
- Specify the rule position (dynamic rules only).
- Configure the rule's conditions.

Note Static rules must include an original destination network. Dynamic rules must include a translated source network.

Step 6 Click **Add**.

Step 7 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

NAT Rule Types

Every NAT rule has an associated type that:

- qualifies network traffic
- specifies how the traffic that matches those qualifications is translated

The following list summarizes the NAT rule types.

Static

Static rules provide one-to-one translations on destination networks and optionally port and protocol. When configuring static translations, you can configure source zones, destination networks, and destination ports. You cannot configure destination zones or source networks.

You **must** specify an original destination network. For destination networks, you can only select network objects and groups containing a single IP address or enter literal IP addresses that represent a single IP address. You can only specify a single original destination network and a single translated destination network.

Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

You can specify a single original destination port and a single translated destination port. You must specify an original destination network before you can specify an original destination port. In addition, you cannot specify a translated destination port unless you also specify an original destination port, and the translated value must match the protocol of the original value.

Caution For static NAT rules on a 7000 or 8000 Series device in a high-availability pair, only select an individual peer interface if all networks affected by the NAT translations are private. Do **not** use this configuration for static NAT rules affecting traffic between public and private networks.

Dynamic IP Only

Dynamic IP Only rules translate many-to-many source networks, but maintain port and protocol. When configuring dynamic IP only translations, you can configure zones, source networks, original destination networks, and original destination ports. You cannot configure translated destination networks or translated destination ports.

You **must** specify at least one translated source network. If the number of translated source network values is less than the number of original source networks, the system displays a warning on the rule that it is possible to run out of translated addresses before all original addresses are matched.

If there are multiple rules with conditions that match the same packet, the low priority rules become dead, meaning they can never be triggered. The system also displays warnings for dead rules. You can view tooltips to determine which rule supersedes the dead rule.

Note You can save and deploy policies with dead rules, but the rules cannot provide any translation.

In some instances, you may want to create rules with limited scope preceding rules with a broader scope. For example:

```
Rule 1: Match on address A and port A/Translate to address B
Rule 2: Match on address A/Translate to Address C
```

In this example, rule 1 matches some packets that also match rule 2. Therefore, rule 2 is not completely dead.

If you specify only original destination ports, you cannot specify translated destination ports.

Dynamic IP + Port

Dynamic IP and port rules translate many-to-one or many-to-many source networks and port and protocol. When configuring dynamic IP and port translations, you can configure zones, source networks, original

destination networks, and original destination ports. You cannot configure translated destination networks or translated destination ports.

You **must** specify at least one translated source network. If there are multiple rules with conditions that match the same packet, the low priority rules become dead, meaning they can never be triggered. The system also displays warnings for dead rules. You can view tool tips to determine which rule supersedes the dead rule.

Note You can save and deploy policies with dead rules, but the rules cannot provide any translation.

If you specify only original destination ports, you cannot specify translated destination ports.

Note If you create a dynamic IP and port rule, and the system passes traffic that does not use a port, no translation occurs for the traffic. For example, a ping (ICMP) from an IP address that matches the source network does not map, because ICMP does not use a port.

NAT Rule Condition Types

The following table summarizes the NAT rule condition types that can be configured based on the specified NAT rule type:

Table 79: Available NAT Rule Condition Types per NAT Rule Type

Condition	Static	Dynamic (IP Only or IP + Port)
Source Zones	Optional	Optional
Destination Zones	Not allowed	Optional
Original Source Networks	Not allowed	Optional
Translated Source Networks	Not allowed	**Required**
Original Destination Networks	**Required**	Optional
Translated Destination Networks	Optional; single address only	Not allowed
Original Destination Ports	Optional; single port only, and only allowed if you define the original destination network	Optional
Translated Destination Ports	Optional; single port only, and only allowed if you define the original destination port	Not allowed

NAT Rule Conditions and Condition Mechanics

You can add conditions to NAT rules to identify the type of traffic that matches the rule. For each condition type, you select conditions you want to add to a rule from a list of available conditions. When applicable, condition filters allow you to constrain available conditions. Lists of available and selected conditions may

be as short as a single condition or many pages long. You can search available conditions and display only those matching a typed name or value in a list that updates as you type.

Depending on the type of condition, lists of available conditions may be comprised of a combination of conditions provided directly by Cisco or configured using other Firepower System features, including objects created using the object manager (**Objects** > **Object Management**), objects created directly from individual conditions pages, and literal conditions.

NAT Rule Conditions

You can set a NAT rule to match traffic meeting any of the conditions described in the following table:

Table 80: NAT Rule Condition Types

Condition	Description
Zones	A configuration of one or more routed interfaces where you can deploy NAT policies. Zones provide a mechanism for classifying traffic on source and destination interfaces, and you can add source and destination zone conditions to rules.
Networks	Any combination of individual IP addresses, CIDR blocks, and prefix lengths, either specified explicitly or using network objects and groups.You can add source and destination network conditions to NAT rules.
Destination Ports	Transport protocol ports, including individual and group port objects you create based on transport protocols.

Adding Conditions to NAT Rules

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Any	Admin/Network Admin

Adding conditions to NAT rules is essentially the same for each type of condition. You choose from a list of available conditions on the left, and add the conditions you chose to one or two lists of selected conditions on the right.

For all condition types, you choose one or more individual available conditions by clicking on them to highlight them. You can either click a button between the two types of lists to add available conditions that you choose to your lists of selected conditions, or drag and drop available conditions that you choose into the list of selected conditions.

You can add up to 50 conditions of each type to a list of selected conditions. For example, you can add up to 50 source zone conditions, up to 50 destination zone conditions, up to 50 source network conditions, and so on, until you reach the upper limit for the appliance.

Procedure

Step 1 Choose **Devices** > **NAT** .

Step 2 Click the edit icon ( ) next to the NAT policy you want to modify.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Add Rule**.

Step 4 Enter a **Name** for the rule.

Step 5 Specify a **Type** for the rule.

Step 6 Click the tab for the type of condition you want to add to the rule.

Step 7 Take any of the following actions:

- To choose available conditions to add to a list of selected conditions, click the available condition.
- To choose all listed available conditions, right-click the row for any available condition, then click **Select All**.
- To choose a list of available conditions or filters, click inside the **Search** field and enter a search string. The list updates as you type to display matching items.

  You can search on object names and on the values configured for objects. For example, if you have an individual network object named `Texas Office` with the configured value `192.168.3.0/24`, and the object is included in the group object `US Offices`, you can display both objects by entering a partial or complete search string such as `Tex`, or by entering a value such as `3`.

- To clear a search when searching available conditions or filters, click the reload icon ( ) above the Search field or the clear icon ( ) in the Search field.
- To add selected zone conditions from a list of available conditions to a list of selected source or destination conditions, click **Add to Source** or **Add to Destination**.
- To add selected network and port conditions from a list of available conditions to a list of selected original or translated conditions, click **Add to Original** or **Add to Translated**.
- To drag and drop selected available conditions into a list of selected conditions, click a selected condition, then drag and drop into the list of selected conditions.
- To add a literal condition to a list of selected conditions using a literal field, click to remove the prompt from the literal field, enter the literal condition, and click **Add**. Network conditions provide a field for adding literal conditions.
- To add a literal condition to a list of selected conditions using a drop-down list, choose a condition from the drop-down list, then click **Add**. Port conditions provide a drop-down list for adding literal conditions.
- To add an individual object or condition filter so you can then choose it from the list of available conditions, click the add icon ( ).
- To delete a single condition from a list of selected conditions, click the delete icon ( ) next to the condition.
- To delete a condition from a list of selected conditions, right-click to highlight the row for a selected condition, then click **Delete**.

Step 8 Click **Add** to save your configuration.

Literal Conditions in NAT Rules

You can add a literal value to the list of original and translated conditions for the following condition types:

- Networks
- Ports

For network conditions, you type the literal value in a configuration field below the list of original or translated conditions.

In the case of port conditions, you choose a protocol from a drop-down list. When the protocol is `All`, or `TCP` or `UDP`, you enter a port number in a configuration field.

Each relevant conditions page provides the controls needed to add literal values. Values you enter in a configuration field appear as red text if the value is invalid, or until it is recognized as valid. Values change to blue text as you type when they are recognized as valid. A grayed **Add** button activates when a valid value is recognized. Literal values you add appear immediately in the list of selected conditions.



Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Objects in NAT Rule Conditions

Objects that you create in the object manager (**Objects** > **Object Management**) are immediately available for you to select from relevant lists of available NAT rule conditions.

You can also create objects on-the-fly from the NAT policy. A control on relevant conditions pages provides access to the same configuration controls that you use in the object manager.

Individual objects created on-the-fly appear immediately in the list of available objects. You can add them to the current rule, and to other existing and future rules. On the relevant conditions page, and also on the policy Edit page, you can hover your pointer over an individual object to display the contents of the object, and over a group object to display the number of individual objects in the group.

Zone Conditions in NAT Rules

The security zones on your system are comprised of interfaces on your managed devices. Zones that you add to a NAT rule target the rule to devices on your network that have routed or hybrid interfaces in those zones. You can only add security zones with routed or hybrid interfaces as conditions for NAT rules.

You can add either zones or standalone interfaces that are currently assigned to a virtual router to NAT rules.

If there are devices with un-deployed device configurations, the Zones page displays a warning icon () at the top of the available zones list, indicating that only deployed zones and interfaces are displayed. You can click the arrow icon () next to a zone to collapse or expand the zone to hide or view its interfaces.

If an interface is on a 7000 or 8000 Series device in a high-availability pair, the available zones list displays an additional branch from that interface with the other interfaces in the high-availability pair as children of

the primary interface on the active device in the high-availability pair. You can also click the arrow icon ( ) to collapse or expand the paired device interfaces to hide or view its interfaces.

Note You can save and deploy policies with disabled interfaces, but the rules cannot provide any translation until the interfaces are enabled.

The two lists on the right are the source and destination zones used for matching purposes by the NAT rules. If the rule already has values configured, these lists display the existing values when you edit the rule. If the source zones list is empty, the rule matches traffic *from* any zone or interface. If the destination zones list is empty, the rule matches traffic *to* any zone or interface.

The system displays warnings for rules with zone combinations that never trigger on a targeted device.

Note You can save and deploy policies with these zone combinations, but the rules will not provide any translation.

You can add individual interfaces by selecting an item in a zone or by selecting a standalone interface. You can only add interfaces in a zone if the zone it is assigned to has not already been added to a source zones or destination zones list. These individually selected interfaces are not affected by changes to zones, even if you remove them and add them to a different zone. If an interface is the primary member of a high-availability pair and you are configuring a dynamic rule, you can add only the primary interface to the source zones or destination zones list. For static rules, you can add individual high-availability pair member interfaces to the source zones list. You can only add a primary high-availability pair interface to a list if none of its children have been added, and you can only add individual high-availability pair interfaces if the primary has not been added.

If you add a zone, the rule uses all interfaces associated with the zone. If you add or remove an interface from the zone, the rule will not use the updated version of the zone until the device configuration has been re-deployed to the devices where the interfaces reside.

Note In a static NAT rule, you can add only source zones. In a dynamic NAT rule, you can add both source and destination zones.

Adding Zone Conditions to NAT Rules

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Any	Admin/Network Admin

Procedure

Step 1 Choose **Devices** > **NAT** .

Step 2 Click the edit icon ( ) next to the NAT policy you want to modify.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Add Rule**.

Step 4 Enter a **Name** for the rule.

Step 5 Specify a **Type** for the rule.

Step 6 Click the **Zones** tab.

Step 7 Click a zone or interface in the **Available Zones** list.

Step 8 You have the following choices:

- To match traffic by source zone, click **Add to Source**.
- To match traffic by destination zone, click **Add to Destination**.

Note You can add only source zones to static NAT rules. Additionally, while you can add disabled interfaces to a NAT rule, the rule does not provide any translation.

Step 9 Click **Add** to save the new rule.

Step 10 Click **Save** to save the changed policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Source Network Conditions in Dynamic NAT Rules

You configure the matching values and translation values of the source IP address for packets. If the original source network is not configured, then any source IP address matches the dynamic NAT rule. Note that you cannot configure source networks for static NAT rules. If a packet matches the NAT rule, the system uses the values in the translated source network to assign the new value for the source IP address. For dynamic rules, you must configure a translated source network with at least one value.

Caution If a network object or object group is being used by a NAT rule, and you change or delete the object or group, it can cause the rule to become invalid.

You can add any of the following kinds of source network conditions to a dynamic NAT rule:

- individual and group network objects that you have created using the object manager
- individual network objects that you add from the Source Network conditions page, and can then add to your rule and to other existing and future rules
- literal, single IP addresses, ranges, or address blocks

Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Adding Network Conditions to a Dynamic NAT Rule

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Any	Admin/Network Admin

When you update the network conditions in a dynamic rule in use in a deployed policy, the system drops any network sessions using the existing translated address pool.

Procedure

Step 1 Choose **Devices** > **NAT** .

Step 2 Click the edit icon ( ) next to the NAT policy you want to modify.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Add Rule**.

Step 4 Enter a **Name** for the rule.

Step 5 Specify a dynamic **Type** for the rule:

- **Dynamic IP Only**
- **Dynamic IP + Port**

Step 6 Click the **Source Networks** tab.

Step 7 Optionally, add an individual network object to the **Available Networks** list by clicking the add icon ( ) above the list.

You can add multiple IP addresses, CIDR blocks, and prefix lengths to each network object.

Step 8 Click a condition in the **Available Networks** list.

Step 9 You have the following choices:

- To match traffic by original source network, click **Add to Original**.
- To specify the translation value for traffic that matches the translated source network, click **Add to Translated**.

Step 10 To add a literal IP address, range, or address block:

a) Click the **Enter an IP address** prompt below the **Original Source Network** or **Translated Source Network** list.

b) Enter an IP address, range, or address block.

You add ranges in the following format: lower IP address-upper IP address. For example: `179.13.1.1-179.13.1.10`.

Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

c) Click **Add** next to the value you entered.

Step 11 Click **Add** to save the rule.

Step 12 Click **Save** to save the changed policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Destination Network Conditions in NAT Rules

You configure the matching values and translation values of the destination IP address for packets. Note that you cannot configure translated destination networks for dynamic NAT rules.

Because static NAT rules are one-to-one translations, the **Available Networks** list contains only network objects and groups that contain only a single IP address. For static translations, you can add only a single object or literal value to both the **Original Destination Network** or **Translated Destination Network** lists.

Caution If a network object or object group is being used by a NAT rule, and you change or delete the object or group, it can cause the rule to become invalid.

You can add any of the following kinds of destination network conditions to a NAT rule:

- individual and group network objects that you have created using the object manager
- individual network objects that you add from the Destination Network conditions page, and can then add to your rule and to other existing and future rules
- literal, single IP addresses, range, or address blocks

  For static NAT rules, you can add only a CIDR with subnet mask `/32`, and only if there is not already a value in the list.

Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Adding Destination Network Conditions to NAT Rules

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Any	Admin/Network Admin

When you update the network conditions in a dynamic rule in use in a deployed policy, the system drops any network sessions using the existing translated address pool.

Procedure

Step 1 Choose **Devices** > **NAT** .

Step 2 Click the edit icon ( ) next to the NAT policy you want to modify.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Add Rule**.

Step 4 Enter a **Name** for the rule.

Step 5 Specify a **Type** for the rule.

Step 6 Click the **Destination Network** tab.

Step 7 Optionally, add an individual network object to the **Available Networks** list by clicking the add icon ( ) above the list.

For dynamic rules, you can add multiple IP addresses, CIDR blocks, and prefix lengths to each network object. For static rules, you can add only a single IP address.

Step 8 Click a condition or object in the **Available Networks** list.

Step 9 You have the following choices:

- To match traffic by original destination network, click **Add to Original**.
- To specify the translation value for traffic that matches the translated destination network, click **Add to Translated**.

Step 10 Optionally, click the **Enter an IP address** prompt below the **Original Destination Network** or **Translated Destination Network** list, enter an IP address or address block, and click **Add**.

Step 11 Click **Add**.

Step 12 Click **Save** to save the changes to the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Port Conditions in NAT Rules

You can add a port condition to a rule to match network traffic based on the original and translated destination port and transport protocol for translation. If the original port is not configured, any destination port matches the rule. If a packet matches the NAT rule and a translated destination port is configured, the system translates the port into that value. Note that for dynamic rules, you can specify only the original destination port. For static rules, you can define a translated destination port, but only with an object with the same protocol as the original destination port object or literal value.

The system matches the destination port against the value of the port object or literal port in the original destination port list for static rules, or multiple values for dynamic rules.

Because static NAT rules are one-to-one translations, the **Available Ports** list contains only port objects and groups that contain only a single port. For static translations, you can add only a single object or literal value to both the **Original Port** or **Translated Port** lists.

For dynamic rules, you can add a range of ports. For example, when specifying the original destination port, you can add `1000-1100` as a literal value.

Caution If a port object or object group is being used by a NAT rule, and you change or delete the object or group, it can cause the rule to become invalid.

You can add any of the following kinds of port conditions to a NAT rule:

- individual and group port objects that you have created using the object manager
- individual port objects that you add from the Destination Ports conditions page, and can then add to your rule and to other existing and future rules
- literal port values, consisting of a TCP, UDP, or All (TCP and UDP) transport protocol and a port

Adding Port Conditions to NAT Rules

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Any	Admin/Network Admin

Procedure

Step 1 Choose **Devices** > **NAT** .

Step 2 Click the edit icon ( ) next to the NAT policy you want to modify.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Add Rule**.

Step 4 Enter a **Name** for the rule.

Step 5 Specify a **Type** for the rule.

Step 6 Click the **Destination Port** tab.

Step 7 Optionally, add an individual port object to the **Available Ports** list by clicking the add icon ( ) above the list.

You can identify a single port or a port range in each port object that you add. You can then choose objects you added as conditions for your rule. For static rules, you can use only port objects with single ports.

Step 8 Click a condition in the **Available Ports** list.

Step 9 You have the following choices:

- Click **Add to Original**.
- Click **Add to Translated**.
- Drag and drop available ports into a list.

Step 10 To add a literal port:

a) Choose an entry from the **Protocol** drop-down list beneath the **Original Port** or **Translated Port** lists.
b) Enter a port.
c) Click **Add**.

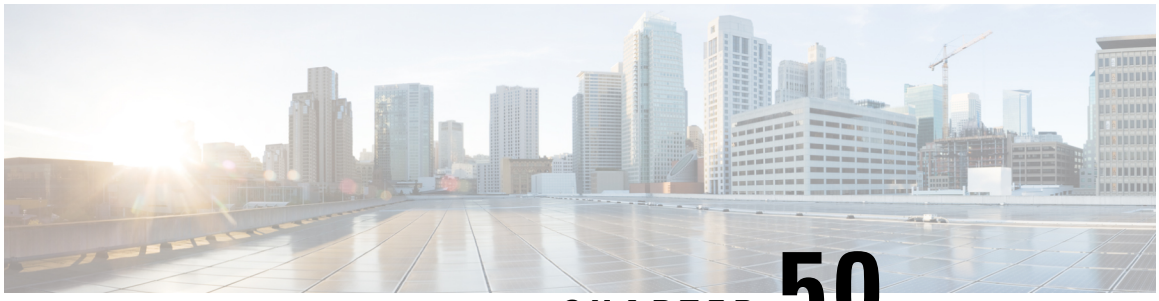
For dynamic rules, you can specify a single port or a range.

Step 11 Click **Add**.

Step 12 Click **Save** to save the changes to the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.



CHAPTER 50

Network Address Translation (NAT) for Firepower Threat Defense

The following topics explain Network Address Translation (NAT) and how to configure it on Firepower Threat Defense devices.

- Why Use NAT?, on page 1005
- NAT Basics, on page 1006
- Guidelines for NAT, on page 1014
- Configure NAT for Threat Defense, on page 1019
- Translating IPv6 Networks, on page 1053
- Monitoring NAT, on page 1063
- Examples for NAT, on page 1063
- History for FTD NAT, on page 1101

Why Use NAT?

Each computer and device within an IP network is assigned a unique IP address that identifies the host. Because of a shortage of public IPv4 addresses, most of these IP addresses are private, not routable anywhere outside of the private company network. RFC 1918 defines the private IP addresses you can use internally that should not be advertised:

- 10.0.0.0 through 10.255.255.255
- 172.16.0.0 through 172.31.255.255
- 192.168.0.0 through 192.168.255.255

One of the main functions of NAT is to enable private IP networks to connect to the Internet. NAT replaces a private IP address with a public IP address, translating the private addresses in the internal private network into legal, routable addresses that can be used on the public Internet. In this way, NAT conserves public addresses because it can be configured to advertise at a minimum only one public address for the entire network to the outside world.

Other functions of NAT include:

- Security—Keeping internal IP addresses hidden discourages direct attacks.
- IP routing solutions—Overlapping IP addresses are not a problem when you use NAT.

- Flexibility—You can change internal IP addressing schemes without affecting the public addresses available externally; for example, for a server accessible to the Internet, you can maintain a fixed IP address for Internet use, but internally, you can change the server address.
- Translating between IPv4 and IPv6 (Routed mode only) —If you want to connect an IPv6 network to an IPv4 network, NAT lets you translate between the two types of addresses.



Note NAT is not required. If you do not configure NAT for a given set of traffic, that traffic will not be translated, but will have all of the security policies applied as normal.

NAT Basics

The following topics explain some of the basics of NAT.

NAT Terminology

This document uses the following terminology:

- Real address/host/network/interface—The real address is the address that is defined on the host, before it is translated. In a typical NAT scenario where you want to translate the inside network when it accesses the outside, the inside network would be the "real" network. Note that you can translate any network connected to the device, not just an inside network. Therefore if you configure NAT to translate outside addresses, "real" can refer to the outside network when it accesses the inside network.
- Mapped address/host/network/interface—The mapped address is the address that the real address is translated to. In a typical NAT scenario where you want to translate the inside network when it accesses the outside, the outside network would be the "mapped" network.



Note During address translation, IP addresses configured for the device interfaces are not translated.

- Bidirectional initiation—Static NAT allows connections to be initiated *bidirectionally*, meaning both to the host and from the host.
- Source and destination NAT—For any given packet, both the source and destination IP addresses are compared to the NAT rules, and one or both can be translated/untranslated. For static NAT, the rule is bidirectional, so be aware that "source" and "destination" are used in commands and descriptions throughout this guide even though a given connection might originate at the "destination" address.

NAT Types

You can implement NAT using the following methods:

- Dynamic NAT—A group of real IP addresses are mapped to a (usually smaller) group of mapped IP addresses, on a first come, first served basis. Only the real host can initiate traffic. See Dynamic NAT, on page 1023.
- Dynamic Port Address Translation (PAT)—A group of real IP addresses are mapped to a single IP address using a unique source port of that IP address. See Dynamic PAT, on page 1028.
- Static NAT—A consistent mapping between a real and mapped IP address. Allows bidirectional traffic initiation. See Static NAT, on page 1035.
- Identity NAT—A real address is statically translated to itself, essentially bypassing NAT. You might want to configure NAT this way when you want to translate a large group of addresses, but then want to exempt a smaller subset of addresses. See Identity NAT, on page 1043.

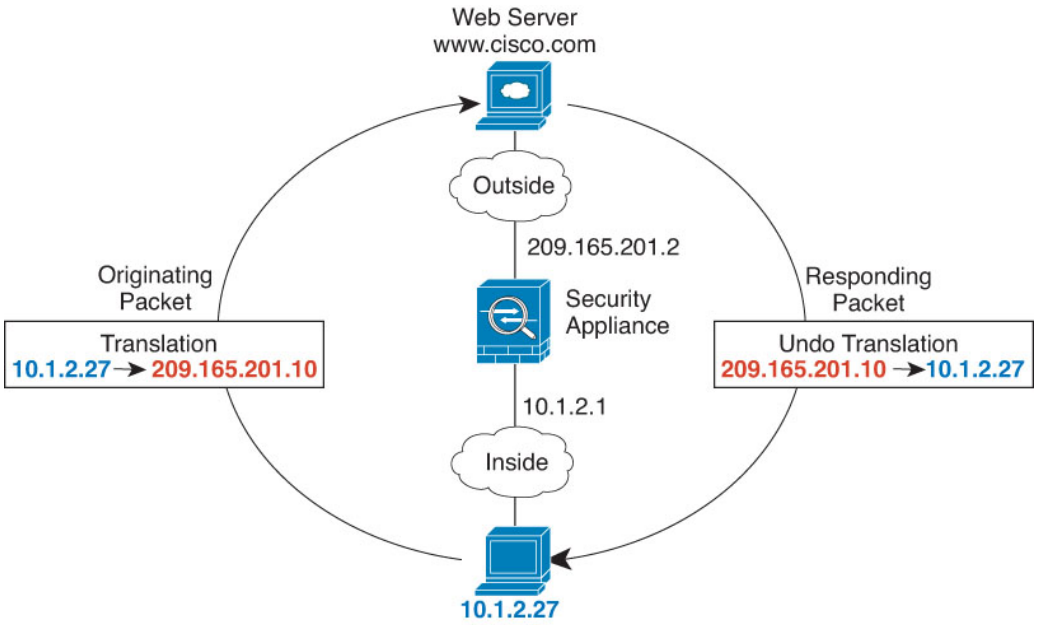
NAT in Routed and Transparent Mode

You can configure NAT in both routed and transparent firewall mode. You cannot configure NAT for interfaces operating in inline, inline tap, or passive modes. The following sections describe typical usage for each firewall mode.

NAT in Routed Mode

The following figure shows a typical NAT example in routed mode, with a private network on the inside.

Figure 32: NAT Example: Routed Mode



1. When the inside host at 10.1.2.27 sends a packet to a web server, the real source address of the packet, 10.1.2.27, is translated to a mapped address, 209.165.201.10.
2. When the server responds, it sends the response to the mapped address, 209.165.201.10, and the FTD device receives the packet because the FTD device performs proxy ARP to claim the packet.

3. The FTD device then changes the translation of the mapped address, 209.165.201.10, back to the real address, 10.1.2.27, before sending it to the host.

NAT in Transparent Mode or Within a Bridge Group

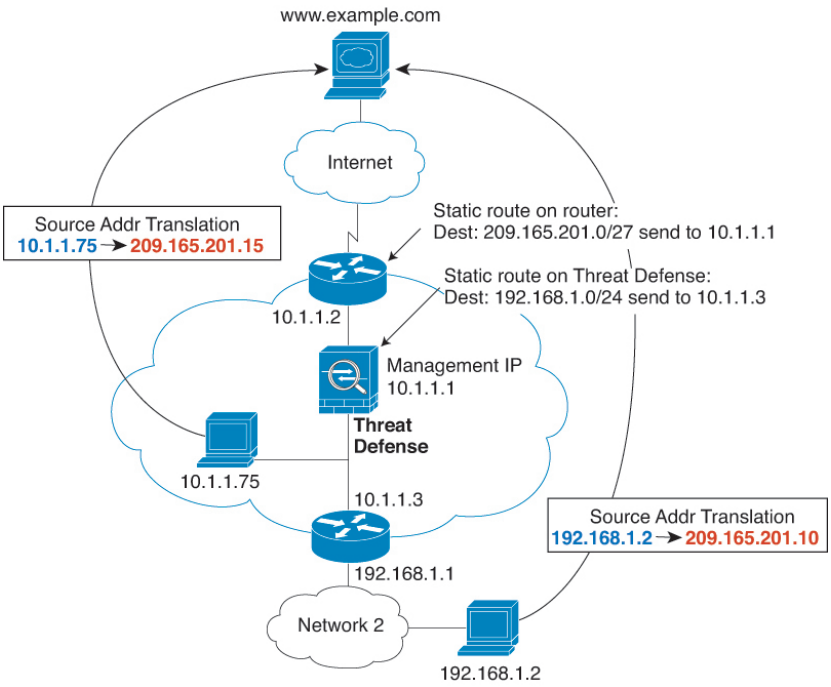
Using NAT in transparent mode eliminates the need for the upstream or downstream routers to perform NAT for their networks. It can perform a similar function within a bridge group in routed mode.

NAT in transparent mode, or in routed mode between members of the same bridge group, has the following requirements and limitations:

- You cannot configure interface PAT when the mapped address is a bridge group member interface, because there is no IP address attached to the interface.
- ARP inspection is not supported. Moreover, if for some reason a host on one side of the FTD sends an ARP request to a host on the other side of the FTD, and the initiating host real address is mapped to a different address on the same subnet, then the real address remains visible in the ARP request.
- Translating between IPv4 and IPv6 networks is not supported. Translating between two IPv6 networks, or between two IPv4 networks is supported.

The following figure shows a typical NAT scenario in transparent mode, with the same network on the inside and outside interfaces. The transparent firewall in this scenario is performing the NAT service so that the upstream router does not have to perform NAT.

Figure 33: NAT Example: Transparent Mode



1. When the inside host at 10.1.1.75 sends a packet to a web server, the real source address of the packet, 10.1.1.75, is changed to a mapped address, 209.165.201.15.

2. When the server responds, it sends the response to the mapped address, 209.165.201.15, and the FTD receives the packet because the upstream router includes this mapped network in a static route directed to the FTD management IP address.
3. The FTD then undoes the translation of the mapped address, 209.165.201.15, back to the real address, 10.1.1.1.75. Because the real address is directly-connected, the FTD sends it directly to the host.
4. For host 192.168.1.2, the same process occurs, except for returning traffic, the FTD looks up the route in its routing table and sends the packet to the downstream router at 10.1.1.3 based on the FTD static route for 192.168.1.0/24.

Auto NAT and Manual NAT

You can implement address translation in two ways: *auto NAT* and *manual NAT*.

We recommend using auto NAT unless you need the extra features that manual NAT provides. It is easier to configure auto NAT, and it might be more reliable for applications such as Voice over IP (VoIP). (For VoIP, you might see a failure in the translation of indirect addresses that do not belong to either of the objects used in the rule.)

Auto NAT

All NAT rules that are configured as a parameter of a network object are considered to be *auto NAT* rules. This is a quick and easy way to configure NAT for a network object. You cannot create these rules for a group object, however.

Although these rules are configured as part of the object itself, you cannot see the NAT configuration in the object definition through the object manager.

When a packet enters an interface, both the source and destination IP addresses are checked against the auto NAT rules. The source and destination address in the packet can be translated by separate rules if separate matches are made. These rules are not tied to each other; different combinations of rules can be used depending on the traffic.

Because the rules are never paired, you cannot specify that sourceA/destinationA should have a different translation than sourceA/destinationB. Use manual NAT for that kind of functionality, where you can identify the source and destination address in a single rule.

Manual NAT

Manual NAT lets you identify both the source and destination address in a single rule. Specifying both the source and destination addresses lets you specify that sourceA/destinationA can have a different translation than sourceA/destinationB.



Note For static NAT, the rule is bidirectional, so be aware that “source” and “destination” are used in commands and descriptions throughout this guide even though a given connection might originate at the “destination” address. For example, if you configure static NAT with port address translation, and specify the source address as a Telnet server, and you want all traffic going to that Telnet server to have the port translated from 2323 to 23, then you must specify the *source* ports to be translated (real: 23, mapped: 2323). You specify the source ports because you specified the Telnet server address as the source address.

The destination address is optional. If you specify the destination address, you can either map it to itself (identity NAT), or you can map it to a different address. The destination mapping is always a static mapping.

Comparing Auto NAT and Manual NAT

The main differences between these two NAT types are:

- How you define the real address.
  - Auto NAT—The NAT rule becomes a parameter for a network object. The network object IP address serves as the original (real) address.
  - Manual NAT—You identify a network object or network object group for both the real and mapped addresses. In this case, NAT is not a parameter of the network object; the network object or group is a parameter of the NAT configuration. The ability to use a network object *group* for the real address means that manual NAT is more scalable.
- How source and destination NAT is implemented.
  - Auto NAT— Each rule can apply to either the source or destination of a packet. So two rules might be used, one for the source IP address, and one for the destination IP address. These two rules cannot be tied together to enforce a specific translation for a source/destination combination.
  - Manual NAT—A single rule translates both the source and destination. A packet matches one rule only, and further rules are not checked. Even if you do not configure the optional destination address, a matching packet still matches one manual NAT rule only. The source and destination are tied together, so you can enforce different translations depending on the source/destination combination. For example, sourceA/destinationA can have a different translation than sourceA/destinationB.
- Order of NAT Rules.
  - Auto NAT—Automatically ordered in the NAT table.
  - Manual NAT—Manually ordered in the NAT table (before or after auto NAT rules).

NAT Rule Order

Auto NAT and manual NAT rules are stored in a single table that is divided into three sections. Section 1 rules are applied first, then section 2, and finally section 3, until a match is found. For example, if a match is found in section 1, sections 2 and 3 are not evaluated. The following table shows the order of rules within each section.

Table 81: NAT Rule Table

Table Section	Rule Type	Order of Rules within the Section
Section 1	Manual NAT	Applied on a first match basis, in the order they appear in the configuration. Because the first match is applied, you must ensure that specific rules come before more general rules, or the specific rules might not be applied as desired. By default, manual NAT rules are added to section 1. By "specific rules first," we mean: • Static rules should come before dynamic rules. • Rules that include destination translation should come before rules with source translation only. If you cannot eliminate overlapping rules, where more than one rule might apply based on the source or destination address, be especially careful to follow these recommendations.
Section 2	Auto NAT	If a match in section 1 is not found, section 2 rules are applied in the following order: **1.** Static rules. **2.** Dynamic rules. **Within each rule type, the following ordering guidelines are used:** **1.** Quantity of real IP addresses—From smallest to largest. For example, an object with one address will be assessed before an object with 10 addresses. **2.** For quantities that are the same, then the IP address number is used, from lowest to highest. For example, 10.1.1.0 is assessed before 11.1.1.0. **3.** If the same IP address is used, then the name of the network object is used, in alphabetical order. For example, abracadabra is assessed before catwoman.
Section 3	Manual NAT	If a match is still not found, section 3 rules are applied on a first match basis, in the order they appear in the configuration. This section should contain your most general rules. You must also ensure that any specific rules in this section come before general rules that would otherwise apply.

For section 2 rules, for example, you have the following IP addresses defined within network objects:

- 192.168.1.0/24 (static)
- 192.168.1.0/24 (dynamic)
- 10.1.1.0/24 (static)

- 192.168.1.1/32 (static)
- 172.16.1.0/24 (dynamic) (object def)
- 172.16.1.0/24 (dynamic) (object abc)

The resultant ordering would be:

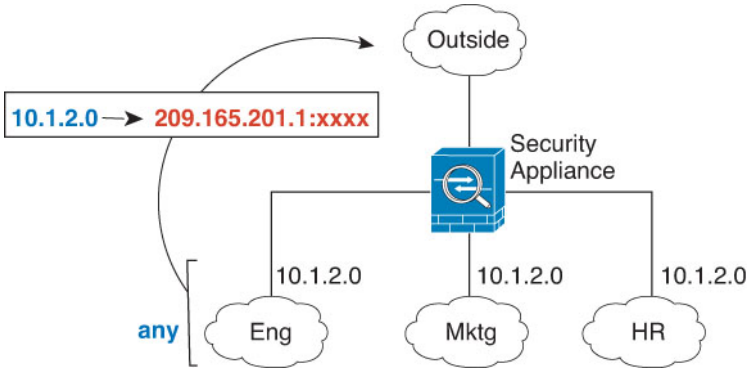
- 192.168.1.1/32 (static)
- 10.1.1.0/24 (static)
- 192.168.1.0/24 (static)
- 172.16.1.0/24 (dynamic) (object abc)
- 172.16.1.0/24 (dynamic) (object def)
- 192.168.1.0/24 (dynamic)

NAT Interfaces

Except for bridge group member interfaces, you can configure a NAT rule to apply to any interface (in other words, all interfaces), or you can identify specific real and mapped interfaces. You can also specify any interface for the real address, and a specific interface for the mapped address, or vice versa.

For example, you might want to specify any interface for the real address and specify the outside interface for the mapped address if you use the same private addresses on multiple interfaces, and you want to translate them all to the same global pool when accessing the outside.

Figure 34: Specifying Any Interface



However, the concept of "any" interface does not apply to bridge group member interfaces. When you specify "any" interface, all bridge group member interfaces are excluded. Thus, to apply NAT to bridge group members, you must specify the member interface. This could result in many similar rules where only one interface is different. You cannot configure NAT for the Bridge Virtual Interface (BVI) itself, you can configure NAT for member interfaces only.

Note You cannot configure NAT for interfaces operating in inline, inline tap, or passive modes. When specifying interfaces, you do so indirectly by selecting the interface object that contains the interface.

Configuring Routing for NAT

The Firepower Threat Defense device needs to be the destination for any packets sent to the translated (mapped) address.

When sending packets, the device uses the destination interface if you specify one, or a routing table lookup if you do not, to determine the egress interface. For identity NAT, you have the option to use a route lookup even if you specify a destination interface.

The type of routing configuration needed depends on the type of mapped address, as explained in the following topics.

Addresses on the Same Network as the Mapped Interface

If you use addresses on the same network as the destination (mapped) interface, the FTD device uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the FTD device does not have to be the gateway for any additional networks. This solution is ideal if the outside network contains an adequate number of free addresses, a consideration if you are using a 1:1 translation like dynamic NAT or static NAT. Dynamic PAT greatly extends the number of translations you can use with a small number of addresses, so even if the available addresses on the outside network is small, this method can be used. For PAT, you can even use the IP address of the mapped interface.



Note If you configure the mapped interface to be any interface, and you specify a mapped address on the same network as one of the mapped interfaces, then if an ARP request for that mapped address comes in on a *different* interface, then you need to manually configure an ARP entry for that network on the ingress interface, specifying its MAC address. Typically, if you specify any interface for the mapped interface, then you use a unique network for the mapped addresses, so this situation would not occur. Configure the ARP table in the ingress interface's **Advanced** settings.

Addresses on a Unique Network

If you need more addresses than are available on the destination (mapped) interface network, you can identify addresses on a different subnet. The upstream router needs a static route for the mapped addresses that points to the FTD device.

Alternatively for routed mode, you can configure a static route on the FTD device for the mapped addresses using any IP address on the destination network as the gateway, and then redistribute the route using your routing protocol. For example, if you use NAT for the inside network (10.1.1.0/24) and use the mapped IP address 209.165.201.5, then you can configure a static route for 209.165.201.5 255.255.255.255 (host address) to the 10.1.1.99 gateway that can be redistributed.

For transparent mode, if the real host is directly-connected, configure the static route on the upstream router to point to the FTD device: specify the bridge group IP address. For remote hosts in transparent mode, in the static route on the upstream router, you can alternatively specify the downstream router IP address.

The Same Address as the Real Address (Identity NAT)

The default behavior for identity NAT has proxy ARP enabled, matching other static NAT rules. You can disable proxy ARP if desired. You can also disable proxy ARP for regular static NAT if desired, in which case you need to be sure to have proper routes on the upstream router.

Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues. For example, if you configure a broad identity NAT rule for “any” IP address, then leaving proxy ARP enabled can cause problems for hosts on the network directly connected to the mapped interface. In this case, when a host on the mapped network wants to communicate with another host on the same network, then the address in the ARP request matches the NAT rule (which matches “any” address). The FTD device will then proxy ARP for the address, even though the packet is not actually destined for the FTD device. (Note that this problem occurs even if you have a manual NAT rule; although the NAT rule must match both the source and destination addresses, the proxy ARP decision is made only on the “source” address). If the FTD device ARP response is received before the actual host ARP response, then traffic will be mistakenly sent to the FTD device.

Guidelines for NAT

The following topics provide detailed guidelines for implementing NAT.

Firewall Mode Guidelines for NAT

NAT is supported in routed and transparent firewall mode.

However, configuring NAT on bridge group member interfaces (interfaces that are part of a Bridge Group Virtual Interface, or BVI) has the following restrictions:

- When configuring NAT for the members of a bridge group, you specify the member interface. You cannot configure NAT for the bridge group interface (BVI) itself.
- When doing NAT between bridge group member interfaces, you must specify the real and mapped addresses. You cannot specify “any” as the interface.
- You cannot configure interface PAT when the mapped address is a bridge group member interface, because there is no IP address attached to the interface.
- You cannot translate between IPv4 and IPv6 networks (NAT64/46) when the source and destination interfaces are members of the same bridge group. Static NAT/PAT 44/66, dynamic NAT44/66, and dynamic PAT44 are the only allowed methods; dynamic PAT66 is not supported. However, you can do NAT64/46 between members of different bridge groups, or between a bridge group member (source) and standard routed interface (destination).



Note You cannot configure NAT for interfaces operating in inline, inline tap, or passive modes.

IPv6 NAT Guidelines

NAT supports IPv6 with the following guidelines and restrictions.

- For standard routed mode interfaces, you can also translate between IPv4 and IPv6.
- You cannot translate between IPv4 and IPv6 for interfaces that are members of the same bridge group. You can translate between two IPv6 or two IPv4 networks only. This restriction does not apply when the interfaces are members of different bridge groups, or between a bridge group member and a standard routed interface.

- You cannot use dynamic PAT for IPv6 (NAT66) when translating between interfaces in the same bridge group. This restriction does not apply when the interfaces are members of different bridge groups, or between a bridge group member and a standard routed interface.
- For static NAT, you can specify an IPv6 subnet up to /64. Larger subnets are not supported.
- When using FTP with NAT46, when an IPv4 FTP client connects to an IPv6 FTP server, the client must use either the extended passive mode (EPSV) or extended port mode (EPRT); PASV and PORT commands are not supported with IPv6.

IPv6 NAT Best Practices

You can use NAT to translate between IPv6 networks, and also to translate between IPv4 and IPv6 networks (routed mode only). We recommend the following best practices:

- NAT66 (IPv6-to-IPv6)—We recommend using static NAT. Although you can use dynamic NAT or PAT, IPv6 addresses are in such large supply, you do not have to use dynamic NAT. If you do not want to allow returning traffic, you can make the static NAT rule unidirectional (manual NAT only).
- NAT46 (IPv4-to-IPv6)—We recommend using static NAT. Because the IPv6 address space is so much larger than the IPv4 address space, you can easily accommodate a static translation. If you do not want to allow returning traffic, you can make the static NAT rule unidirectional (manual NAT only). When translating to an IPv6 subnet (/96 or lower), the resulting mapped address is by default an IPv4-embedded IPv6 address, where the 32-bits of the IPv4 address is embedded after the IPv6 prefix. For example, if the IPv6 prefix is a /96 prefix, then the IPv4 address is appended in the last 32-bits of the address. For example, if you map 192.168.1.0/24 to 201b::0/96, then 192.168.1.4 will be mapped to 201b::0.192.168.1.4 (shown with mixed notation). If the prefix is smaller, such as /64, then the IPv4 address is appended after the prefix, and a suffix of 0s is appended after the IPv4 address. You can also optionally translate the addresses net-to-net, where the first IPv4 address maps to the first IPv6 address, the second to the second, and so on.
- NAT64 (IPv6-to-IPv4)—You may not have enough IPv4 addresses to accommodate the number of IPv6 addresses. We recommend using a dynamic PAT pool to provide a large number of IPv4 translations.

NAT Support for Inspected Protocols

Some application layer protocols that open secondary connections, or that embedded IP addresses in packets, are inspected to provide the following services:

- Pinhole creation—Some application protocols open secondary TCP or UDP connections either on standard or negotiated ports. Inspection opens pinholes for these secondary ports so that you do not need to create access control rules to allow them.
- NAT rewrite— Protocols such as FTP embed IP addresses and ports for the secondary connections in packet data as part of the protocol. If there is NAT translation involved for either of the endpoints, the inspection engines rewrite the packet data to reflect the NAT translation of the embedded addresses and ports. The secondary connections would not work without NAT rewrite.
- Protocol enforcement—Some inspections enforce some degree of conformance to the RFCs for the inspected protocol.

The following table lists the inspected protocols that apply NAT rewrite and their NAT limitations. Keep these limitations in mind when writing NAT rules that include these protocols. Inspected protocols not listed here do not apply NAT rewrite. These inspections include GTP, HTTP, IMAP, POP, SMTP, SSH, and SSL.



Note NAT rewrite is supported on the listed ports only. For some of these protocols, you can extend inspection to other ports using Network Analysis Policies, but NAT rewrite is not extended to those ports. This includes DCERPC, DNS, FTP, and Sun RPC inspection. If you use these protocols on non-standard ports, do not use NAT on the connections.

Table 82: NAT Supported Application Inspection

Application	Inspected Protocol, Port	NAT Limitations	Pinholes Created
DCERPC	TCP/135	No NAT64.	Yes
DNS over UDP	UDP/53	No NAT support is available for name resolution through WINS.	No
ESMTP	TCP/25	No NAT64.	No
FTP	TCP/21	(Clustering) No static PAT.	Yes
H.323 H.225 (Call signaling) H.323 RAS	TCP/1720 UDP/1718 For RAS, UDP/1718-1719	(Clustering) No static PAT. No extended PAT. No NAT64.	Yes
ICMP ICMP Error	ICMP (ICMP traffic directed to a device interface is never inspected.)	No limitations.	No
IP Options	RSVP	No NAT64.	No
NetBIOS Name Server over IP	UDP/137, 138 (Source ports)	No extended PAT. No NAT64.	No
RSH	TCP/514	No PAT. No NAT64. (Clustering) No static PAT.	Yes
RTSP	TCP/554 (No handling for HTTP cloaking.)	No extended PAT. No NAT64. (Clustering) No static PAT.	Yes

Application	Inspected Protocol, Port	NAT Limitations	Pinholes Created
SIP	TCP/5060 UDP/5060	No extended PAT. No NAT64 or NAT46. (Clustering) No static PAT.	Yes
Skinny (SCCP)	TCP/2000	No extended PAT. No NAT64, NAT46, or NAT66. (Clustering) No static PAT.	Yes
SQL*Net (versions 1, 2)	TCP/1521	No extended PAT. No NAT64. (Clustering) No static PAT.	Yes
Sun RPC	TCP/111 UDP/111	No extended PAT. No NAT64.	Yes
TFTP	UDP/69	No NAT64. (Clustering) No static PAT. Payload IP addresses are not translated.	Yes
XDMCP	UDP/177	No extended PAT. No NAT64. (Clustering) No static PAT.	Yes

Additional Guidelines for NAT

- For interfaces that are members of a bridge group, you write NAT rules for the member interfaces. You cannot write NAT rules for the Bridge Virtual Interface (BVI) itself.
- (Auto NAT only.) You can only define a single NAT rule for a given object; if you want to configure multiple NAT rules for an object, you need to create multiple objects with different names that specify the same IP address.
- If a VPN is defined on an interface, inbound ESP traffic on the interface is not subject to the NAT rules. The system allows the ESP traffic for established VPN tunnels only, dropping traffic not associated with an existing tunnel. This restriction applies to ESP and UDP ports 500 and 4500.
- If you define a site-to-site VPN on a device that is behind a device that is applying dynamic PAT, so that UDP ports 500 and 4500 are not the ones actually used, you must initiate the connection from the device that is behind the PAT device. The responder cannot initiate the security association (SA) because it does not know the correct port numbers.
- If you change the NAT configuration, and you do not want to wait for existing translations to time out before the new NAT configuration is used, you can clear the translation table using the **clear xlate** command in the device CLI. However, clearing the translation table disconnects all current connections that use translations.

If you create a new NAT rule that should apply to an existing connection (such as a VPN tunnel), you need to use **clear conn** to end the connection. Then, the attempt to re-establish the connection should hit the NAT rule and the connection should be NAT'ed correctly.

Note If you remove a dynamic NAT or PAT rule, and then add a new rule with mapped addresses that overlap the addresses in the removed rule, then the new rule will not be used until all connections associated with the removed rule time out or are cleared using the **clear xlate** or **clear conn** commands. This safeguard ensures that the same address is not assigned to multiple hosts.

- You cannot use an object group with both IPv4 and IPv6 addresses; the object group must include only one type of address.
- (Manual NAT only.) When using **any** as the source address in a NAT rule, the definition of "any" traffic (IPv4 vs. IPv6) depends on the rule. Before the FTD device performs NAT on a packet, the packet must be IPv6-to-IPv6 or IPv4-to-IPv4; with this prerequisite, the FTD device can determine the value of **any** in a NAT rule. For example, if you configure a rule from "any" to an IPv6 server, and that server was mapped from an IPv4 address, then **any** means "any IPv6 traffic." If you configure a rule from "any" to "any," and you map the source to the interface IPv4 address, then **any** means "any IPv4 traffic" because the mapped interface address implies that the destination is also IPv4.
- You can use the same mapped object or group in multiple NAT rules.
- The mapped IP address pool cannot include:
  - The mapped interface IP address. If you specify "any" interface for the rule, then all interface IP addresses are disallowed. For interface PAT (routed mode only), specify the interface name instead of the interface address.
  - The failover interface IP address.
  - (Transparent mode.) The management IP address.
  - (Dynamic NAT.) The standby interface IP address when VPN is enabled.
- Avoid using overlapping addresses in static and dynamic NAT policies. For example, with overlapping addresses, a PPTP connection can fail to get established if the secondary connection for PPTP hits the static instead of dynamic xlate.
- You cannot use overlapping addresses in the source address of a NAT rule and a remote access VPN address pool.
- If you specify a destination interface in a rule, then that interface is used as the egress interface rather than looking up the route in the routing table. However, for identity NAT, you have the option to use a route lookup instead.
- If you use PAT on Sun RPC traffic, which is used to connect to NFS servers, be aware that the NFS server might reject connections if the PAT'ed port is above 1024. The default configuration of NFS servers is to reject connections from ports higher than 1024. The error is typically "Permission Denied." Mapping ports above 1024 might happen if you use the "flat range" option to use the higher port numbers if a port in the lower range is not available, especially if you do not select the option to include the lower range in the flat range. You can avoid this problem by changing the NFS server configuration to allow all port numbers.

- NAT applies to through traffic only. Traffic generated by the system is not subject to NAT.
- Do not name a network object or group pat-pool, using any combination of upper- or lower-case letters.
- The unidirectional option is mainly useful for testing purposes and might not work with all protocols. For example, SIP requires protocol inspection to translate SIP headers using NAT, but this will not occur if you make the translation unidirectional.
- You cannot use NAT on the internal payload of Protocol Independent Multicast (PIM) registers.
- (Manual NAT) When writing NAT rules for a dual ISP interface setup (primary and backup interfaces using service level agreements in the routing configuration), do not specify destination criteria in the rule. Ensure the rule for the primary interface comes before the rule for the backup interface. This allows the device to choose the correct NAT destination interface based on the current routing state when the primary ISP is unavailable. If you specify destination objects, the NAT rule will always select the primary interface for the otherwise duplicate rules.
- If you get the ASP drop reason nat-no-xlate-to-pat-pool for traffic that should not match the NAT rules defined for the interface, configure identity NAT rules for the affected traffic so the traffic can pass untranslated.
- If you configure NAT for GRE tunnel endpoints, you must disable keepalives on the endpoints or the tunnel cannot be established. The endpoints send keepalives to the original addresses.

Configure NAT for Threat Defense

Network address translation can be very complex. We recommend that you keep your rules as simple as possible to avoid translation problems and difficult troubleshooting situations. Careful planning before you implement NAT is critical. The following procedure provides the basic approach.

The NAT policy is a shared policy. You assign the policy to devices that should have similar NAT rules.

Whether a given rule in the policy applies to an assigned device is determined by the interface objects (security zones or interface groups) used in the rule. If the interface objects include one or more interface for the device, the rule is deployed to the device. Thus, you can configure rules that apply to subsets of devices within a single shared policy by carefully designing your interface objects. Rules that apply to “any” interface object are deployed to all devices.

You can configure multiple NAT policies if groups of your devices require significantly different rules.

Procedure

Step 1 Select **Devices** > **NAT**.

- Click **New Policy** > **Threat Defense NAT** to create a new policy. Give the policy a name, optionally assign devices to it, and click **Save**.

  You can change device assignments later by editing the policy and clicking **Policy Assignments**.

- Click **Edit** () to edit an existing Threat Defense NAT policy. Note that the page also shows Firepower NAT policies, which are not used by Firepower Threat Defense devices.

Step 2 Decide what kinds of rules you need.

You can create dynamic NAT, dynamic PAT, static NAT, and identity NAT rules. For an overview, see NAT Types, on page 1006.

Step 3 Decide which rules should be implemented as manual or auto NAT.

For a comparison of these two implementation options, see Auto NAT and Manual NAT, on page 1009.

Step 4 Decide which rules should be custom per device.

Because you can assign a NAT policy to multiple devices, you can configure a single rule on many devices. However, you might have rules that should be interpreted differently by each device, or some rules that should apply to a subset of devices only.

Use interface objects to control on which devices a rule is configured. Then, use object overrides on network objects to customize the addresses used per device.

For detailed information, see Customizing NAT Rules for Multiple Devices, on page 1021.

Step 5 Create the rules as explained in the following sections.

- Dynamic NAT, on page 1023
- Dynamic PAT, on page 1028
- Static NAT, on page 1035
- Identity NAT, on page 1043

Step 6 Manage the NAT policy and rules.

You can do the following to manage the policy and its rules.

- To edit the policy name or description, click in those fields, type in your changes, and click outside the fields.
- To view only those rules that apply to a specific device, click **Filter by Device** and select the desired device. A rule applies to a device if it uses an interface object that includes an interface on the device.
- To change the devices to which the policy is assigned, click the **Policy Assignments** link and modify the selected devices list as desired.
- To change whether a rule is enabled or disabled, right click the rule and select the desired option from the **State** command. You can temporarily disable a rule without deleting it using these controls.
- To edit a rule, click **Edit** () for the rule.
- To delete a rule, click **Delete** () for the rule.

Step 7 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Customizing NAT Rules for Multiple Devices

Because the NAT policy is shared, you can assign a given policy to more than one device. However, you can configure at most one auto NAT rule for a given object. Thus, if you want to configure different translations for an object based on the specific device doing the translation, you need to carefully configure the interface objects (security zones or interface groups) and define network object overrides for the translated address.

The interface objects determine on which devices a rule gets configured. The network object overrides determine what IP addresses are used by a given device for that object.

Consider the following scenario:

- FTD-A and FTD-B have inside networks 192.168.1.0/24 attached to the interface named “inside.”
- On FTD-A, you want to translate all 192.168.1.0/24 addresses to a NAT pool in the 10.100.10.10 - 10.100.10.200 range when going to the “outside” interface.
- On FTD-B, you want to translate all 192.168.1.0/24 addresses to a NAT pool in the 10.200.10.10 - 10.200.10.200 range when going to the “outside” interface.

To accomplish the above, you would do the following. Although this example rule is for dynamic auto NAT, you can generalize the technique for any type of NAT rule.

Procedure

Step 1 Create the security zones for the inside and outside interfaces.

a) Choose **Objects** > **Object Management**.
b) Select **Interface Objects** from the table of contents and click **Add** > **Security Zone**. (You can use interface groups instead of zones.)
c) Configure the inside zone properties.
   - **Name**—Enter a name, for example, **inside-zone**.
   - **Type**—Select **Routed** for routed-mode devices, **Switched** for transparent mode.
   - **Selected Interfaces**—Add the FTD-A/inside and FTD-B/inside interfaces to the selected list.
d) Click **Save**.
e) Click **Add** > **Security Zone** and define the outside zone properties.
   - **Name**—Enter a name, for example, **outside-zone**.
   - **Interface Type**—Select **Routed** for routed-mode devices, **Switched** for transparent mode.
   - **Selected Interfaces**—Add the FTD-A/outside and FTD-B/outside interfaces to the selected list.
f) Click **Save**.

Step 2 Create the network object for the original inside network on the Object Management page.

a) Select **Network** from the table of contents and click **Add Network** > **Add Object**.
b) Configure the inside network properties.
   - **Name**—Enter a name, for example, **inside-network**.
   - **Network**—Enter the network address, for example, **192.168.1.0/24**.

c) Click **Save**.

Step 3 Create the network object for the translated NAT pool and define overrides.

a) Click **Add Network** > **Add Object**.

b) Configure the NAT pool properties for FTD-A.

- **Name**—Enter a name, for example, **NAT-pool**.
- **Network**—Enter the range of addresses to include in the pool for FTD-A, for example, **10.100.10.10-10.100.10.200**.

c) Select **Allow Overrides**.

d) Click the **Overrides** heading to open the list of object overrides.

e) Click **Add** to open the Add Object Override dialog box.

f) Select FTD-B and **Add** it to the Selected Devices list.

g) Click **Override** and change **Network** to **10.200.10.10-10.200.10.200**

h) Click **Add** to add the override to the device.

By defining an override for FTD-B, whenever the system configures this object on FTD-B, it will use the override value instead of the value defined in the original object.

i) Click **Save**.

Step 4 Configure the NAT rule.

a) Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Dynamic.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside-zone.
- **Destination Interface Objects** = outside-zone.

Note The interface objects control on which devices the rule is configured. Because in this example the zones contain interfaces for FTD-A and FTD-B only, even if the NAT policy were assigned to additional devices, the rule would be deployed to those 2 devices only.

e) On **Translation**, configure the following:

- **Original Source** = inside-network object.
- **Translated Source** > **Address**= NAT-pool object.

f) Click **Save**.

You now have a single rule that will be interpreted differently for FTD-A and FTD-B, providing unique translations for the inside networks protected by each firewall.

Dynamic NAT

The following topics explain dynamic NAT and how to configure it.

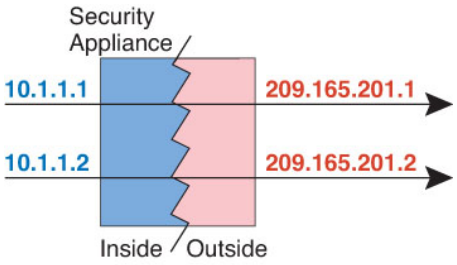
About Dynamic NAT

Dynamic NAT translates a group of real addresses to a pool of mapped addresses that are routable on the destination network. The mapped pool typically includes fewer addresses than the real group. When a host you want to translate accesses the destination network, NAT assigns the host an IP address from the mapped pool. The translation is created only when the real host initiates the connection. The translation is in place only for the duration of the connection, and a given user does not keep the same IP address after the translation times out. Users on the destination network, therefore, cannot initiate a reliable connection to a host that uses dynamic NAT, even if the connection is allowed by an access rule.

Note For the duration of the translation, a remote host can initiate a connection to the translated host if an access rule allows it. Because the address is unpredictable, a connection to the host is unlikely. Nevertheless, in this case you can rely on the security of the access rule.

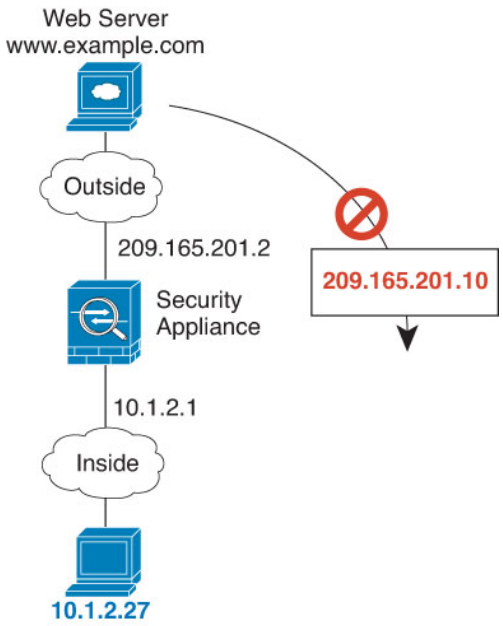
The following figure shows a typical dynamic NAT scenario. Only real hosts can create a NAT session, and responding traffic is allowed back.

Figure 35: Dynamic NAT



The following figure shows a remote host attempting to initiate a connection to a mapped address. This address is not currently in the translation table; therefore, the packet is dropped.

Figure 36: Remote Host Attempts to Initiate a Connection to a Mapped Address



Dynamic NAT Disadvantages and Advantages

Dynamic NAT has these disadvantages:

- If the mapped pool has fewer addresses than the real group, you could run out of addresses if the amount of traffic is more than expected.

  Use PAT or a PAT fall-back method if this event occurs often because PAT provides over 64,000 translations using ports of a single address.

- You have to use a large number of routable addresses in the mapped pool, and routable addresses may not be available in large quantities.

The advantage of dynamic NAT is that some protocols cannot use PAT. PAT does not work with the following:

- IP protocols that do not have a port to overload, such as GRE version 0.
- Some multimedia applications that have a data stream on one port, the control path on another port, and are not open standard.

Configure Dynamic Auto NAT

Use dynamic auto NAT rules to translate addresses to different IP addresses that are routable on the destination network.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Alternatively, you can create the objects while defining the NAT rule. The objects must meet the following requirements:

- **Original Source**—This must be a network object (not a group), and it can be a host, range, or subnet.

- **Translated Source**—This can be a network object or group, but it cannot include a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. If a group contains both ranges and host IP addresses, then the ranges are used for dynamic NAT, and then the host IP addresses are used as a PAT fallback.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** () to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Auto NAT Rule**.
- **Type**—Select **Dynamic**.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 On **Translation**, configure the following options:

- **Original Source**—The network object that contains the addresses you are translating.
- **Translated Source**—The network object or group that contains the mapped addresses.

Step 6 (Optional.) On **Advanced**, select the desired options:

- **Translate DNS replies that match this rule**—Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see Rewriting DNS Queries and Responses Using NAT, on page 1089.
- **Fallthrough to Interface PAT (Destination Interface)**—Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. To use the IPv6 address of the interface, also check the **IPv6** option.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.

Step 7 Click **Save** to add the rule.

Step 8 Click **Save** on the NAT page to save your changes.

Configure Dynamic Manual NAT

Use dynamic manual NAT rules when auto NAT does not meet your needs. For example, if you want to do different translations based on the destination. Dynamic NAT translates addresses to different IP addresses that are routable on the destination network.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Groups cannot contain both IPv4 and IPv6 addresses; they must contain one type only. Alternatively, you can create the objects while defining the NAT rule. The objects must also meet the following requirements:

- **Original Source**—This can be a network object or group, and it can contain a host, range, or subnet. If you want to translate all original source traffic, you can skip this step and specify **Any** in the rule.
- **Translated Source**—This can be a network object or group, but it cannot include a subnet. If a group contains both ranges and host IP addresses, then the ranges are used for dynamic NAT, and then the host IP addresses are used as a PAT fallback.

You can also create network objects for the **Original Destination** and **Translated Destination** if you are configuring a static translation for those addresses in the rule.

For dynamic NAT, you can also perform port translation on the destination. In the Object Manager, ensure that there are port objects you can use for the **Original Destination Port** and **Translated Destination Port**. If you specify the source port, it will be ignored.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** ( ) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

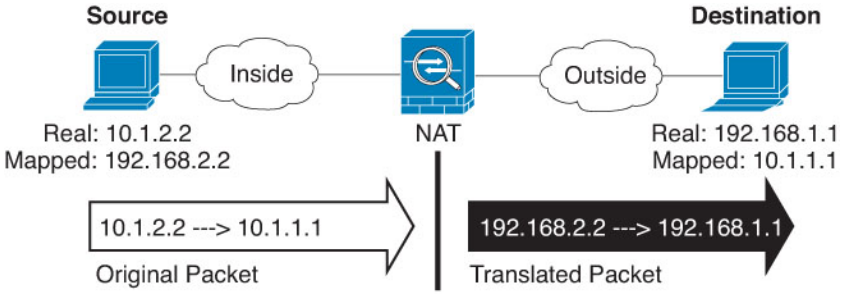
- **NAT Rule**—Select **Manual NAT Rule**.
- **Type**—Select **Dynamic**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.
- **Enable**—Whether you want the rule to be active. You can later activate or deactivate the rule using the right-click menu on the rules page.
- **Insert**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule number you specify.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 (On the **Translation** page.) Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear in the original packet.

See the following figure for an example of the original packet vs. the translated packet.



- **Original Source**—The network object or group that contains the addresses you are translating.
- **Original Destination**—(Optional.) The network object that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

  You can select **Source Interface IP** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network. You can translate between IPv4 and IPv6 if desired.

- **Translated Source**—The network object or group that contains the mapped addresses.
- **Translated Destination**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

Step 7 (Optional.) Identify the destination service ports for service translation: **Original Destination Port**, **Translated Destination Port**.

Dynamic NAT does not support port translation, so leave the **Original Source Port** and **Translated Source Port** fields empty. However, because the destination translation is always static, you can perform port translation for the destination port.

NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports.

Step 8 (Optional.) On **Advanced**, select the desired options:

- (For source translation only.) **Translate DNS replies that match this rule**—Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see Rewriting DNS Queries and Responses Using NAT, on page 1089.
- **Fallthrough to Interface PAT (Destination Interface)**—Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. To use the IPv6 address of the interface, also check the **IPv6** option.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.

Step 9 Click **Save** to add the rule.

Step 10 Click **Save** on the NAT page to save your changes.

Dynamic PAT

The following topics describe dynamic PAT.

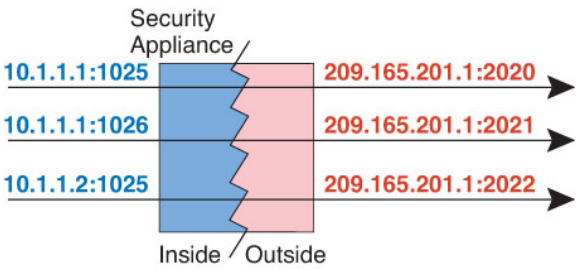
About Dynamic PAT

Dynamic PAT translates multiple real addresses to a single mapped IP address by translating the real address and source port to the mapped address and a unique port.

Each connection requires a separate translation session because the source port differs for each connection. For example, 10.1.1.1:1025 requires a separate translation from 10.1.1.1:1026.

The following figure shows a typical dynamic PAT scenario. Only real hosts can create a NAT session, and responding traffic is allowed back. The mapped address is the same for each translation, but the port is dynamically assigned.

Figure 37: Dynamic PAT



For the duration of the translation, a remote host on the destination network can initiate a connection to the translated host if an access rule allows it. Because the port address (both real and mapped) is unpredictable, a connection to the host is unlikely. Nevertheless, in this case you can rely on the security of the access rule.

After the connection expires, the port translation also expires.

Note We recommend that you use different PAT pools for each interface. If you use the same pool for multiple interfaces, especially if you use it for "any" interface, the pool can be quickly exhausted, with no ports available for new translations.

Dynamic PAT Disadvantages and Advantages

Dynamic PAT lets you use a single mapped address, thus conserving routable addresses. You can even use the FTD device interface IP address as the PAT address.

You cannot use dynamic PAT for IPv6 (NAT66) when translating between interfaces in the same bridge group. This restriction does not apply when the interfaces are members of different bridge groups, or between a bridge group member and a standard routed interface.

Dynamic PAT does not work with some multimedia applications that have a data stream that is different from the control path. For more information, see NAT Support for Inspected Protocols, on page 1015.

Dynamic PAT might also create a large number of connections appearing to come from a single IP address, and servers might interpret the traffic as a DoS attack. You can configure a PAT pool of addresses and use a round-robin assignment of PAT addresses to mitigate this situation.

PAT Pool Object Guidelines

When creating network objects for a PAT pool, follow these guidelines.

For a PAT pool

- If available, the real source port number is used for the mapped port. However, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 0 to 511, 512 to 1023, and 1024 to 65535. Therefore, ports below 1024 have only a small PAT pool that can be used. If you have a lot of traffic that uses the lower port ranges, you can specify a flat range of ports to be used instead of the three unequal-sized tiers: either 1024 to 65535, or 1 to 65535.
- If you use the same PAT pool object in two separate rules, then be sure to specify the same options for each rule. For example, if one rule specifies extended PAT, then the other rule must also specify extended PAT.
- If a host has an existing connection, then subsequent connections from that host use the same PAT IP address. If no ports are available, this can prevent the connection. Use the round robin option to avoid this problem.
- For best performance, limit the number of IP addresses within a PAT pool to 10,000.

For extended PAT for a PAT pool

- Many application inspections do not support extended PAT.
- If you enable extended PAT for a dynamic PAT rule, then you cannot also use an address in the PAT pool as the PAT address in a separate static NAT with port translation rule. For example, if the PAT pool includes 10.1.1.1, then you cannot create a static NAT-with-port-translation rule using 10.1.1.1 as the PAT address.
- If you use a PAT pool and specify an interface for fallback, you cannot specify extended PAT.

- For VoIP deployments that use ICE or TURN, do not use extended PAT. ICE and TURN rely on the PAT binding to be the same for all destinations.
- Extended PAT increases memory usage on the device.

For round robin for a PAT pool

- If a host has an existing connection, then subsequent connections from that host will use the same PAT IP address if ports are available. However, this "stickiness" does not survive a failover. If the device fails over, then subsequent connections from a host might not use the initial IP address.
- IP address "stickiness" is also impacted if you mix PAT pool/round robin rules with interface PAT rules on the same interface. For any given interface, choose either a PAT pool or interface PAT; do not create competing PAT rules.
- Round robin, especially when combined with extended PAT, can consume a large amount of memory. Because NAT pools are created for every mapped protocol/IP address/port range, round robin results in a large number of concurrent NAT pools, which use memory. Extended PAT results in an even larger number of concurrent NAT pools.

Configure Dynamic Auto PAT

Use dynamic auto PAT rules to translate addresses to unique IP address/port combinations, rather than to multiple IP addresses only. You can translate to a single address (either the destination interface's address or another address), or use a PAT pool of addresses to provide a larger number of possible translations.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Alternatively, you can create the objects while defining the NAT rule. The objects must meet the following requirements:

- **Original Source**—This must be a network object (not a group), and it can be a host, range, or subnet.
- **Translated Source**—You have the following options to specify the PAT address:
  - **Destination Interface**—To use the destination interface address, you do not need a network object.
  - **Single PAT address**—Create a network object containing a single host.
  - **PAT pool**—Create a network object that includes a range, or create a network object group that contains hosts, ranges, or both. You cannot include subnets. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** ( ) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Auto NAT Rule**.
- **Type**—Select **Dynamic**.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 On **Translation**, configure the following options:

- **Original Source**—The network object that contains the addresses you are translating.
- **Translated Source**—One of the following:
  - (Interface PAT.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. Skip the step for configuring a PAT pool.
  - To use a single address other than the destination interface address, select the host network object you created for this purpose. Skip the step for configuring a PAT pool.
  - To use a PAT pool, leave **Translated Source** empty.

Step 6 If you are using a PAT pool, select the **PAT Pool** page and do the following:

a) Select **Enable PAT pool**.

b) Select the network object group that contains the addresses for the pool in the **PAT** > **Address** field.

You can alternatively select **Destination Interface IP**, which is another way to implement interface PAT.

c) (Optional) Select the following options as needed:

- **Use Round Robin Allocation**—To assign addresses/ports in a round-robin fashion. By default without round robin, all ports for a PAT address will be allocated before the next PAT address is used. The round-robin method assigns one address/port from each PAT address in the pool before returning to use the first address again, and then the second address, and so on.
- **Extended PAT Table**—To use extended PAT. Extended PAT uses 65535 ports per *service*, as opposed to per IP address, by including the destination address and port in the translation information. Normally, the destination port and address are not considered when creating PAT translations, so you are limited to 65535 ports per PAT address. For example, with extended PAT, you can create a translation of 10.1.1.1:1027 when going to 192.168.1.7:23 as well as a translation of 10.1.1.1:1027 when going to 192.168.1.7:80. You cannot use this option with interface PAT or interface PAT fallback.
- **Flat Port Range**, **Include Reserved Ports**—To use the 1024 to 65535 port range as a single flat range when allocating TCP/UDP ports. When choosing the mapped port number for a translation, PAT uses the real source port number if it is available. However, without this option, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 1 to 511, 512 to 1023, and 1024 to 65535. To avoid running out of ports at the low ranges,

configure this setting. To use the entire range of 1 to 65535, also check the **Include Reserved Ports** option.

Step 7 (Optional.) On **Advanced**, select the desired options:

- **Fallthrough to Interface PAT (Destination Interface)**—Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. To use the IPv6 address of the interface, also check the **IPv6** option. You cannot select this option if you already configured interface PAT as the translated address or PAT pool.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.

Step 8 Click **Save** to add the rule.

Step 9 Click **Save** on the NAT page to save your changes.

Configure Dynamic Manual PAT

Use dynamic manual PAT rules when auto PAT does not meet your needs. For example, if you want to do different translations based on the destination. Dynamic PAT translates addresses to unique IP address/port combinations, rather than to multiple IP addresses only. You can translate to a single address (either the destination interface's address or another address), or use a PAT pool of addresses to provide a larger number of possible translations.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Groups cannot contain both IPv4 and IPv6 addresses; they must contain one type only. Alternatively, you can create the objects while defining the NAT rule. The objects must also meet the following requirements:

- **Original Source**—This can be a network object or group, and it can contain a host, range, or subnet. If you want to translate all original source traffic, you can skip this step and specify **Any** in the rule.
- **Translated Source**—You have the following options to specify the PAT address:
  - **Destination Interface**—To use the destination interface address, you do not need a network object.
  - **Single PAT address**—Create a network object containing a single host.
  - **PAT pool**—Create a network object that includes a range, or create a network object group that contains hosts, ranges, or both. You cannot include subnets.

You can also create network objects for the **Original Destination** and **Translated Destination** if you are configuring a static translation for those addresses in the rule.

For dynamic NAT, you can also perform port translation on the destination. In the Object Manager, ensure that there are port objects you can use for the **Original Destination Port** and **Translated Destination Port**. If you specify the source port, it will be ignored.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** ( ) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

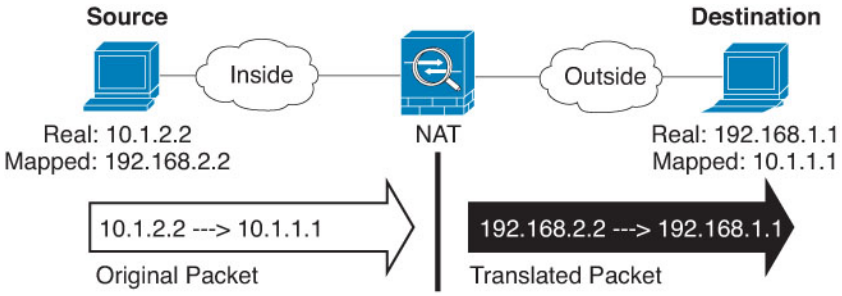
- **NAT Rule**—Select **Manual NAT Rule**.
- **Type**—Select **Dynamic**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.
- **Enable**—Whether you want the rule to be active. You can later activate or deactivate the rule using the right-click menu on the rules page.
- **Insert**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule number you specify.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 (On the **Translation** page.) Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear in the original packet.

See the following figure for an example of the original packet vs. the translated packet.



- **Original Source**—The network object or group that contains the addresses you are translating.
- **Original Destination**—(Optional.) The network object that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

  You can select **Source Interface IP** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network. You can translate between IPv4 and IPv6 if desired.

- **Translated Source**—One of the following:
  - (Interface PAT.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. Skip the step for configuring a PAT pool.
  - To use a single address other than the destination interface address, select the host network object you created for this purpose. Skip the step for configuring a PAT pool.
  - To use a PAT pool, leave **Translated Source** empty.
- **Translated Destination**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

Step 7 (Optional.) Identify the destination service ports for service translation: **Original Destination Port**, **Translated Destination Port**.

Dynamic NAT does not support port translation, so leave the **Original Source Port** and **Translated Source Port** fields empty. However, because the destination translation is always static, you can perform port translation for the destination port.

NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports.

Step 8 If you are using a PAT pool, select the **PAT Pool** page and do the following:

a) Select **Enable PAT pool**.

b) Select the network object group that contains the addresses for the pool in the **PAT** > **Address** field.

You can alternatively select **Destination Interface IP**, which is another way to implement interface PAT.

c) (Optional) Select the following options as needed:

- **Use Round Robin Allocation**—To assign addresses/ports in a round-robin fashion. By default without round robin, all ports for a PAT address will be allocated before the next PAT address is used. The round-robin method assigns one address/port from each PAT address in the pool before returning to use the first address again, and then the second address, and so on.
- **Extended PAT Table**—To use extended PAT. Extended PAT uses 65535 ports per *service*, as opposed to per IP address, by including the destination address and port in the translation information. Normally, the destination port and address are not considered when creating PAT translations, so you are limited to 65535 ports per PAT address. For example, with extended PAT, you can create a translation of 10.1.1.1:1027 when going to 192.168.1.7:23 as well as a translation of 10.1.1.1:1027 when going to 192.168.1.7:80. You cannot use this option with interface PAT or interface PAT fallback.
- **Flat Port Range**, **Include Reserved Ports**—To use the 1024 to 65535 port range as a single flat range when allocating TCP/UDP ports. When choosing the mapped port number for a translation, PAT uses the real source port number if it is available. However, without this option, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 1 to 511, 512 to 1023, and 1024 to 65535. To avoid running out of ports at the low ranges,

configure this setting. To use the entire range of 1 to 65535, also check the **Include Reserved Ports** option.

Step 9 (Optional.) On **Advanced**, select the desired options:

- **Fallthrough to Interface PAT (Destination Interface)**—Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. To use the IPv6 address of the interface, also check the **IPv6** option.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.

Step 10 Click **Save** to add the rule.

Step 11 Click **Save** on the NAT page to save your changes.

Static NAT

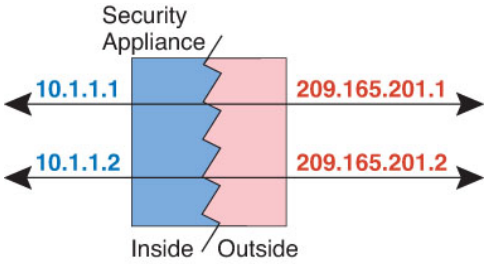
The following topics explain static NAT and how to implement it.

About Static NAT

Static NAT creates a fixed translation of a real address to a mapped address. Because the mapped address is the same for each consecutive connection, static NAT allows bidirectional connection initiation, both to and from the host (if an access rule exists that allows it). With dynamic NAT and PAT, on the other hand, each host uses a different address or port for each subsequent translation, so bidirectional initiation is not supported.

The following figure shows a typical static NAT scenario. The translation is always active so both real and remote hosts can initiate connections.

Figure 38: Static NAT



Note You can disable bidirectionality if desired.

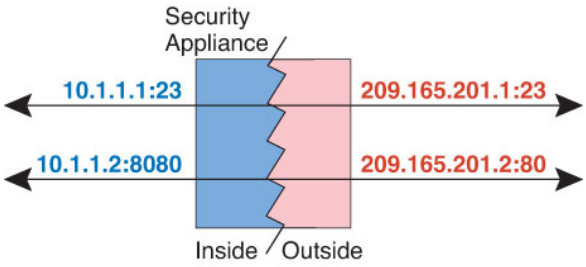
Static NAT with Port Translation

Static NAT with port translation lets you specify a real and mapped protocol and port.

When you specify the port with static NAT, you can choose to map the port and/or the IP address to the same value or to a different value.

The following figure shows a typical static NAT with port translation scenario showing both a port that is mapped to itself and a port that is mapped to a different value; the IP address is mapped to a different value in both cases. The translation is always active so both translated and remote hosts can initiate connections.

Figure 39: Typical Static NAT with Port Translation Scenario



Static NAT-with-port-translation rules limit access to the destination IP address for the specified port only. If you try to access the destination IP address on a different port not covered by a NAT rule, then the connection is blocked. In addition, for manual NAT, traffic that does not match the source IP address of the NAT rule will be dropped if it matches the destination IP address, regardless of the destination port. Therefore, you must add additional rules for all other traffic allowed to the destination IP address. For example, you can configure a static NAT rule for the IP address, without port specification, and place it after the port translation rule.

Note For applications that require application inspection for secondary channels (for example, FTP and VoIP), NAT automatically translates the secondary ports.

Following are some other uses of static NAT with port translation.

Static NAT with Identity Port Translation

You can simplify external access to internal resources. For example, if you have three separate servers that provide services on different ports (such as FTP, HTTP, and SMTP), you can give external users a single IP address to access those services. You can then configure static NAT with identity port translation to map the single external IP address to the correct IP addresses of the real servers based on the port they are trying to access. You do not need to change the port, because the servers are using the standard ones (21, 80, and 25 respectively).

Static NAT with Port Translation for Non-Standard Ports

You can also use static NAT with port translation to translate a well-known port to a non-standard port or vice versa. For example, if inside web servers use port 8080, you can allow outside users to connect to port 80, and then undo translation to the original port 8080. Similarly, to provide extra security, you can tell web users to connect to non-standard port 6785, and then undo translation to port 80.

Static Interface NAT with Port Translation

You can configure static NAT to map a real address to an interface address/port combination. For example, if you want to redirect Telnet access for the device's outside interface to an inside host, then you can map the inside host IP address/port 23 to the outside interface address/port 23.

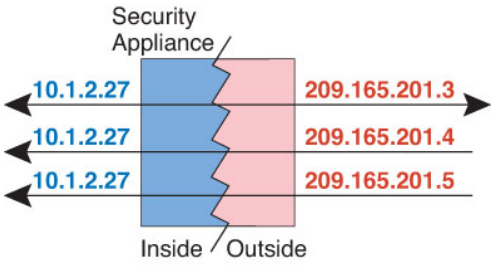
One-to-Many Static NAT

Typically, you configure static NAT with a one-to-one mapping. However, in some cases, you might want to configure a single real address to several mapped addresses (one-to-many). When you configure one-to-many

static NAT, when the real host initiates traffic, it always uses the first mapped address. However, for traffic initiated to the host, you can initiate traffic to any of the mapped addresses, and they will be untranslated to the single real address.

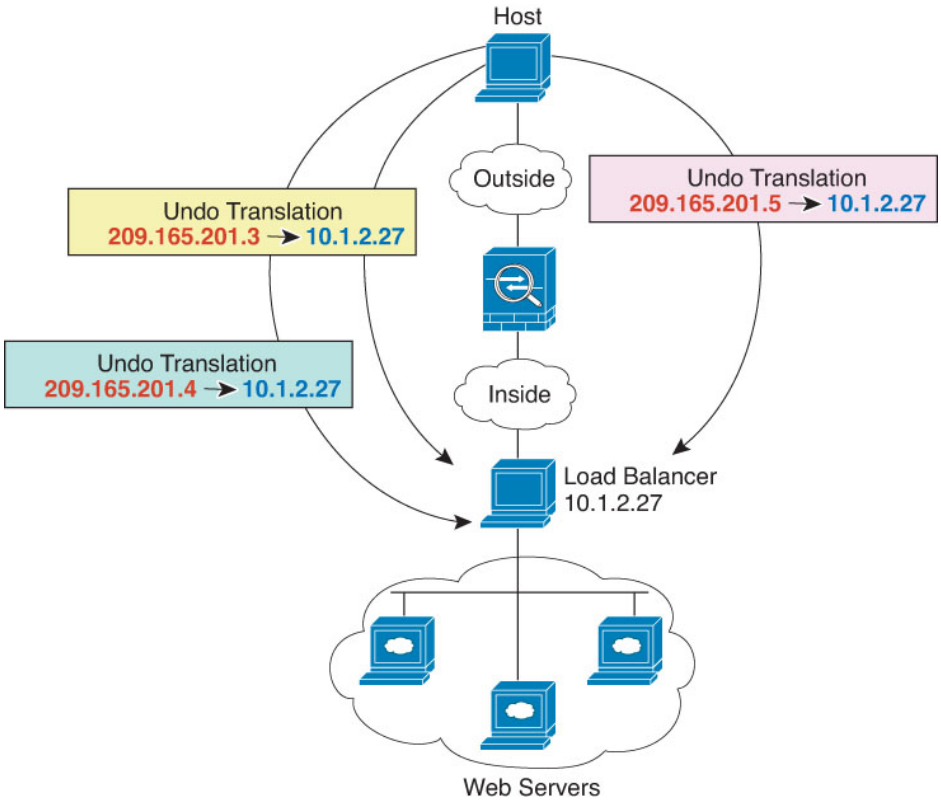
The following figure shows a typical one-to-many static NAT scenario. Because initiation by the real host always uses the first mapped address, the translation of real host IP/first mapped IP is technically the only bidirectional translation.

Figure 40: One-to-Many Static NAT



For example, you have a load balancer at 10.1.2.27. Depending on the URL requested, it redirects traffic to the correct web server.

Figure 41: One-to-Many Static NAT Example



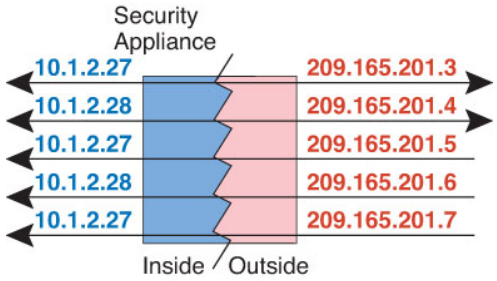
Other Mapping Scenarios (Not Recommended)

NAT has the flexibility to allow any kind of static mapping scenario: one-to-one, one-to-many, but also few-to-many, many-to-few, and many-to-one mappings. We recommend using only one-to-one or one-to-many mappings. These other mapping options might result in unintended consequences.

Functionally, few-to-many is the same as one-to-many; but because the configuration is more complicated and the actual mappings may not be obvious at a glance, we recommend creating a one-to-many configuration for each real address that requires it. For example, for a few-to-many scenario, the few real addresses are mapped to the many mapped addresses in order (A to 1, B to 2, C to 3). When all real addresses are mapped, the next mapped address is mapped to the first real address, and so on until all mapped addresses are mapped (A to 4, B to 5, C to 6). This results in multiple mapped addresses for each real address. Just like a one-to-many configuration, only the first mappings are bidirectional; subsequent mappings allow traffic to be initiated *to* the real host, but all traffic *from* the real host uses only the first mapped address for the source.

The following figure shows a typical few-to-many static NAT scenario.

Figure 42: Few-to-Many Static NAT

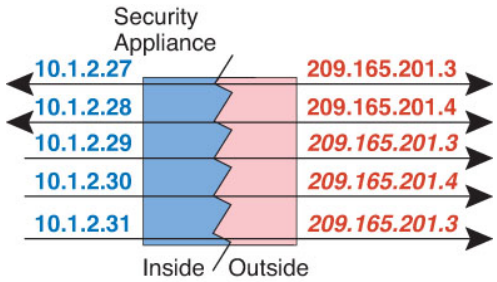


For a many-to-few or many-to-one configuration, where you have more real addresses than mapped addresses, you run out of mapped addresses before you run out of real addresses. Only the mappings between the lowest real IP addresses and the mapped pool result in bidirectional initiation. The remaining higher real addresses can initiate traffic, but traffic cannot be initiated to them (returning traffic for a connection is directed to the correct real address because of the unique 5-tuple (source IP, destination IP, source port, destination port, protocol) for the connection).

Note Many-to-few or many-to-one NAT is not PAT. If two real hosts use the same source port number and go to the same outside server and the same TCP destination port, and both hosts are translated to the same IP address, then both connections will be reset because of an address conflict (the 5-tuple is not unique).

The following figure shows a typical many-to-few static NAT scenario.

Figure 43: Many-to-Few Static NAT



Instead of using a static rule this way, we suggest that you create a one-to-one rule for the traffic that needs bidirectional initiation, and then create a dynamic rule for the rest of your addresses.

Configure Static Auto NAT

Use static auto NAT rules to translate addresses to different IP addresses that are routable on the destination network. You can also do port translation with the static NAT rule.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Alternatively, you can create the objects while defining the NAT rule. The objects must meet the following requirements:

- **Original Source**—This must be a network object (not a group), and it can be a host, range, or subnet.
- **Translated Source**—You have the following options to specify the translated address:
  - **Destination Interface**—To use the destination interface address, you do not need a network object. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
  - **Address**—Create a network object or group containing hosts, ranges, or subnets. A group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** ( ) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Auto NAT Rule**.
- **Type**—Select **Static**.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 On **Translation**, configure the following options:

- **Original Source**—The network object that contains the addresses you are translating.
- **Translated Source**—One of the following:
  - To use a set group of addresses, select **Address** and the network object or group that contains the mapped addresses. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.
  - (Static interface NAT with port translation.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
- (Optional.) **Original Port**, **Translated Port**—If you need to translate a TCP or UDP port, select the protocol in **Original Port**, and type the original and translated port numbers. For example, you can translate TCP/80 to 8080 if necessary.

Step 6 (Optional.) On **Advanced**, select the desired options:

- **Translate DNS replies that match this rule**—Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see Rewriting DNS Queries and Responses Using NAT, on page 1089. This option is not available if you are doing port translation.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.
- **Net to Net Mapping**—For NAT 46, select this option to translate the first IPv4 address to the first IPv6 address, the second to the second, and so on. Without this option, the IPv4-embedded method is used. For a one-to-one translation, you must use this option.
- **Do not proxy ARP on Destination Interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.

Step 7 Click **Save** to add the rule.

Step 8 Click **Save** on the NAT page to save your changes.

Configure Static Manual NAT

Use static manual NAT rules when auto NAT does not meet your needs. For example, if you want to do different translations based on the destination. Static NAT translates addresses to different IP addresses that are routable on the destination network. You can also do port translation with the static NAT rule.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Groups cannot contain both IPv4 and IPv6 addresses; they must contain one type only. Alternatively, you can create the objects while defining the NAT rule. The objects must also meet the following requirements:

- **Original Source**—This can be a network object or group, and it can contain a host, range, or subnet. If you want to translate all original source traffic, you can skip this step and specify **Any** in the rule.
- **Translated Source**—You have the following options to specify the translated address:
  - **Destination Interface**—To use the destination interface address, you do not need a network object. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
  - **Address**—Create a network object or group containing hosts, range, or subnets. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.

You can also create network objects for the **Original Destination** and **Translated Destination** if you are configuring a static translation for those addresses in the rule. If you want to configure destination static interface NAT with port translation only, you can skip adding an object for the destination mapped addresses and specify the interface in the rule.

You can also perform port translation on the source, destination, or both. In the Object Manager, ensure that there are port objects you can use for the original and translated ports.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** () to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Manual NAT Rule**.
- **Type**—Select **Static**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.
- **Enable**—Whether you want the rule to be active. You can later activate or deactivate the rule using the right-click menu on the rules page.
- **Insert**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule number you specify.

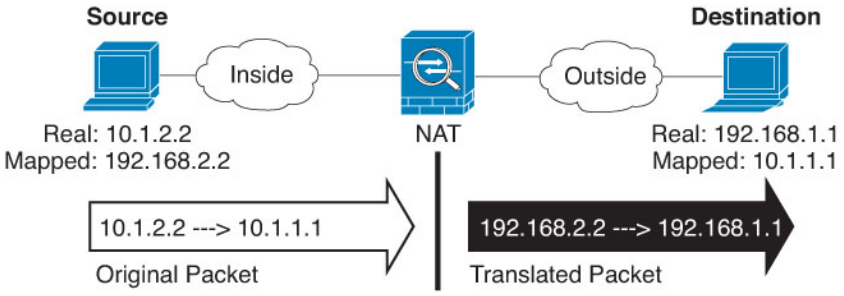
Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic

exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 (On the **Translation** page.) Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear in the original packet.

See the following figure for an example of the original packet vs. the translated packet.



- **Original Source**—The network object or group that contains the addresses you are translating.
- **Original Destination**—(Optional.) The network object that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

  You can select **Source Interface IP** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network. You can translate between IPv4 and IPv6 if desired.

- **Translated Source**—One of the following:
  - To use a set group of addresses, select **Address** and the network object or group that contains the mapped addresses. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.
  - (Static interface NAT with port translation.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
- **Translated Destination**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

Step 7 (Optional.) Identify the source or destination service ports for service translation.

If you are configuring static NAT with port translation, you can translate ports for the source, destination, or both. For example, you can translate between TCP/80 and TCP/8080.

NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports.

- **Original Source Port**, **Translated Source Port**—Defines a port translation for the source address.
- **Original Destination Port**, **Translated Destination Port**—Defines a port translation for the destination address.

Step 8 (Optional.) On **Advanced**, select the desired options:

- **Translate DNS replies that match this rule**—Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see Rewriting DNS Queries and Responses Using NAT, on page 1089. This option is not available if you are doing port translation.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.
- **Net to Net Mapping**—For NAT 46, select this option to translate the first IPv4 address to the first IPv6 address, the second to the second, and so on. Without this option, the IPv4-embedded method is used. For a one-to-one translation, you must use this option.
- **Do not proxy ARP on Destination Interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.
- **Unidirectional**—Select this option to prevent the destination addresses from initiating traffic to the source addresses. The unidirectional option is mainly useful for testing purposes and might not work with all protocols. For example, SIP requires protocol inspection to translate SIP headers using NAT, but this will not occur if you make the translation unidirectional.

Step 9 Click **Save** to add the rule.

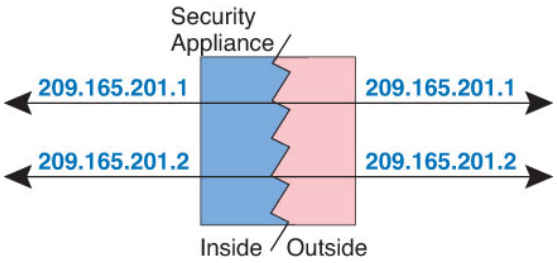
Step 10 Click **Save** on the NAT page to save your changes.

Identity NAT

You might have a NAT configuration in which you need to translate an IP address to itself. For example, if you create a broad rule that applies NAT to every network, but want to exclude one network from NAT, you can create a static NAT rule to translate an address to itself.

The following figure shows a typical identity NAT scenario.

Figure 44: Identity NAT



The following topics explain how to configure identity NAT.

Configure Identity Auto NAT

Use static identity auto NAT rules to prevent the translation of an address. That is, to translate the address to itself.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Alternatively, you can create the objects while defining the NAT rule. The objects must meet the following requirements:

- **Original Source**—This must be a network object (not a group), and it can be a host, range, or subnet.
- **Translated Source**—A network object or group with the exact same contents as the original source object. You can use the same object.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** ( ) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Auto NAT Rule**.
- **Type**—Select **Static**.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 On **Translation**, configure the following options:

- **Original Source**—The network object that contains the addresses you are translating.
- **Translated Source**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.

Do not configure the **Original Port** and **Translated Port** options for identity NAT.

Step 6 (Optional.) On **Advanced**, select the desired options:

- **Translate DNS replies that match this rule**—Do not configure this option for identity NAT.
- **IPv6**—Do not configure this option for identity NAT.
- **Net to Net Mapping**—Do not configure this option for identity NAT.
- **Do not proxy ARP on Destination Interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.
- **Perform Route Lookup for Destination Interface**— If you select source and destination interfaces when selecting the same object for original and translated source address, you can select this option to have the system determine the destination interface based on the routing table rather than using the destination interface configured in the NAT rule.

Step 7 Click **Save** to add the rule.

Step 8 Click **Save** on the NAT page to save your changes.

Configure Identity Manual NAT

Use static identity manual NAT rules when auto NAT does not meet your needs. For example, if you want to do different translations based on the destination. Use static identity NAT rules to prevent the translation of an address. That is, to translate the address to itself.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Groups cannot contain both IPv4 and IPv6 addresses; they must contain one type only. Alternatively, you can create the objects while defining the NAT rule. The objects must also meet the following requirements:

- **Original Source**—This can be a network object or group, and it can contain a host, range, or subnet. If you want to translate all original source traffic, you can skip this step and specify **Any** in the rule.
- **Translated Source**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.

You can also create network objects for the **Original Destination** and **Translated Destination** if you are configuring a static translation for those addresses in the rule. If you want to configure destination static interface NAT with port translation only, you can skip adding an object for the destination mapped addresses and specify the interface in the rule.

You can also perform port translation on the source, destination, or both. In the Object Manager, ensure that there are port objects you can use for the original and translated ports. You can use the same object for identity NAT.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** ( ) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

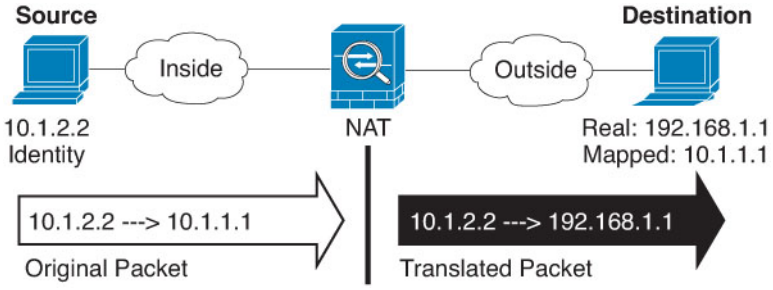
- **NAT Rule**—Select **Manual NAT Rule**.
- **Type**—Select **Static**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.
- **Enable**—Whether you want the rule to be active. You can later activate or deactivate the rule using the right-click menu on the rules page.
- **Insert**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule number you specify.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear in the original packet.

See the following figure for an example of the original packet vs. the translated packet where you perform identity NAT on the inside host but translate the outside host.



- **Original Source**—The network object or group that contains the addresses you are translating.
- **Original Destination**—(Optional.) The network object that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify

the destination address, you can configure a static translation for that address or just use identity NAT for it.

You can select **Interface Object** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network. You can translate between IPv4 and IPv6 if desired.

- **Translated Source**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.
- **Translated Destination**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

Step 7 (Optional.) Identify the source or destination service ports for service translation.

If you are configuring static NAT with port translation, you can translate ports for the source, destination, or both. For example, you can translate between TCP/80 and TCP/8080.

NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports.

- **Original Source Port**, **Translated Source Port**—Defines a port translation for the source address.
- **Original Destination Port**, **Translated Destination Port**—Defines a port translation for the destination address.

Step 8 (Optional.) On **Advanced**, select the desired options:

- **Translate DNS replies that match this rule**—Do not configure this option for identity NAT.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.
- **Do not proxy ARP on Destination Interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.
- **Perform Route Lookup for Destination Interface**— If you select source and destination interfaces when selecting the same object for original and translated source address, you can select this option to have the system determine the destination interface based on the routing table rather than using the destination interface configured in the NAT rule.
- **Unidirectional**—Select this option to prevent the destination addresses from initiating traffic to the source addresses. The unidirectional option is mainly useful for testing purposes and might not work with all protocols. For example, SIP requires protocol inspection to translate SIP headers using NAT, but this will not occur if you make the translation unidirectional.

Step 9 Click **Save** to add the rule.

Step 10 Click **Save** on the NAT page to save your changes.

NAT Rule Properties for Firepower Threat Defense

Use Network Address Translation (NAT) rules to translate IP addresses to other IP addresses. You would typically use NAT rules to convert private addresses to publically routable addresses. The translation can be from one address to another, or you can use Port Address Translation (PAT) to translate many addresses to one or a few addresses, using port numbers to distinguish among the source addresses.

NAT rules include the following basic properties. The properties are the same for auto NAT and manual NAT rules except where indicated.

NAT Type

Whether you want to configure a **Manual NAT Rule** or an **Auto NAT Rule**. Auto NAT translates the source address only, and you cannot make different translations based on the destination address. Because auto NAT is more simple to configure, use it unless you need the added features of manual NAT. For more information on the differences, see Auto NAT and Manual NAT, on page 1009.

Type

Whether the translation rule is **Dynamic** or **Static**. Dynamic translation automatically chooses the mapped address from a pool of addresses, or an address/port combination when implementing PAT. Use static translation if you want to precisely define the mapped address/port.

Enable (Manual NAT only.)

Whether you want the rule to be active. You can later activate or deactivate the rule using the right-click menu on the rules page. You cannot disable auto NAT rules.

Insert (Manual NAT only.)

Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule number you specify.

Description (Optional. Manual NAT only.)

A description of the purpose of the rule.

The following topics describe the tabs for the NAT rules properties.

Interface Objects NAT Properties

Interface objects (security zones or interface groups) define the interfaces to which a NAT rule applies. In routed mode, you can use the default "any" for both source and destination to apply to all interfaces of all assigned devices. However, you typically want to select specific source and destination interfaces.



Note The concept of “any” interface does not apply to bridge group member interfaces. When you specify “any” interface, all bridge group member interfaces are excluded. Thus, to apply NAT to bridge group members, you must specify the member interface. You cannot configure NAT for the Bridge Virtual Interface (BVI) itself, you can configure NAT for member interfaces only.

If you select interface objects, a NAT rule will be configured on an assigned device only if the device has interfaces included in all selected objects. For example, if you select both source and destination security zones, both zones must contain one or more interface for a given device.

Source Interface Objects, Destination Interface Objects

(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Translation Properties for Auto NAT

Use the options on **Translation** to define the source addresses and the mapped translated addresses. The following properties apply to auto NAT only.

Original Source (Always required.)

The network object that contains the addresses you are translating. This must be a network object (not a group), and it can be a host, range, or subnet.

You cannot create auto NAT rules for the system-defined any-ipv4 or any-ipv6 objects.

Translated Source (Usually required.)

The mapped addresses, the ones to which you are translating. What you select here depends on the type of translation rule you are defining.

- **Dynamic NAT**—The network object or group that contains the mapped addresses. This can be a network object or group, but it cannot include a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. If a group contains both ranges and host IP addresses, then the ranges are used for dynamic NAT, and then the host IP addresses are used as a PAT fallback.
- **Dynamic PAT**—One of the following:
  - (Interface PAT.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. Do not configure a PAT pool.
  - To use a single address other than the destination interface address, select the host network object you created for this purpose. Do not configure a PAT pool.
  - To use a PAT pool, leave **Translated Source** empty. Select the PAT pool object on **PAT Pool**.
- **Static NAT**—One of the following:
  - To use a set group of addresses, select **Address** and the network object or group that contains the mapped addresses. The object or group can contain hosts, ranges, or subnets. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.
  - (Static interface NAT with port translation.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on the **Advanced** tab. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
- **Identity NAT**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.

Original Port, Translated Port (Static NAT only.)

If you need to translate a TCP or UDP port, select the protocol in **Original Port**, and type the original and translated port numbers. For example, you can translate TCP/80 to 8080 if necessary. Do not configure these options for identity NAT.

Translation Properties for Manual NAT

Use the options on **Translation** to define the source addresses and the mapped translated addresses. The following properties apply to manual NAT only. All are optional except as indicated.

Original Source (Always required.)

The network object or group that contains the addresses you are translating. This can be a network object or group, and it can contain a host, range, or subnet. If you want to translate all original source traffic, you can specify **Any** in the rule.

Translated Source (Usually required.)

The mapped addresses, the ones to which you are translating. What you select here depends on the type of translation rule you are defining.

- **Dynamic NAT**—The network object or group that contains the mapped addresses. This can be a network object or group, but it cannot include a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. If a group contains both ranges and host IP addresses, then the ranges are used for dynamic NAT, and then the host IP addresses are used as a PAT fallback.
- **Dynamic PAT**—One of the following:
  - (Interface PAT.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. Do not configure a PAT pool.
  - To use a single address other than the destination interface address, select the host network object you created for this purpose. Do not configure a PAT pool.
  - To use a PAT pool, leave **Translated Source** empty. Select the PAT pool object on **PAT Pool**.
- **Static NAT**—One of the following:
  - To use a set group of addresses, select **Address** and the network object or group that contains the mapped addresses. The object or group can contain hosts, ranges, or subnets. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.
  - (Static interface NAT with port translation.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on the **Advanced** tab. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
- **Identity NAT**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.

Original Destination

The network object that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

You can select **Source Interface IP** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Translated Destination

The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

Original Source Port, Translated Source Port, Original Destination Port, Translated Destination Port

The port objects that define the source and destination services for the original and translated packets. You can translate the ports, or select the same object to make the rule sensitive to the service without translating the ports. Keep the following rules in mind when configuring services:

- (Dynamic NAT or PAT.) You cannot do translation on the **Original Source Port** and **Translated Source Port**. You can do translation on the destination port only.
- NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports.

PAT Pool NAT Properties

When you configure dynamic NAT, you can define a pool of addresses to use for Port Address Translation using the properties on the **PAT Pool** tab.

Enable PAT Pool

Select this option to configure a pool of addresses for PAT.

PAT

The addresses to use for the PAT pool, one of the following:

- **Address**—The object that defines the PAT pool addresses, either a network object that includes a range, or a network object group that contains hosts, ranges, or both. You cannot include subnets. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.
- **Destination Interface IP**—Indicates that you want to use the destination interface as the PAT address. For this option, you must select a specific **Destination Interface Object**; you cannot use **Any** as the destination interface. This is another way to implement interface PAT.

Round Robin

To assign addresses/ports in a round-robin fashion. By default without round robin, all ports for a PAT address will be allocated before the next PAT address is used. The round-robin method assigns one address/port from each PAT address in the pool before returning to use the first address again, and then the second address, and so on.

Extended PAT Table

To use extended PAT. Extended PAT uses 65535 ports per *service*, as opposed to per IP address, by including the destination address and port in the translation information. Normally, the destination port and address are not considered when creating PAT translations, so you are limited to 65535 ports per PAT address. For example, with extended PAT, you can create a translation of 10.1.1.1:1027 when going to 192.168.1.7:23 as well as a translation of 10.1.1.1:1027 when going to 192.168.1.7:80. You cannot use this option with interface PAT or interface PAT fallback.

Flat Port Range; Include Reserved Ports

To use the 1024 to 65535 port range as a single flat range when allocating TCP/UDP ports. When choosing the mapped port number for a translation, PAT uses the real source port number if it is available. However, without this option, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 1 to 511, 512 to 1023, and 1024 to 65535. To avoid running out of ports at the low ranges, configure this setting. To use the entire range of 1 to 65535, also check the **Include Reserved Ports** option.

Advanced NAT Properties

When you configure NAT, you can configure properties that provide specialized services in the **Advanced** options. All of these properties are optional: configure them only if you need the service.

Translate DNS replies that match this rule

Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see Rewriting DNS Queries and Responses Using NAT, on page 1089. This option is not available if you are doing port translation in a static NAT rule.

Fallthrough to Interface PAT (Destination Interface) (Dynamic NAT only.)

Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. To use the IPv6 address of the interface, also check the **IPv6** option. You cannot select this option if you already configured interface PAT as the translated address. You also cannot select the option if you configure a PAT pool.

IPv6

Whether to use the IPv6 address of the destination interface for interface PAT.

Net to Net Mapping (Static NAT only.)

For NAT 46, select this option to translate the first IPv4 address to the first IPv6 address, the second to the second, and so on. Without this option, the IPv4-embedded method is used. For a one-to-one translation, you must use this option.

Do not proxy ARP on Destination Interface (Static NAT only.)

Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if

desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.

Perform Route Lookup for Destination Interface (Static Identity NAT only. Routed mode only.)

If you select source and destination interfaces when selecting the same object for original and translated source address, you can select this option to have the system determine the destination interface based on the routing table rather than using the destination interface configured in the NAT rule.

Unidirectional (Manual NAT only, static NAT only.)

Select this option to prevent the destination addresses from initiating traffic to the source addresses. The unidirectional option is mainly useful for testing purposes and might not work with all protocols. For example, SIP requires protocol inspection to translate SIP headers using NAT, but this will not occur if you make the translation unidirectional.

Translating IPv6 Networks

In cases where you need to pass traffic between IPv6-only and IPv4-only networks, you need to use NAT to convert between the address types. Even with two IPv6 networks, you might want to hide internal addresses from the outside network.

You can use the following translation types with IPv6 networks:

- NAT64, NAT46—Translates IPv6 packets into IPv4 and vice versa. You need to define two policies, one for the IPv6 to IPv4 translation, and one for the IPv4 to IPv6 translation. Although you can accomplish this with a single manual NAT rule, if the DNS server is on the external network, you probably need to rewrite the DNS response. Because you cannot enable DNS rewrite on a manual NAT rule when you specify a destination, creating two auto NAT rules is the better solution.



Note NAT46 supports static mappings only.

- NAT66—Translates IPv6 packets to a different IPv6 address. We recommend using static NAT. Although you can use dynamic NAT or PAT, IPv6 addresses are in such large supply, you do not have to use dynamic NAT.



Note NAT64 and NAT 46 are possible on standard routed interfaces only. NAT66 is possible on both routed and bridge group member interfaces.

NAT64/46: Translating IPv6 Addresses to IPv4

When traffic goes from an IPv6 network to an IPv4-only network, you need to convert the IPv6 address to IPv4, and return traffic from IPv4 to IPv6. You need to define two address pools, an IPv4 address pool to bind IPv6 addresses in the IPv4 network, and an IPv6 address pool to bind IPv4 addresses in the IPv6 network.

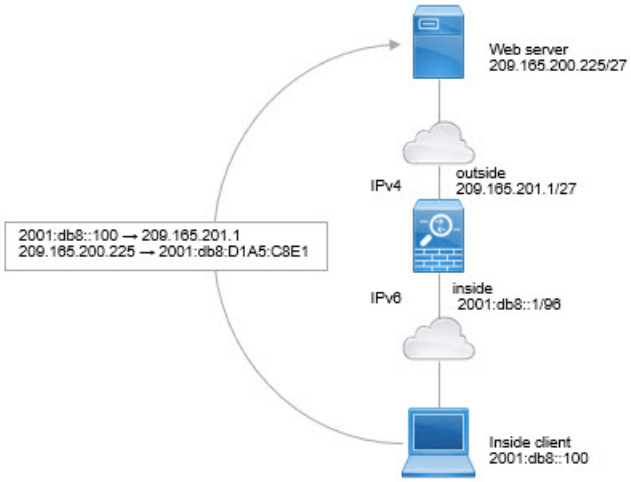
- The IPv4 address pool for the NAT64 rule is normally small and typically might not have enough addresses to map one-to-one with the IPv6 client addresses. Dynamic PAT might more easily meet the possible large number of IPv6 client addresses compared to dynamic or static NAT.

- The IPv6 address pool for the NAT46 rule can be equal to or larger than the number of IPv4 addresses to be mapped. This allows each IPv4 address to be mapped to a different IPv6 address. NAT46 supports static mappings only, so you cannot use dynamic PAT.

You need to define two policies, one for the source IPv6 network, and one for the destination IPv4 network. Although you can accomplish this with a single manual NAT rule, if the DNS server is on the external network, you probably need to rewrite the DNS response. Because you cannot enable DNS rewrite on a manual NAT rule when you specify a destination, creating two auto NAT rules is the better solution.

NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet

Following is a straight-forward example where you have an inside IPv6-only network, and you want to convert to IPv4 for traffic sent to the Internet. This example assumes you do not need DNS translation, so you can perform both the NAT64 and NAT46 translations in a single manual NAT rule.



In this example, you translate the inside IPv6 network to IPv4 using dynamic interface PAT with the IP address of the outside interface. Outside IPv4 traffic is statically translated to addresses on the 2001:db8::/96 network, allowing transmission on the inside network.

Procedure

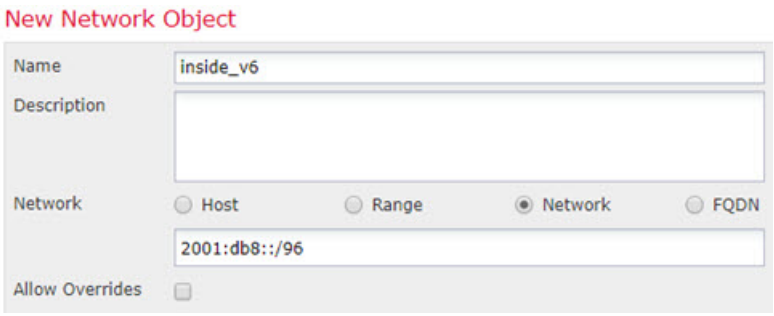
Step 1 Create the network object that defines the inside IPv6 network.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the inside IPv6 network.

Name the network object (for example, inside_v6) and enter the network address, 2001:db8::/96.



d) Click **Save**.

Step 2 Create the manual NAT rule to translate the IPv6 network to IPv4 and back again.

a) Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

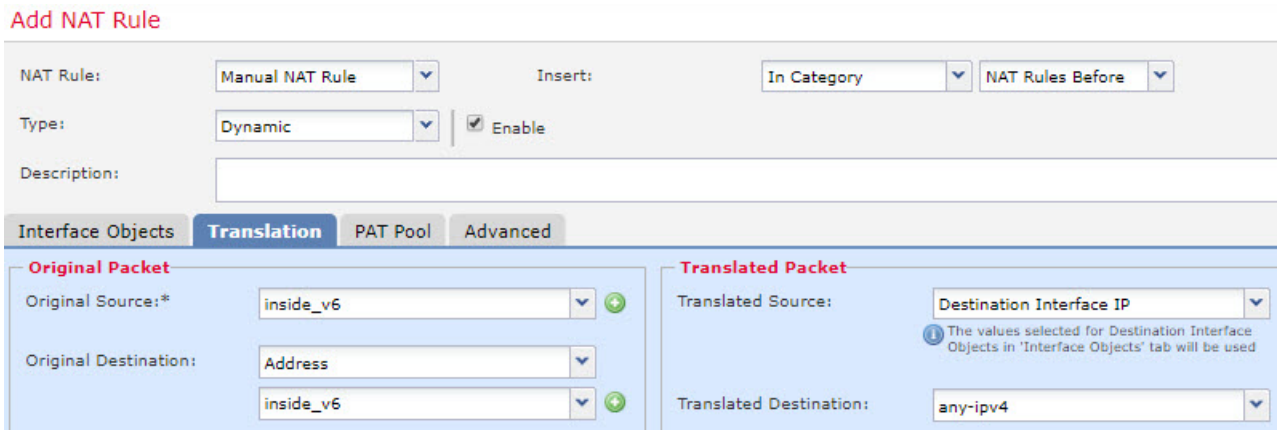
- **NAT Rule** = Manual NAT Rule.
- **Type** = Dynamic.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = inside_v6 network object.
- **Translated Source** = **Destination Interface IP**.
- **Original Destination** = inside_v6 network object.
- **Translated Destination** = any-ipv4 network object.



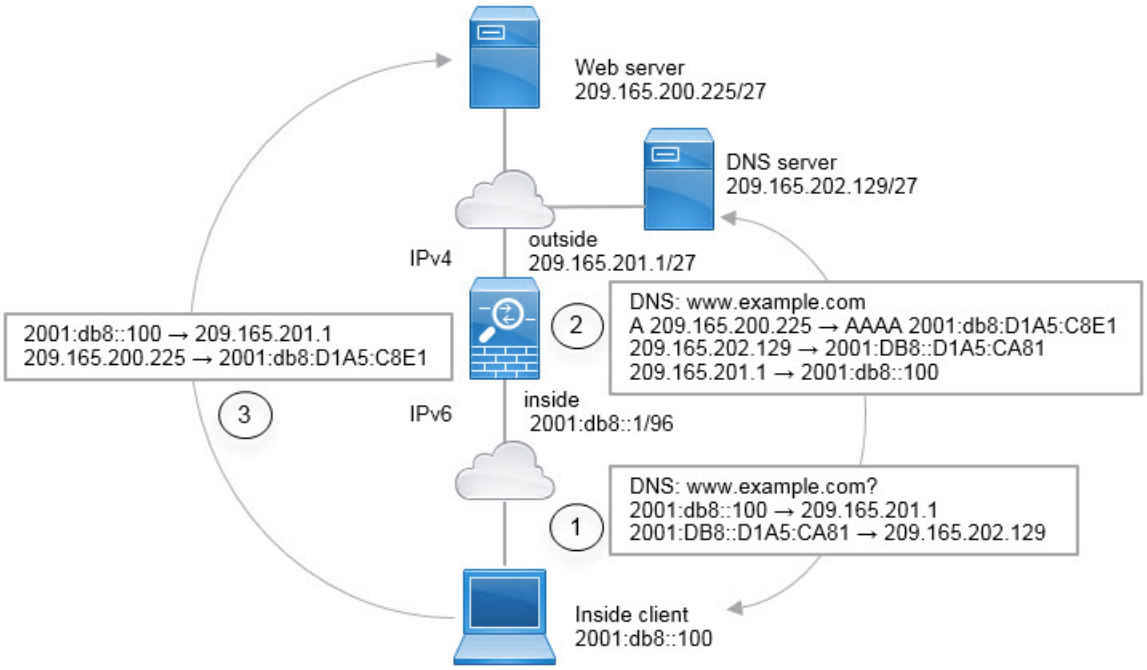
f) Click **OK**.

With this rule, any traffic from the 2001:db8::/96 subnet on the inside interface going to the outside interface gets a NAT64 PAT translation using the IPv4 address of the outside interface. Conversely, any IPv4 address on the outside network coming to the inside interface is translated to an address on the 2001:db8::/96 network using the embedded IPv4 address method.

g) Click **Save** on the NAT rules page.

NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet and DNS Translation

Following is a typical example where you have an inside IPv6-only network, but there are some IPv4-only services on the outside Internet that internal users need.



In this example, you translate the inside IPv6 network to IPv4 using dynamic interface PAT with the IP address of the outside interface. Outside IPv4 traffic is statically translated to addresses on the 2001:db8::/96 network, allowing transmission on the inside network. You enable DNS rewrite on the NAT46 rule, so that replies from the external DNS server can be converted from A (IPv4) to AAAA (IPv6) records, and the addresses converted from IPv4 to IPv6.

Following is a typical sequence for a web request where a client at 2001:DB8::100 on the internal IPv6 network tries to open www.example.com.

1. The client's computer sends a DNS request to the DNS server at 2001:DB8::D1A5:CA81. The NAT rules make the following translations to the source and destination in the DNS request:
   - 2001:DB8::100 to a unique port on 209.165.201.1 (The NAT64 interface PAT rule.)
   - 2001:DB8::D1A5:CA81 to 209.165.202.129 (The NAT46 rule. D1A5:CA81 is the IPv6 equivalent of 209.165.202.129.)

2. The DNS server responds with an A record indicating that www.example.com is at 209.165.200.225. The NAT46 rule, with DNS rewrite enabled, converts the A record to the IPv6-equivalent AAAA record, and translates 209.165.200.225 to 2001:db8:D1A5:C8E1in the AAAA record. In addition, the source and destination addresses in the DNS response are untranslated:
   - 209.165.202.129 to 2001:DB8::D1A5:CA81
   - 209.165.201.1 to 2001:db8::100
3. The IPv6 client now has the IP address of the web server, and makes an HTTP request to www.example.com at 2001:db8:D1A5:C8E1. (D1A5:C8E1 is the IPv6 equivalent of 209.165.200.225.) The source and destination of the HTTP request are translated:
   - 2001:DB8::100 to a unique port on 209.156.101.54 (The NAT64 interface PAT rule.)
   - 2001:db8:D1A5:C8E1 to 209.165.200.225 (The NAT46 rule.)

The following procedure explains how to configure this example.

Before you begin

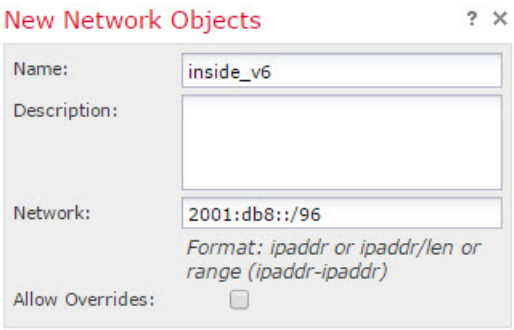
Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

Step 1 Create the network objects that define the inside IPv6 and outside IPv4 networks.

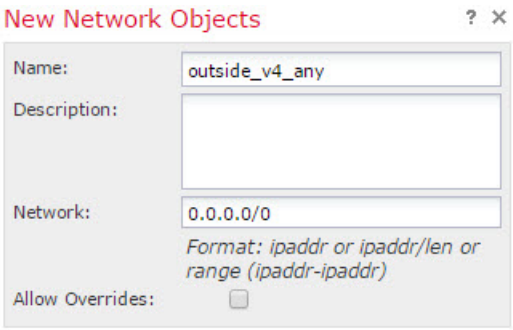
a) Choose **Objects** > **Object Management**.
b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.
c) Define the inside IPv6 network.

Name the network object (for example, inside_v6) and enter the network address, 2001:db8::/96.



d) Click **Save**.
e) Click **Add Network** > **Add Object** and define the outside IPv4 network.

Name the network object (for example, outside_v4_any) and enter the network address 0.0.0.0/0.



f) Click **Save**.

Step 2 Configure the NAT64 dynamic PAT rule for the inside IPv6 network.

Step 3 Configure the static NAT46 rule for the outside IPv4 network.

a) Click **Add Rule**.

b) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

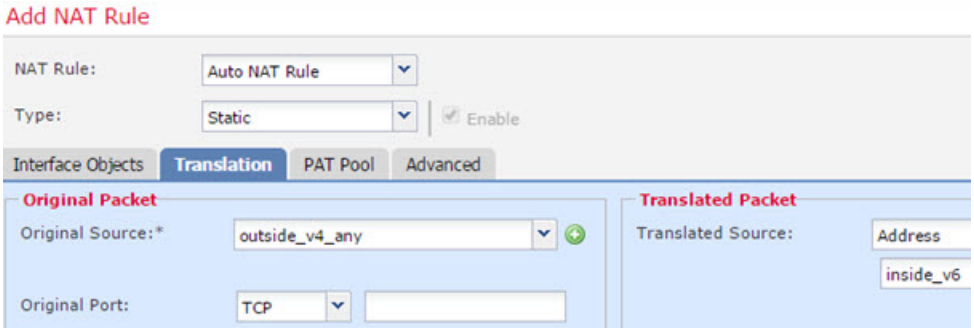
c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = outside.
- **Destination Interface Objects** = inside.

d) On **Translation**, configure the following:

- **Original Source** = outside_v4_any network object.
- **Translated Source** > **Address** = inside_v6 network object.

e) On **Advanced**, select **Translate DNS replies that match this rule**.



f) Click **OK**.

With this rule, any IPv4 address on the outside network coming to the inside interface is translated to an address on the 2001:db8::/96 network using the embedded IPv4 address method. In addition, DNS responses are converted from A (IPv4) to AAAA (IPv6) records, and the addresses converted from IPv4 to IPv6.

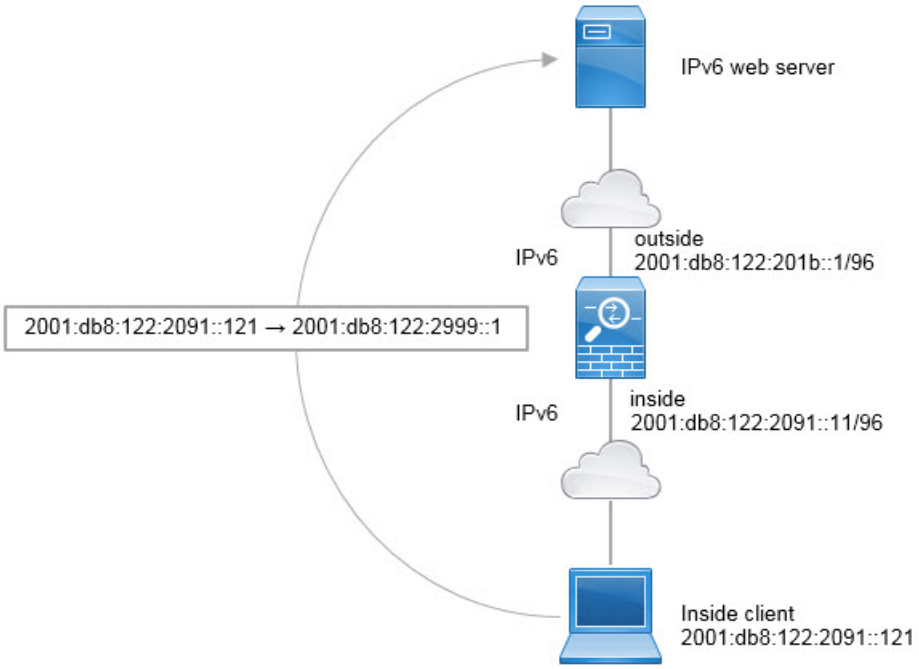
NAT66: Translating IPv6 Addresses to Different IPv6 Addresses

When going from an IPv6 network to another IPv6 network, you can translate the addresses to different IPv6 addresses on the outside network. We recommend using static NAT. Although you can use dynamic NAT or PAT, IPv6 addresses are in such large supply, you do not have to use dynamic NAT.

Because you are not translating between different address types, you need a single rule for NAT66 translations. You can easily model these rules using auto NAT. However, if you do not want to allow returning traffic, you can make the static NAT rule unidirectional using manual NAT only.

NAT66 Example, Static Translation between Networks

You can configure a static translation between IPv6 address pools using auto NAT. The following example explains how to convert inside addresses on the 2001:db8:122:2091::/96 network to outside addresses on the 2001:db8:122:2999::/96 network.



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

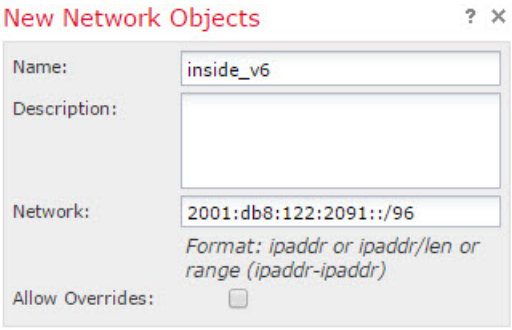
Step 1 Create the network objects that define the inside IPv6 and outside IPv6 NAT networks.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the inside IPv6 network.

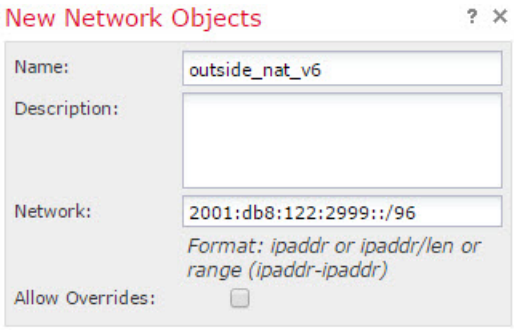
Name the network object (for example, inside_v6) and enter the network address, 2001:db8:122:2091::/96.



d) Click **Save**.

e) Click **Add Network** > **Add Object** and define the outside IPv6 NAT network.

Name the network object (for example, outside_nat_v6) and enter the network address 2001:db8:122:2999::/96.



f) Click **Save**.

Step 2 Configure the static NAT rule for the inside IPv6 network.

a) Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

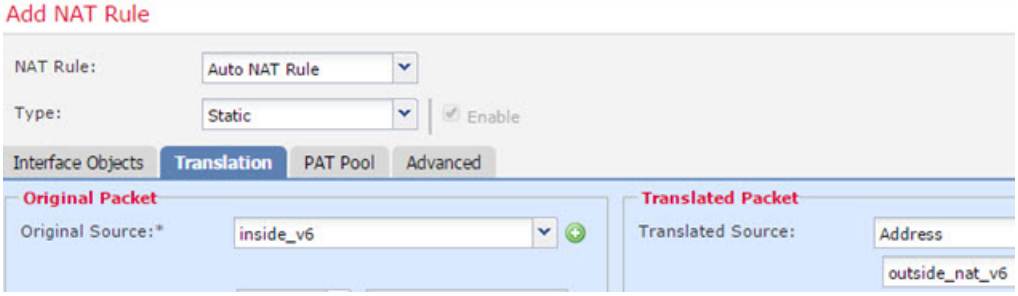
- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = inside_v6 network object.
- **Translated Source** > **Address** = outside_nat_v6 network object.



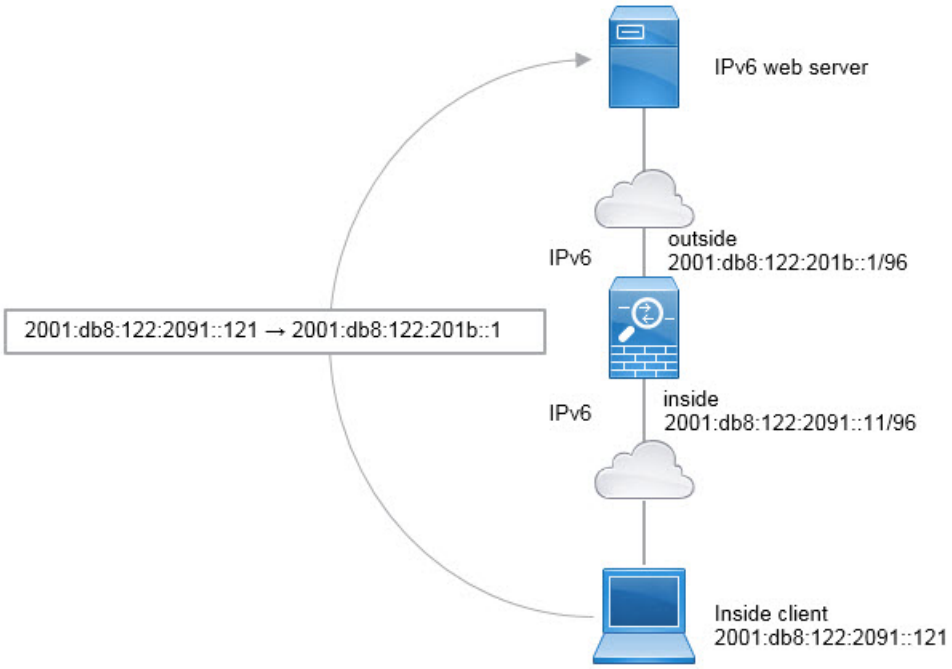
f) Click **OK**.

With this rule, any traffic from the 2001:db8:122:2091::/96 subnet on the inside interface going to the outside interface gets a static NAT66 translation to an address on the 2001:db8:122:2999::/96 network.

NAT66 Example, Simple IPv6 Interface PAT

A simple approach for implementing NAT66 is to dynamically assign internal addresses to different ports on the outside interface IPv6 address.

When you configure an interface PAT rule for NAT66, all the global addresses that are configured on that interface are used for PAT mapping. Link-local or site-local addresses for the interface are not used for PAT.



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

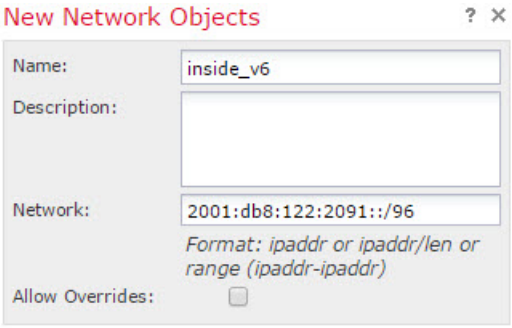
Step 1 Create the network object that defines the inside IPv6 network.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the inside IPv6 network.

Name the network object (for example, inside_v6) and enter the network address, 2001:db8:122:2091::/96.



d) Click **Save**.

Step 2 Configure the dynamic PAT rule for the inside IPv6 network.

a) Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Dynamic.

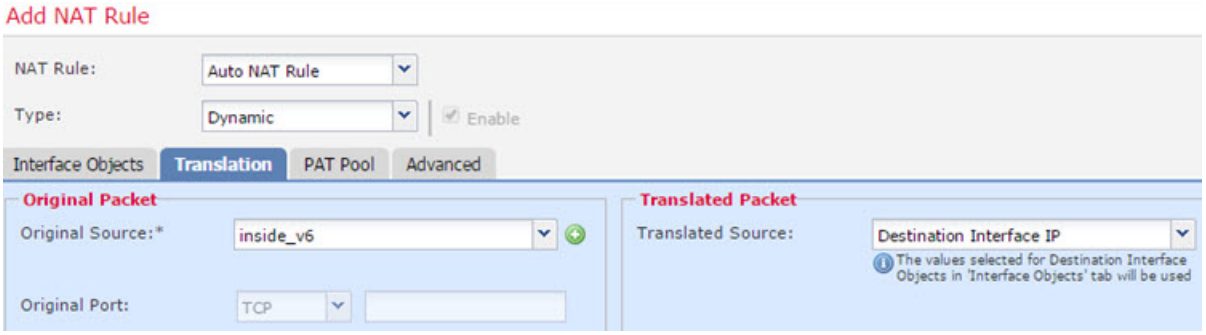
d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = inside_v6 network object.
- **Translated Source** = **Destination Interface IP**.

f) On t**Advanced**, select **IPv6**, which indicates that the IPv6 address of the destination interface should be used.



g) Click **OK**.

With this rule, any traffic from the 2001:db8:122:2091::/96 subnet on the inside interface going to the outside interface gets a NAT66 PAT translation to one of the IPv6 global addresses configured for the outside interface.

Monitoring NAT

To monitor and troubleshoot NAT connections, log into the device CLI and use the following commands.

- **show nat** displays the NAT rules and per-rule hit counts. There are additional keywords to show other aspects of NAT.
- **show xlate** displays the actual NAT translations that are currently active.
- **clear xlate** lets you remove an active NAT translation. You might need to remove active translations if you alter NAT rules, because existing connections continue to use the old translation slot until the connection ends. Clearing a translation allows the system to build a new translation for a client on the client's next connection attempt based on your new rules.

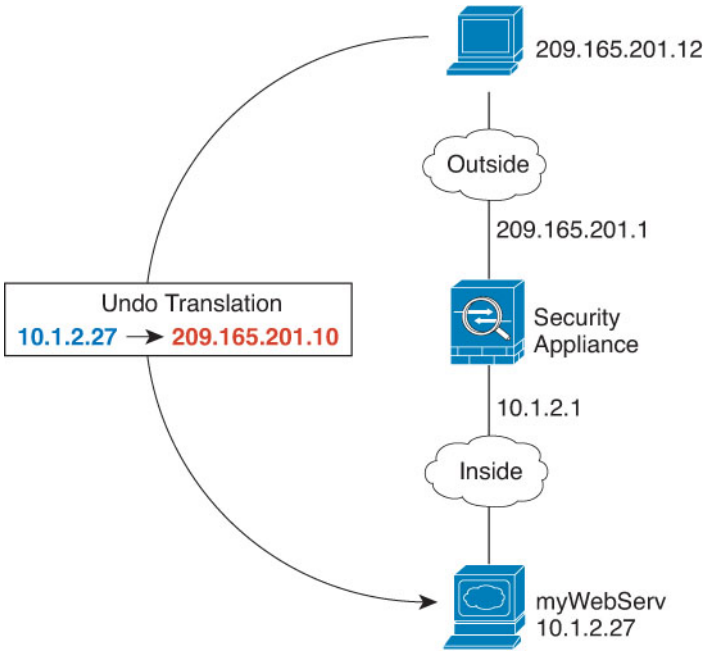
Examples for NAT

The following topics provide examples for configuring NAT on Threat Defense devices.

Providing Access to an Inside Web Server (Static Auto NAT)

The following example performs static NAT for an inside web server. The real address is on a private network, so a public address is required. Static NAT is necessary so hosts can initiate traffic to the web server at a fixed address.

Figure 45: Static NAT for an Inside Web Server



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the web server. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

Step 1 Create the network objects that define the server's private and public host addresses.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the web server's private address.

Name the network object (for example, WebServerPrivate) and enter the real host IP address, 10.1.2.27.



d) Click **Save**.

e) Click **Add Network** > **Add Object** and define the public address.

Name the network object (for example, WebServerPublic) and enter the host address 209.165.201.10.



f) Click **Save**.

Step 2 Configure static NAT for the object.

a) Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

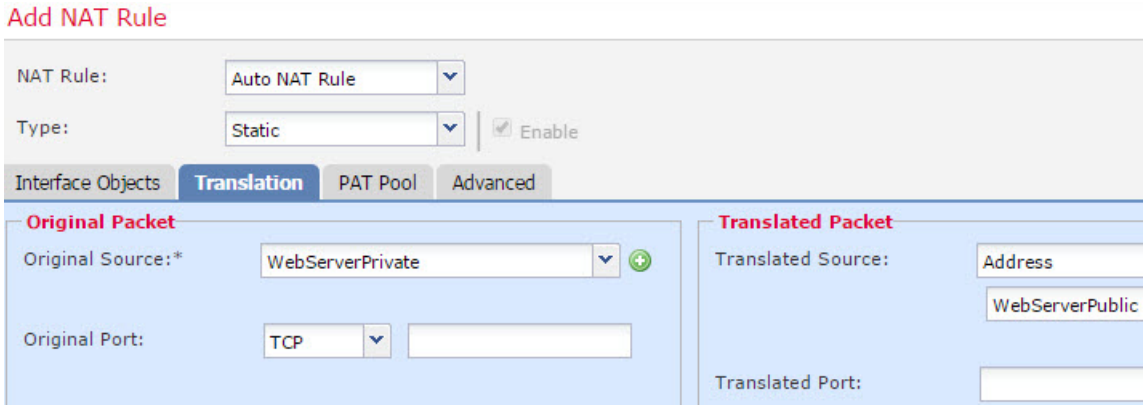
- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = WebServerPrivate network object.
- **Translated Source** > **Address**= WebServerPublic network object.



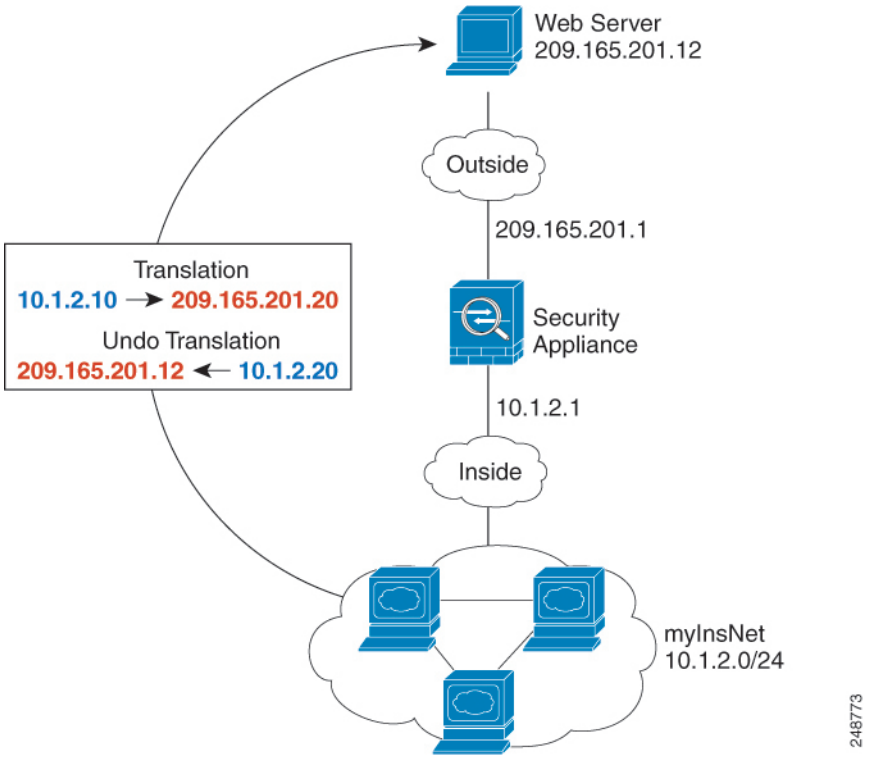
f) Click **Save**.

Step 3 Click **Save** on the NAT rule page.

Dynamic Auto NAT for Inside Hosts and Static NAT for an Outside Web Server

The following example configures dynamic NAT for inside users on a private network when they access the outside. Also, when inside users connect to an outside web server, that web server address is translated to an address that appears to be on the inside network.

Figure 46: Dynamic NAT for Inside, Static NAT for Outside Web Server



Before you begin

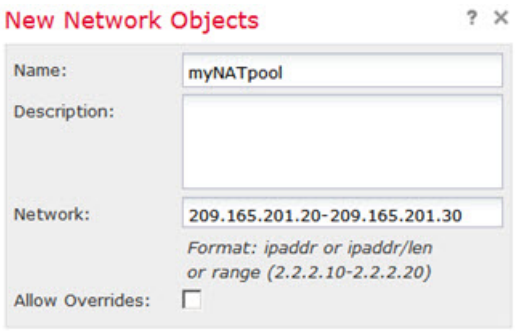
Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the web server. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

Step 1 Create a network object for the dynamic NAT pool to which you want to translate the inside addresses.

a) Choose **Objects** > **Object Management**.
b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.
c) Define the dynamic NAT pool.

Name the network object (for example, myNATpool) and enter the network range 209.165.201.20-209.165.201.30.

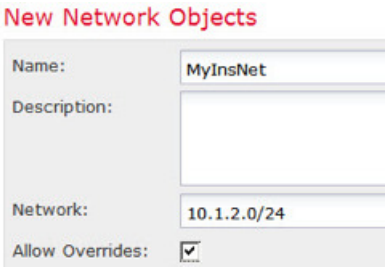


d) Click **Save**.

Step 2 Create a network object for the inside network.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, MyInsNet) and enter the network address 10.1.2.0/24.



c) Click **Save**.

Step 3 Create a network object for the outside web server.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, MyWebServer) and enter the host address 209.165.201.12.

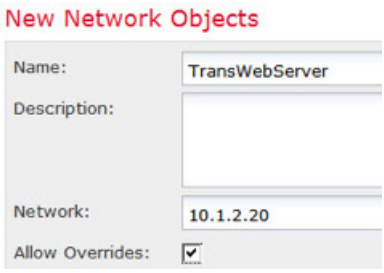


c) Click **Save**.

Step 4 Create a network object for the translated web server address.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, TransWebServer) and enter the host address 10.1.2.20.



c) Click **Save**.

Step 5 Configure dynamic NAT for the inside network using the dynamic NAT pool object.

a) Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

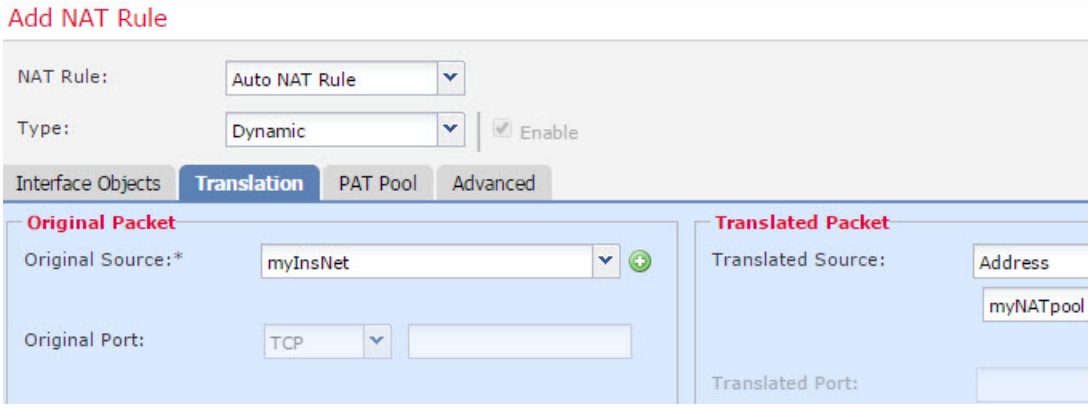
- **NAT Rule** = Auto NAT Rule.
- **Type** = Dynamic.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = myInsNet network object.
- **Translated Source** > **Address**= myNATpool network group.



f) Click **Save**.

Step 6 Configure static NAT for the web server.

a) Click **Add Rule**.

b) Configure the following properties:

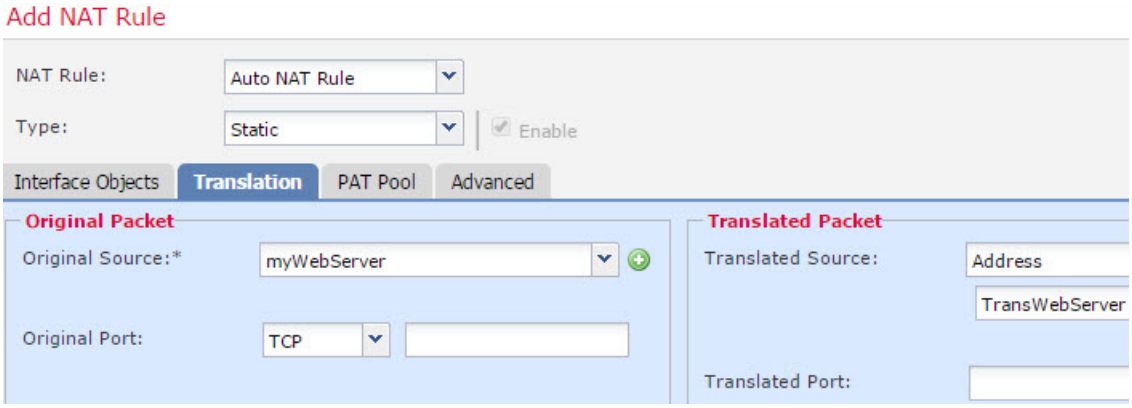
- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = outside.
- **Destination Interface Objects** = inside.

d) On **Translation**, configure the following:

- **Original Source** = myWebServer network object.
- **Translated Source** > **Address**= TransWebServer network object.



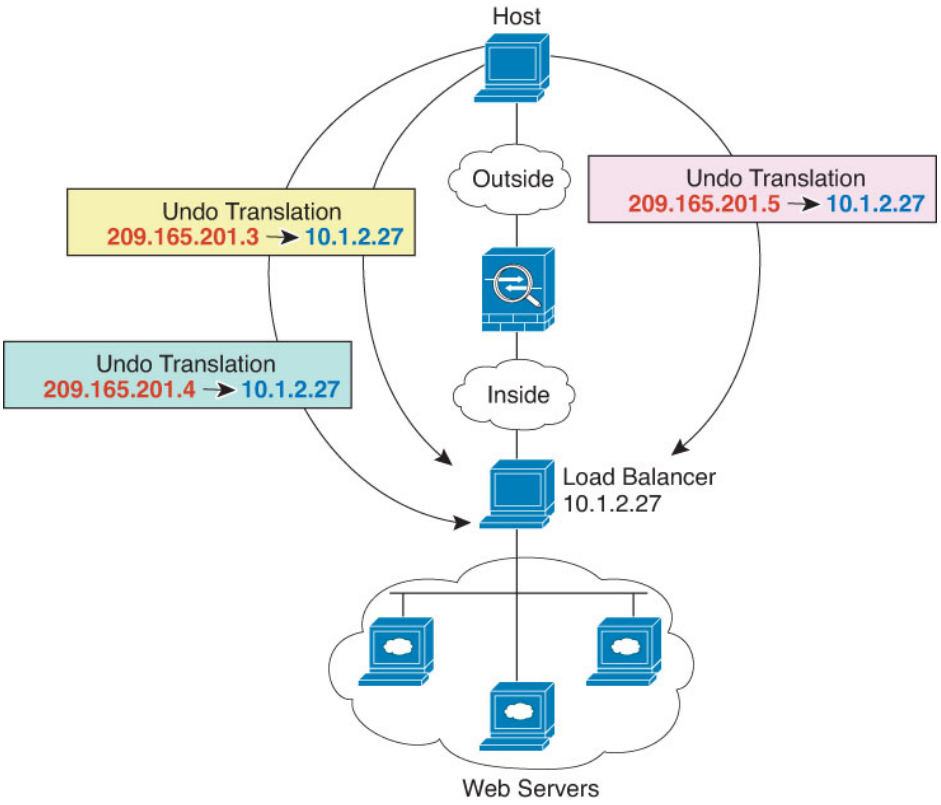
e) Click **Save**.

Step 7 Click **Save** on the NAT rule page.

Inside Load Balancer with Multiple Mapped Addresses (Static Auto NAT, One-to-Many)

The following example shows an inside load balancer that is translated to multiple IP addresses. When an outside host accesses one of the mapped IP addresses, it is untranslated to the single load balancer address. Depending on the URL requested, it redirects traffic to the correct web server.

Figure 47: Static NAT with One-to-Many for an Inside Load Balancer



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the web server. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

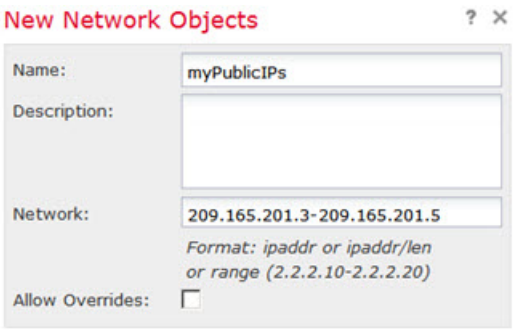
Step 1 Create a network object for the addresses to which you want to map the load balancer.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the addresses.

Name the network object (for example, myPublicIPs) and enter the network range 209.165.201.3-209.165.201.5.

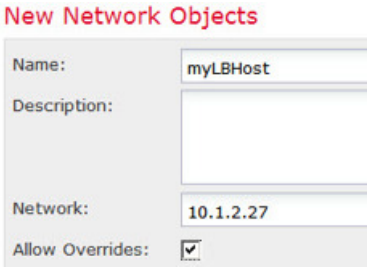


d) Click **Save**.

Step 2 Create a network object for the load balancer.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, myLBHost), enter the host address 10.1.2.27.



c) Click **Save**.

Step 3 Configure static NAT for the load balancer.

a) Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

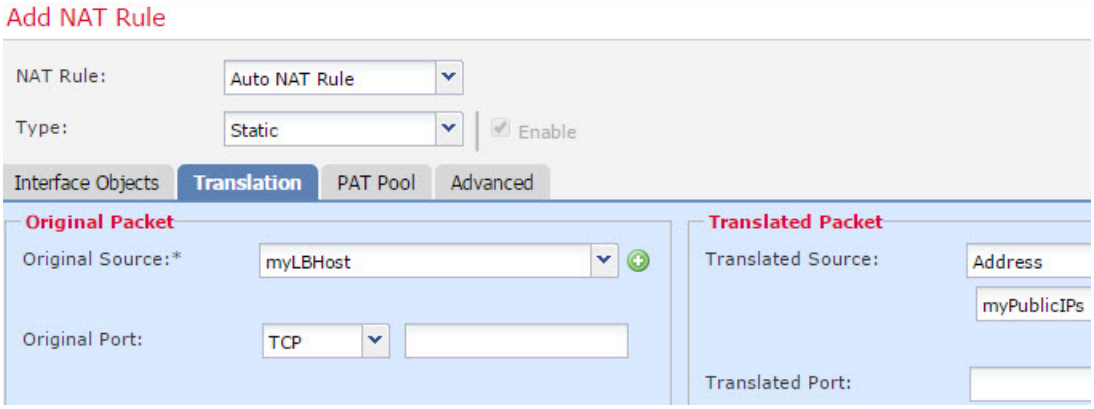
- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = myLBHost network object.
- **Translated Source** > **Address**= myPublicIPs network group.



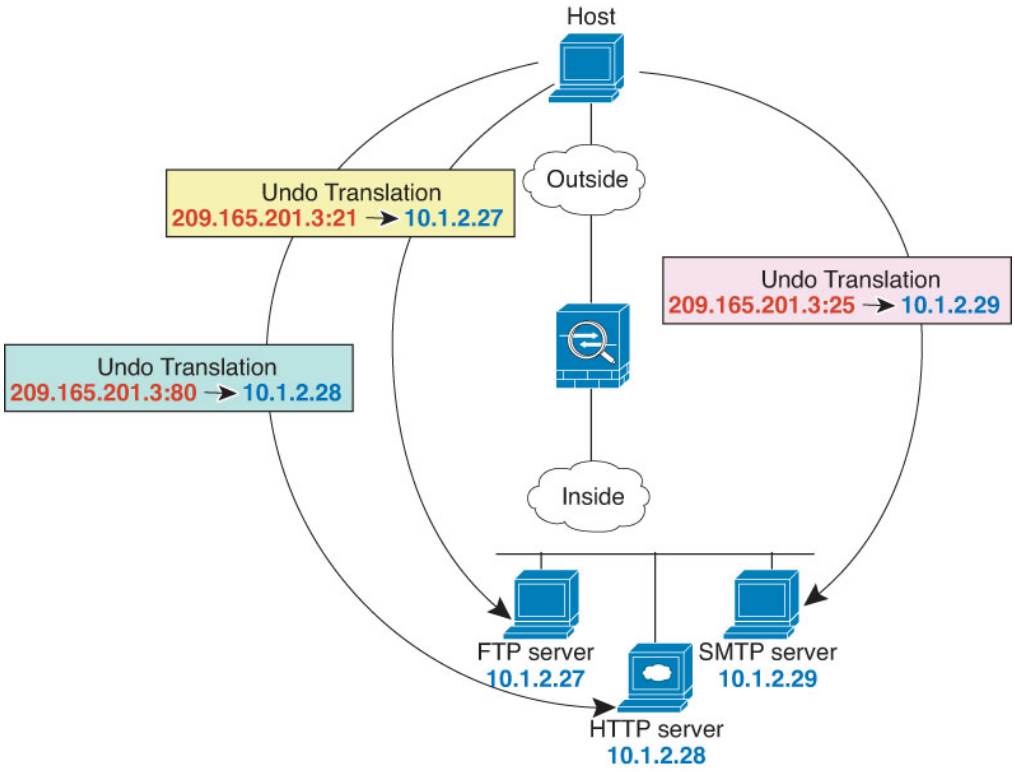
f) Click **Save**.

Step 4 Click **Save** on the NAT rule page.

Single Address for FTP, HTTP, and SMTP (Static Auto NAT-with-Port-Translation)

The following static NAT-with-port-translation example provides a single address for remote users to access FTP, HTTP, and SMTP. These servers are actually different devices on the real network, but for each server, you can specify static NAT-with-port-translation rules that use the same mapped IP address, but different ports.

Figure 48: Static NAT-with-Port-Translation



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the servers. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

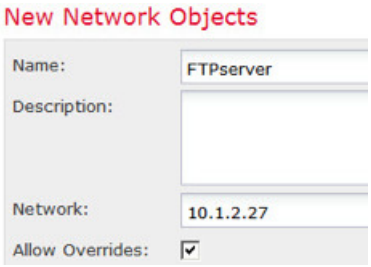
Procedure

Step 1 Create a network object for the FTP server.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Name the network object (for example, FTPserver), and enter the real IP address for the FTP server, 10.1.2.27.



d) Click **Save**.

Step 2 Create a network object for the HTTP server.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, HTTPserver), enter the host address 10.1.2.28.

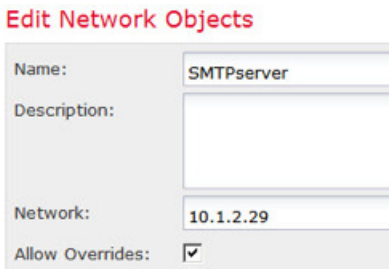


c) Click **Save**.

Step 3 Create a network object for the SMTP server.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, SMTPserver), enter the host address 10.1.2.29.

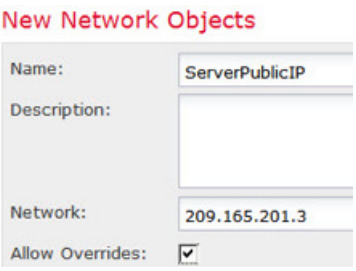


c) Click **Save**.

Step 4 Create a network object for the public IP address used for the three servers.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, ServerPublicIP) and enter the host address 209.165.201.3.



c) Click **Save**.

Step 5 Configure static NAT with port translation for the FTP server, mapping the FTP port to itself.

a) Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.

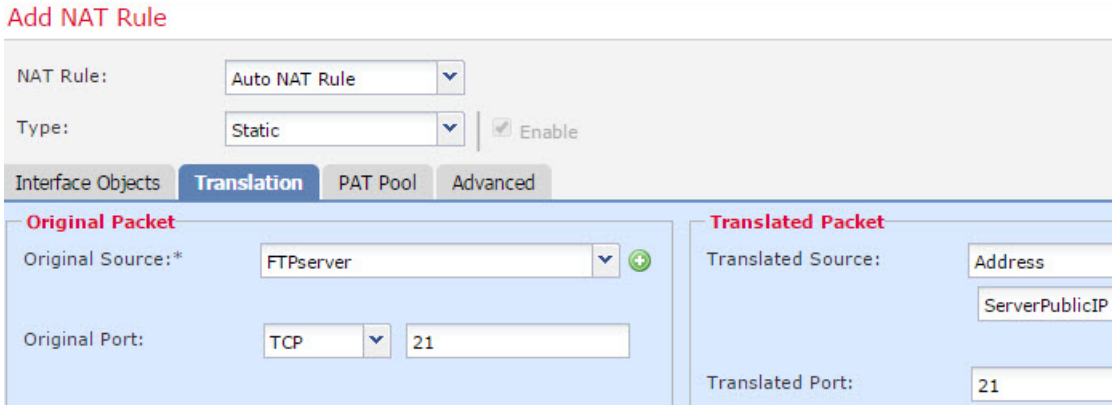
- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = FTPserver network object.
- **Translated Source** > **Address**= ServerPublicIP network object.
- **Original Port** > **TCP** = 21.
- **Translated Port** = 21.



f) Click **Save**.

Step 6 Configure static NAT with port translation for the HTTP server, mapping the HTTP port to itself.

a) Click **Add Rule**.

b) Configure the following properties:

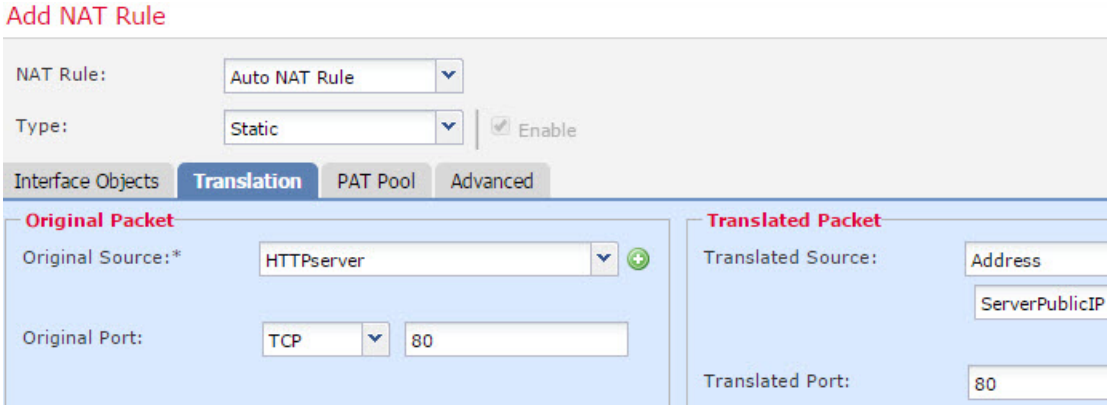
- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

d) On **Translation**, configure the following:

- **Original Source** = HTTPserver network object.
- **Translated Source** > **Address**= ServerPublicIP network object.
- **Original Port** > **TCP** = 80.
- **Translated Port** = 80.



e) Click **Save**.

Step 7 Configure static NAT with port translation for the SMTP server, mapping the SMTP port to itself.

a) Click **Add Rule**.

b) Configure the following properties:

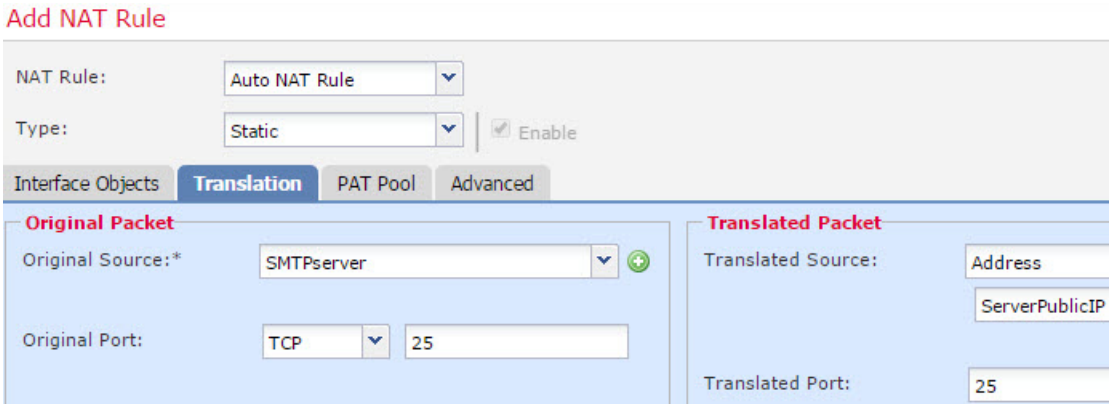
- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

d) On **Translation**, configure the following:

- **Original Source** = SMTPserver network object.
- **Translated Source** > **Address**= ServerPublicIP network object.
- **Original Port** > **TCP** = 25.
- **Translated Port** = 25.



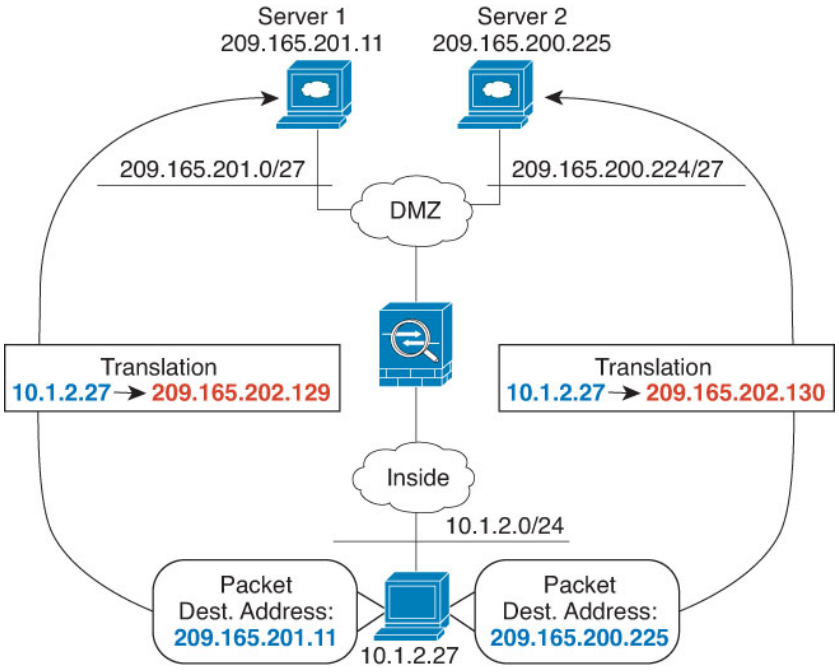
e) Click **Save**.

Step 8 Click **Save** on the NAT rule page.

Different Translation Depending on the Destination (Dynamic Manual PAT)

The following figure shows a host on the 10.1.2.0/24 network accessing two different servers. When the host accesses the server at 209.165.201.11, the real address is translated to 209.165.202.129:*port*. When the host accesses the server at 209.165.200.225, the real address is translated to 209.165.202.130:*port*.

Figure 49: Manual NAT with Different Destination Addresses



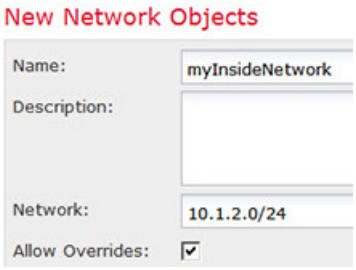
Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the servers. In this example, we will assume the interface objects are security zones named **inside** and **dmz**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

Step 1 Create a network object for the inside network.

a) Choose **Objects** > **Object Management**.
b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.
c) Name the network object (for example, myInsideNetwork), and enter the real network address, 10.1.2.0/24.

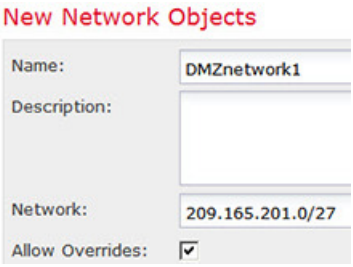


d) Click **Save**.

Step 2 Create a network object for the DMZ network 1.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, DMZnetwork1) and enter the network address 209.165.201.0/27 (subnet mask of 255.255.255.224).

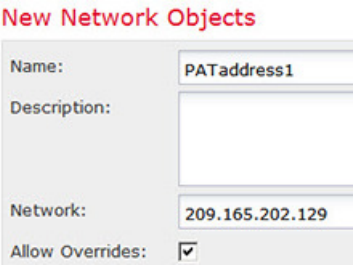


c) Click **Save**.

Step 3 Create a network object for the PAT address for DMZ network 1.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, PATaddress1) and enter the host address 209.165.202.129.

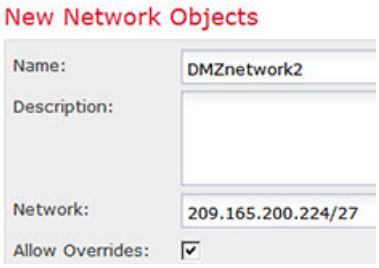


c) Click **Save**.

Step 4 Create a network object for the DMZ network 2.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, DMZnetwork2) and enter the network address 209.165.200.224/27 (subnet mask of 255.255.255.224).

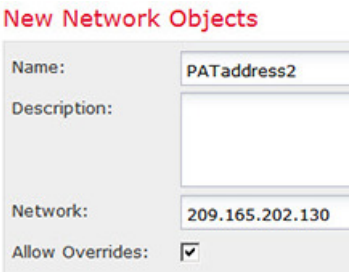


c) Click **Save**.

Step 5 Create a network object for the PAT address for DMZ network 2.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, PATaddress2) and enter the host address 209.165.202.130.



c) Click **Save**.

Step 6 Configure dynamic manual PAT for DMZ network 1.

a) Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Manual NAT Rule.
- **Type** = Dynamic.

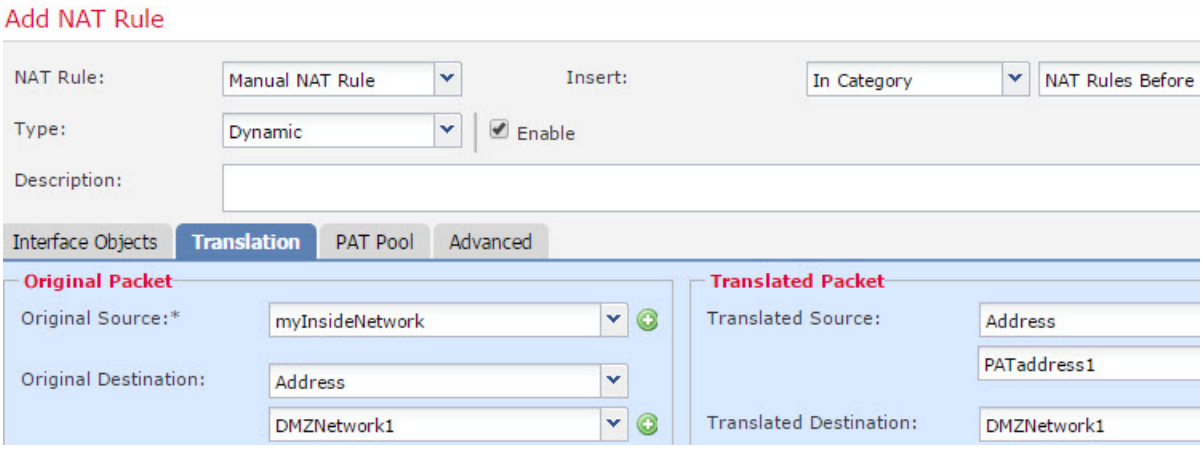
d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = dmz.

e) On **Translation**, configure the following:

- **Original Source** = myInsideNetwork network object.
- **Translated Source** > **Address**= PATaddress1 network object.
- **Original Destination** > **Address** = DMZnetwork1 network object.
- **Translated Destination** = DMZnetwork1 network object.

Note Because you do not want to translate the destination address, you need to configure identity NAT for it by specifying the same address for the original and translated destination addresses. Leave all of the port fields blank.



f) Click **Save**.

Step 7 Configure dynamic manual PAT for DMZ network 2.

a) Click **Add Rule**.

b) Configure the following properties:

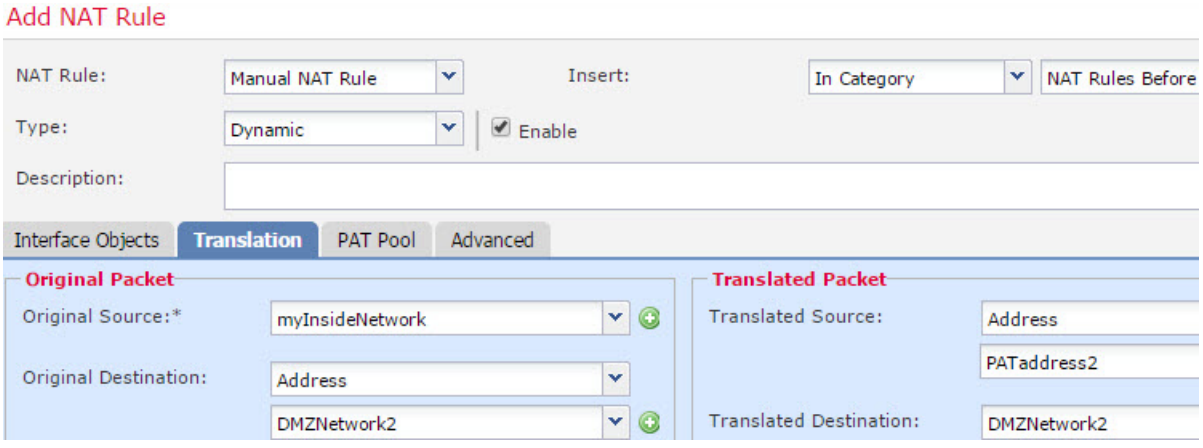
- **NAT Rule** = Manual NAT Rule.
- **Type** = Dynamic.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = dmz.

d) On **Translation**, configure the following:

- **Original Source** = myInsideNetwork network object.
- **Translated Source** > **Address**= PATaddress2 network object.
- **Original Destination** > **Address** = DMZnetwork2 network object.
- **Translated Destination** = DMZnetwork2 network object.



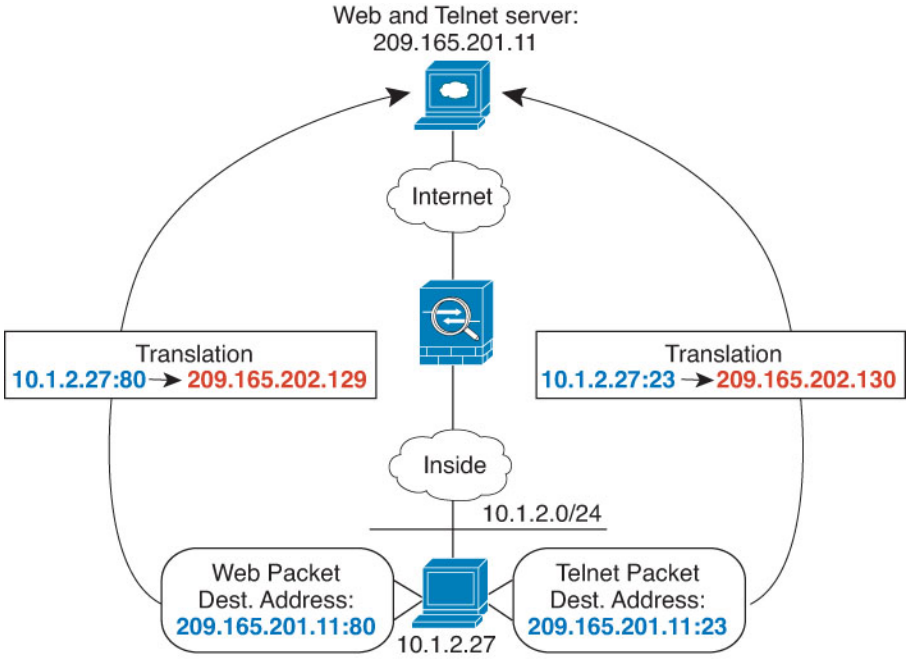
e) Click **Save**.

Step 8 Click **Save** on the NAT rule page.

Different Translation Depending on the Destination Address and Port (Dynamic Manual PAT)

The following figure shows the use of source and destination ports. The host on the 10.1.2.0/24 network accesses a single host for both web services and Telnet services. When the host accesses the server for Telnet services, the real address is translated to 209.165.202.129:*port*. When the host accesses the same server for web services, the real address is translated to 209.165.202.130:*port*.

Figure 50: Manual NAT with Different Destination Ports



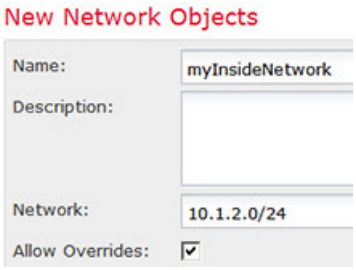
Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the servers. In this example, we will assume the interface objects are security zones named **inside** and **dmz**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

Step 1 Create a network object for the inside network.

a) Choose **Objects** > **Object Management**.
b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.
c) Name the network object (for example, myInsideNetwork) and enter the real network address, 10.1.2.0/24.

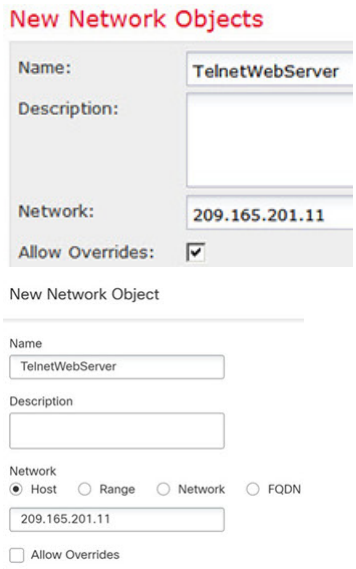


d) Click **Save**.

Step 2 Create a network object for the Telnet/Web server.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, TelnetWebServer) and enter the host address 209.165.201.11.

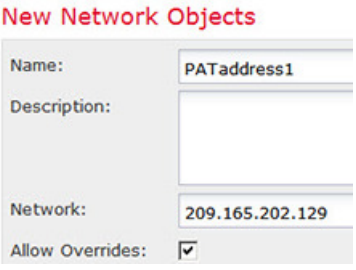


c) Click **Save**.

Step 3 Create a network object for the PAT address when using Telnet.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, PATaddress1) and enter the host address 209.165.202.129.

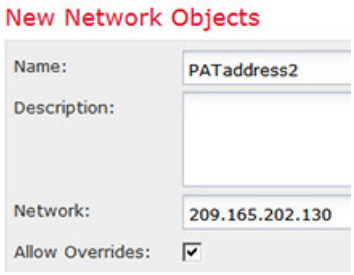


c) Click **Save**.

Step 4 Create a network object for the PAT address when using HTTP.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, PATaddress2) and enter the host address 209.165.202.130.



c) Click **Save**.

Step 5 Configure dynamic manual PAT for Telnet access.

a) Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Manual NAT Rule.
- **Type** = Dynamic.

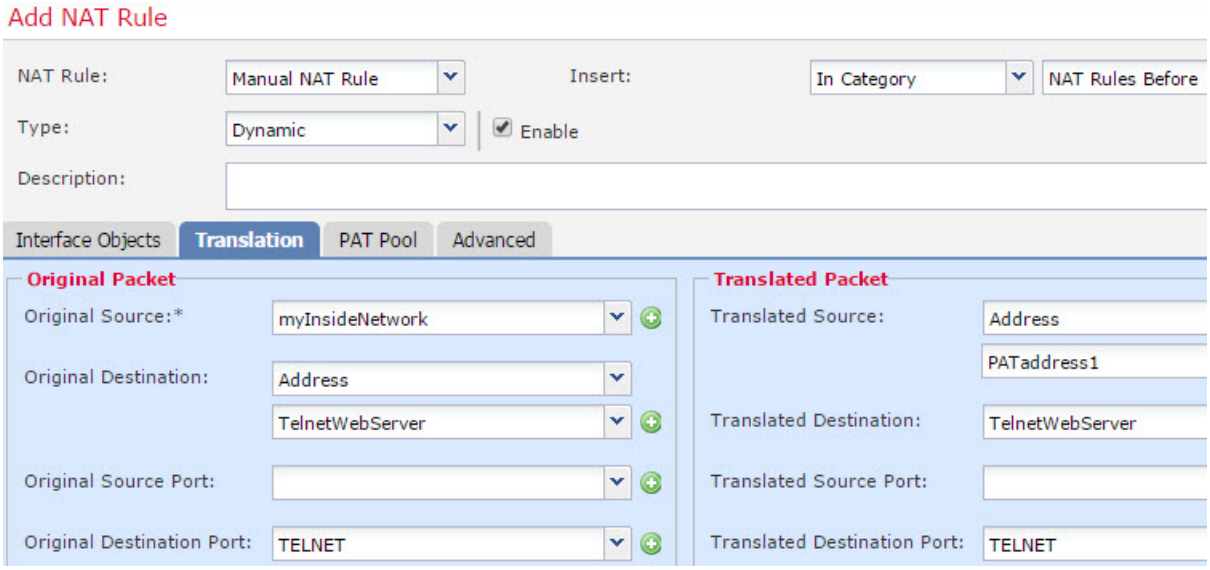
d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = dmz.

e) On **Translation**, configure the following:

- **Original Source** = myInsideNetwork network object.
- **Translated Source** > **Address**= PATaddress1 network object.
- **Original Destination** > **Address** = TelnetWebServer network object.
- **Translated Destination** = TelnetWebServer network object.
- **Original Destination Port** = TELNET port object (system-defined).
- **Translated Destination Port** = TELNET port object (system-defined).

Note Because you do not want to translate the destination address or port, you need to configure identity NAT for them by specifying the same address for the original and translated destination addresses, and the same port for the original and translated port.



f) Click **Save**.

Step 6 Configure dynamic manual PAT for web access.

a) Click **Add Rule**.

b) Configure the following properties:

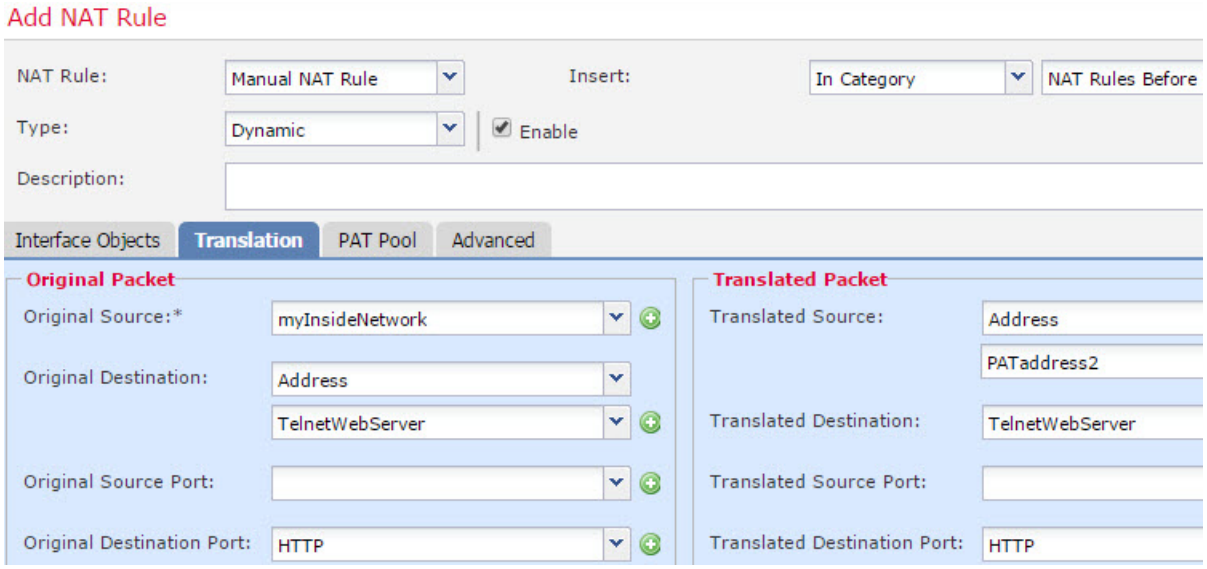
- **NAT Rule** = Manual NAT Rule.
- **Type** = Dynamic.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = dmz.

d) On **Translation**, configure the following:

- **Original Source** = myInsideNetwork network object.
- **Translated Source** > **Address**= PATaddress2 network object.
- **Original Destination** > **Address** = TelnetWebServer network object.
- **Translated Destination** = TelnetWebServer network object.
- **Original Destination Port** = HTTP port object (system-defined).
- **Translated Destination Port** = HTTP port object (system-defined).



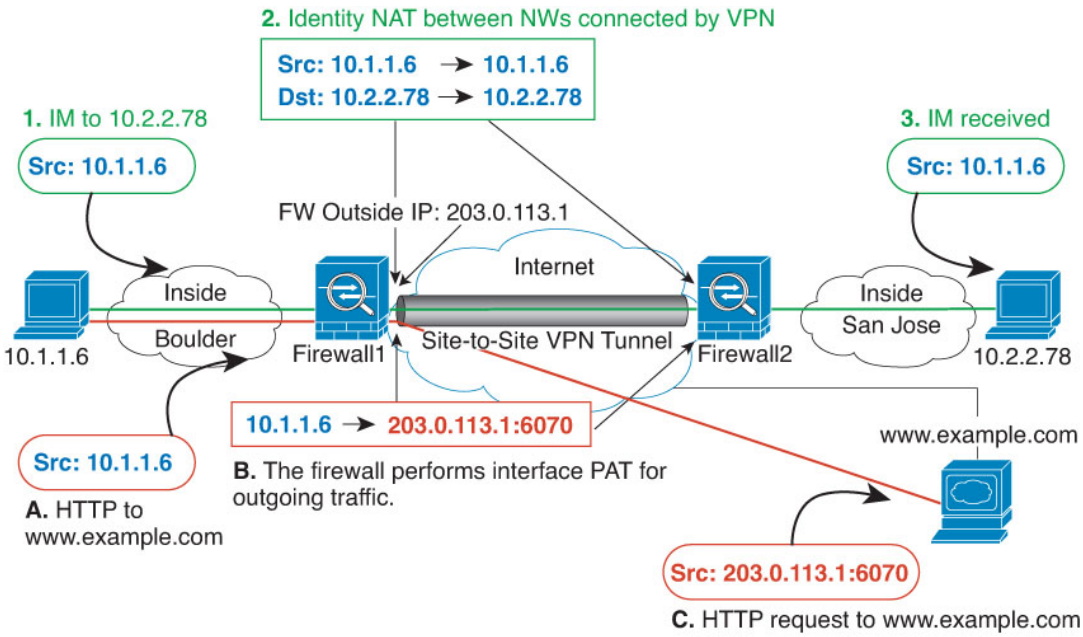
e) Click **Save**.

Step 7 Click **Save** on the NAT rule page.

NAT and Site-to-Site VPN

The following figure shows a site-to-site tunnel connecting the Boulder and San Jose offices. For traffic that you want to go to the Internet (for example from 10.1.1.6 in Boulder to www.example.com), you need a public IP address provided by NAT to access the Internet. The below example uses interface PAT rules. However, for traffic that you want to go over the VPN tunnel (for example from 10.1.1.6 in Boulder to 10.2.2.78 in San Jose), you do not want to perform NAT; you need to exempt that traffic by creating an identity NAT rule. Identity NAT simply translates an address to the same address.

Figure 51: Interface PAT and Identity NAT for Site-to-Site VPN



The following example explains the configuration for Firewall1 (Boulder).

Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the devices in the VPN. In this example, we will assume the interface objects are security zones named **inside-boulder** and **outside-boulder** for the Firewall1 (Boulder) interfaces. To configure interface objects, select **Objects** > **Object Management**, then select **Interfaces**.

Procedure

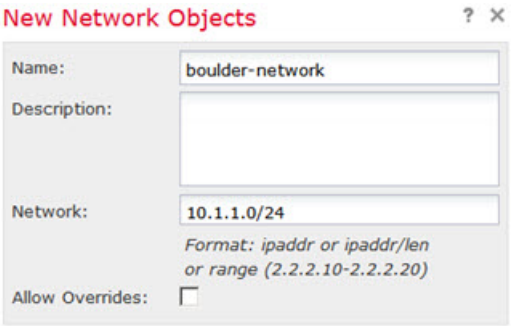
Step 1 Create the objects to define the various networks.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Identify the Boulder inside network.

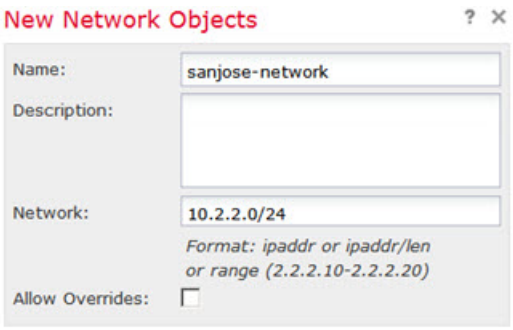
Name the network object (for example, boulder-network) and enter the network address, 10.1.1.0/24.



d) Click **Save**.

e) Click **Add Network** > **Add Object** and define the inside San Jose network.

Name the network object (for example, sanjose-network) and enter the network address 10.2.2.0/24.



f) Click **Save**.

Step 2 Configure manual identity NAT for the Boulder network when going over the VPN to San Jose on Firewall1 (Boulder).

a) Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Manual NAT Rule.
- **Type** = Static.

d) On **Interface Objects**, configure the following:

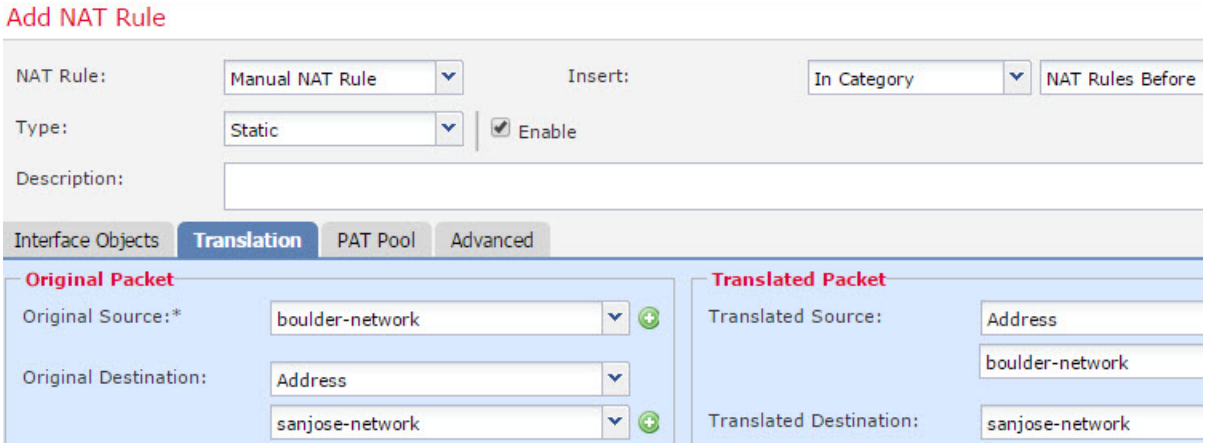
- **Source Interface Objects** = inside-boulder.
- **Destination Interface Objects** = outside-boulder.

e) On **Translation**, configure the following:

- **Original Source** = boulder-network object.
- **Translated Source** > **Address** = boulder-network object.
- **Original Destination** > **Address** = sanjose-network object.
- **Translated Destination** = sanjose-network object.

Note Because you do not want to translate the destination address, you need to configure identity NAT for it by specifying the same address for the original and translated destination addresses. Leave all of the port fields blank. This rule configures identity NAT for both source and destination.

f) On **Advanced**, select **Do not proxy ARP on Destination interface**.



g) Click **Save**.

Step 3 Configure manual dynamic interface PAT when going to the Internet for the inside Boulder network on Firewall1 (Boulder).

a) Click **Add Rule**.

b) Configure the following properties:

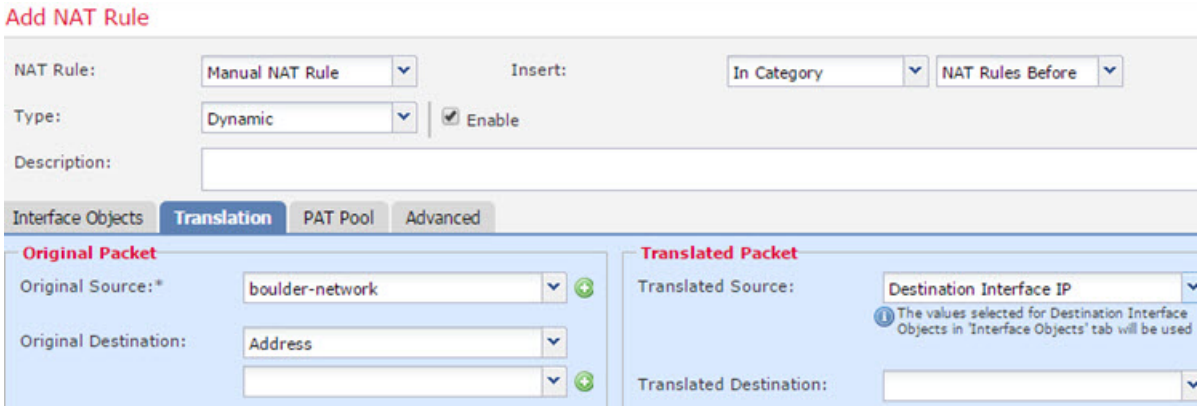
- **NAT Rule** = Manual NAT Rule.
- **Type** = Dynamic.
- **Insert Rule** = any position after the first rule. Because this rule will apply to any destination address, the rule that uses sanjose-network as the destination must come before this rule, or the sanjose-network rule will never be matched. The default is to place new manual NAT rules at the end of the "NAT Rules Before Auto NAT" section.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside-boulder.
- **Destination Interface Objects** = outside-boulder.

d) On **Translation**, configure the following:

- **Original Source** = boulder-network object.
- **Translated Source** = **Destination Interface IP**. This option configures interface PAT using the interface contained in the destination interface object.
- **Original Destination** > **Address** = any (leave blank).
- **Translated Destination** = any (leave blank).



e) Click **Save**.

Step 4 If you are also managing Firewall2 (San Jose), you can configure similar rules for that device.

- The manual identity NAT rule would be for sanjose-network when the destination is boulder-network. Create new interface objects for the Firewall2 inside and outside networks.
- The manual dynamic interface PAT rule would be for sanjose-network when the destination is "any."

Rewriting DNS Queries and Responses Using NAT

You might need to configure the FTD device to modify DNS replies by replacing the address in the reply with an address that matches the NAT configuration. You can configure DNS modification when you configure each translation rule. DNS modification is also known as DNS doctoring.

This feature rewrites the address in DNS queries and replies that match a NAT rule (for example, the A record for IPv4, the AAAA record for IPv6, or the PTR record for reverse DNS queries). For DNS replies traversing from a mapped interface to any other interface, the record is rewritten from the mapped value to the real value. Inversely, for DNS replies traversing from any interface to a mapped interface, the record is rewritten from the real value to the mapped value. This feature works with NAT44,NAT 66, NAT46, and NAT64.

Following are the main circumstances when you would need to configure DNS rewrite on a NAT rule.

- The rule is NAT64 or NAT46, and the DNS server is on the outside network. You need DNS rewrite to convert between DNS A records (for IPv4) and AAAA records (for IPv6).
- The DNS server is on the outside, clients are on the inside, and some of the fully-qualified domain names that the clients use resolve to other inside hosts.
- The DNS server is on the inside and responds with private IP addresses, clients are on the outside, and the clients access fully-qualified domain names that point to servers that are hosted on the inside.

DNS Rewrite Limitations

Following are some limitations with DNS rewrite:

- DNS rewrite is not applicable for PAT because multiple PAT rules are applicable for each A or AAAA record, and the PAT rule to use is ambiguous.

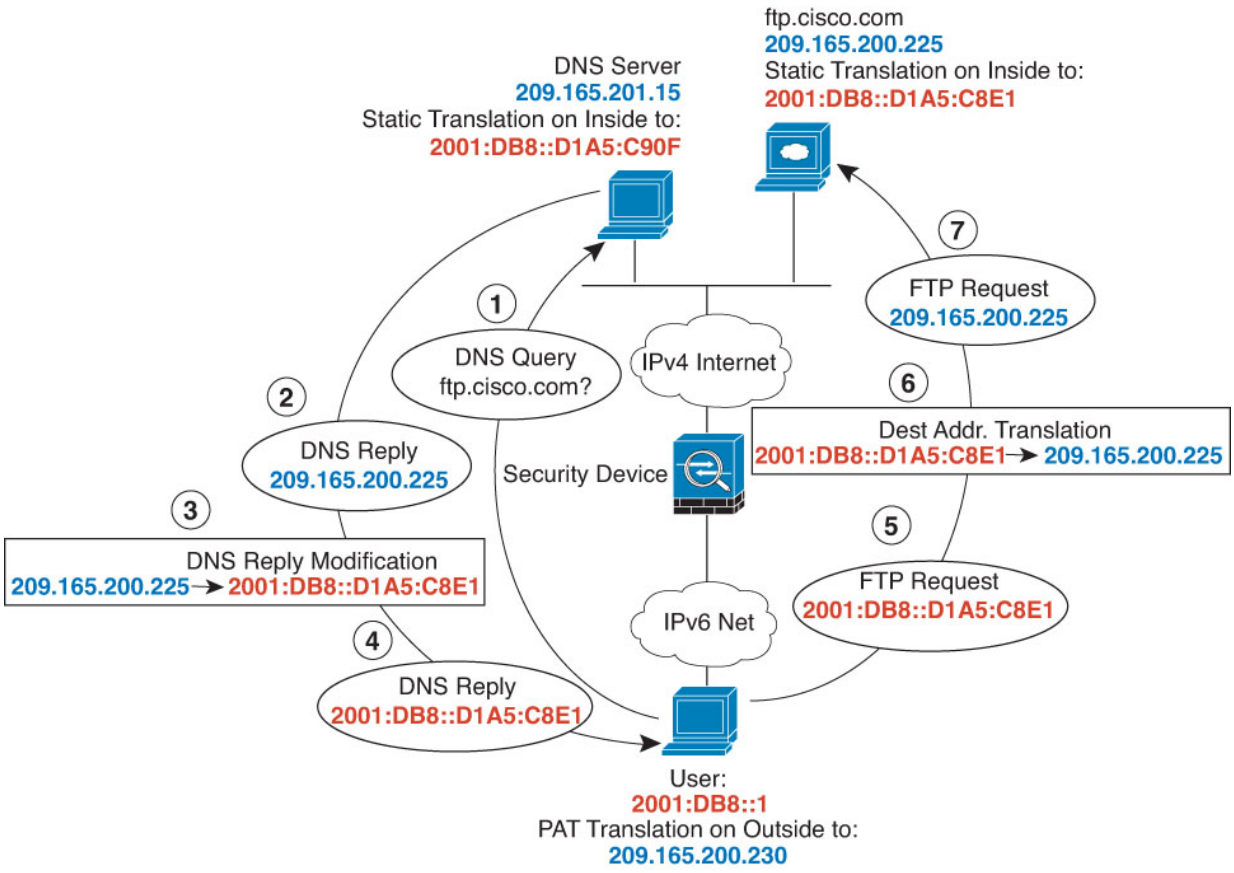
- If you configure a manual NAT rule, you cannot configure DNS modification if you specify the destination address as well as the source address. These kinds of rules can potentially have a different translation for a single address when going to A vs. B. Therefore, the can not accurately match the IP address inside the DNS reply to the correct twice NAT rule; the DNS reply does not contain information about which source/destination address combination was in the packet that prompted the DNS request.
- You must enable DNS application inspection with DNS NAT rewrite enabled for NAT rules to rewrite DNS queries and responses. By default, DNS inspection with DNS NAT rewrite enabled is globally applied, so you probably do not need to change the inspection configuration.
- DNS rewrite is actually done on the xlate entry, not the NAT rule. Thus, if there is no xlate for a dynamic rule, rewrite cannot be done correctly. The same problem does not occur for static NAT.
- DNS rewrite does not rewrite DNS Dynamic Update messages (opcode 5).

The following topics provide examples of DNS rewrite in NAT rules.

DNS64 Reply Modification

The following figure shows an FTP server and DNS server on the outside IPv4 network. The system has a static translation for the outside server. In this case, when an inside IPv6 user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.200.225.

Because you want inside users to use the mapped address for ftp.cisco.com (2001:DB8::D1A5:C8E1, where D1A5:C8E1 is the IPv6 equivalent of 209.165.200.225) you need to configure DNS reply modification for the static translation. This example also includes a static NAT translation for the DNS server, and a PAT rule for the inside IPv6 hosts.



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

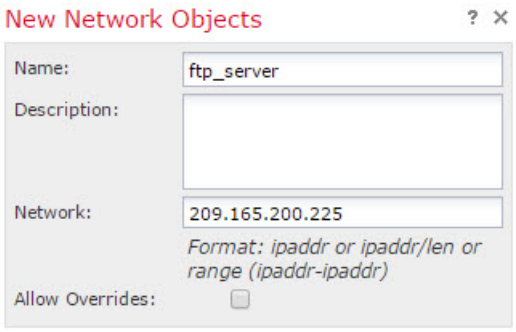
Step 1 Create the network objects for the FTP server, DNS server, inside network, and PAT pool.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the real FTP server address.

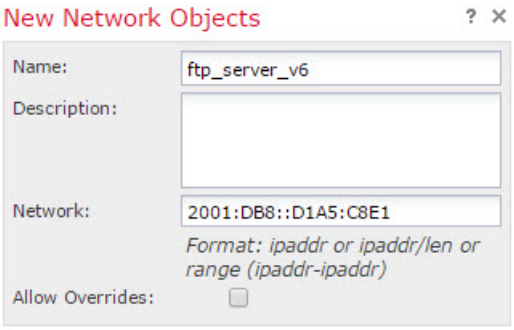
Name the network object (for example, ftp_server) and enter the host address, 209.165.200.225.



d) Click **Save**.

e) Click **Add Network** > **Add Object** and define the FTP server's translated IPv6 address.

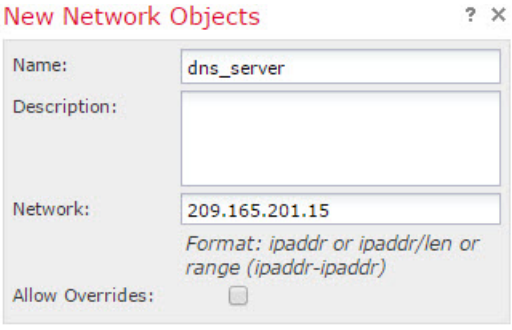
Name the network object (for example, ftp_server_v6) and enter the host address, 2001:DB8::D1A5:C8E1.



f) Click **Save**.

g) Click **Add Network** > **Add Object** and define the DNS server's real address.

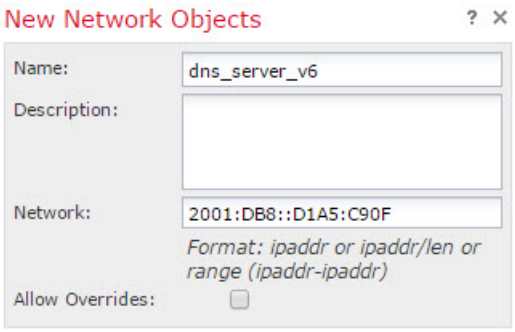
Name the network object (for example, dns_server) and enter the host address, 209.165.201.15.



h) Click **Save**.

i) Click **Add Network** > **Add Object** and define the DNS server's translated IPv6 address.

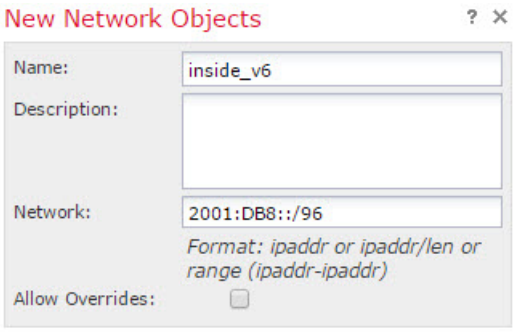
Name the network object (for example, dns_server_v6) and enter the host address, 2001:DB8::D1A5:C90F (where D1A5:C90F is the IPv6 equivalent of 209.165.201.15).



j) Click **Save**.

k) Click **Add Network** > **Add Object** and define the inside IPv6 network.

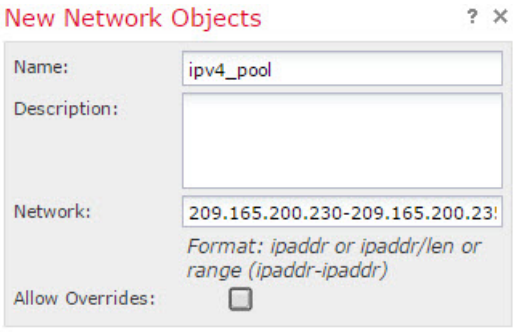
Name the network object (for example, inside_v6) and enter the network address, 2001:DB8::/96.



l) Click **Save**.

m) Click **Add Network** > **Add Object** and define the IPv4 PAT pool for the inside IPv6 network.

Name the network object (for example, ipv4_pool) and enter the range 209.165.200.230-209.165.200.235.



n) Click **Save**.

Step 2 Configure the static NAT rule with DNS modification for the FTP server.

a) Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

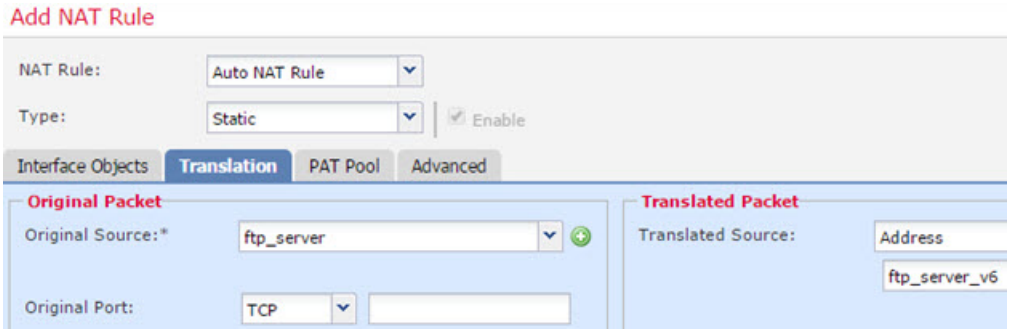
- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = outside.
- **Destination Interface Objects** = inside.

e) On **Translation**, configure the following:

- **Original Source** = ftp_server network object.
- **Translated Source** > **Address** = ftp_server_v6 network object.



f) On **Advanced**, select the following options:

- **Translate DNS replies that match this rule**.
- **Net to Net Mapping**, because this is a one-to-one NAT46 translation.

g) Click **OK**.

Step 3 Configure the static NAT rule for the DNS server.

a) Click **Add Rule**.

b) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

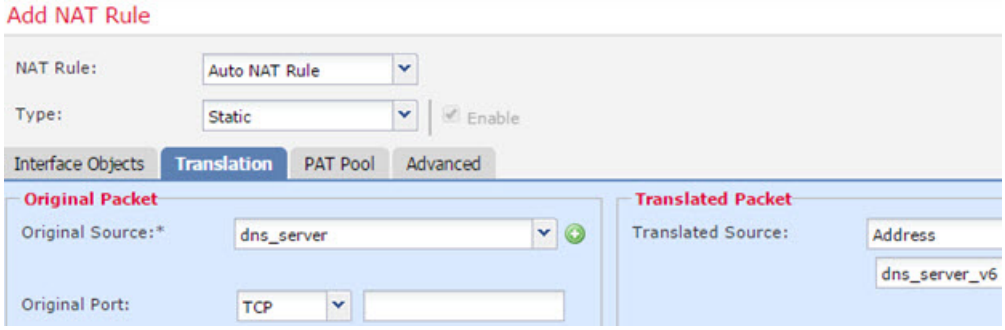
c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = outside.
- **Destination Interface Objects** = inside.

d) On **Translation**, configure the following:

- **Original Source** = dns_server network object.
- **Translated Source** > **Address** = dns_server_v6 network object.

e) On **Advanced**, select **Net to Net Mapping**, because this is a one-to-one NAT46 translation.



f) Click **OK**.

Step 4 Configure the dynamic NAT with a PAT pool rule for the inside IPv6 network.

a) Click **Add Rule**.

b) Configure the following properties:

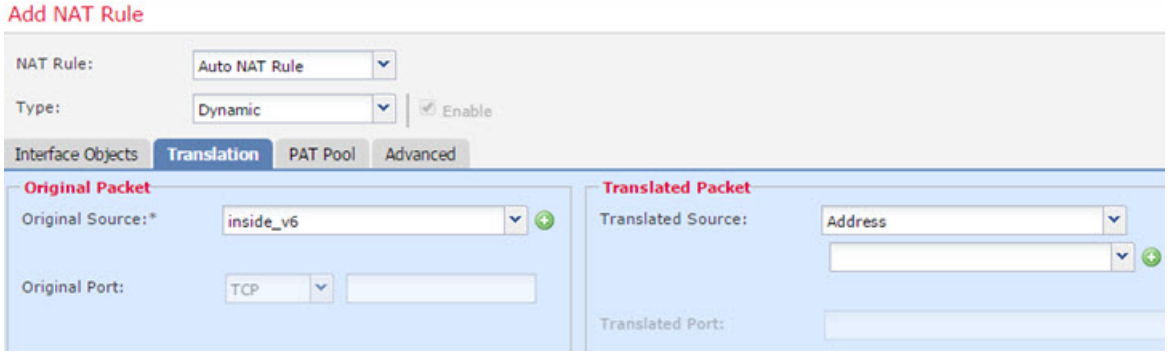
- **NAT Rule** = Auto NAT Rule.
- **Type** = Dynamic.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

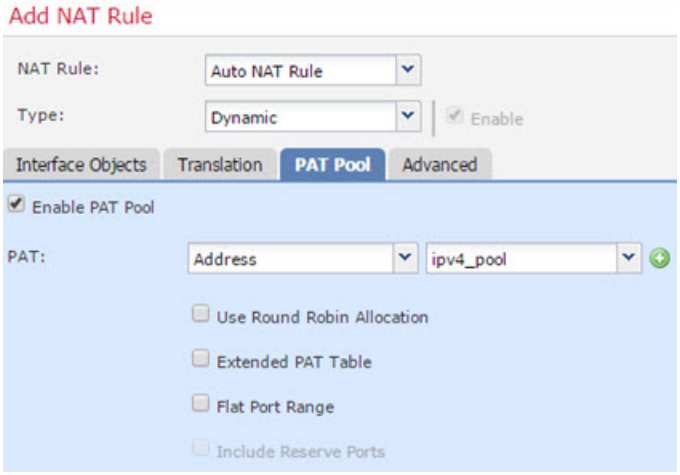
d) On **Translation**, configure the following:

- **Original Source** = inside_v6 network object.
- **Translated Source** > **Address** = leave this field empty.



e) On **PAT Pool**, configure the following:

- **Enable PAT Pool** = select this option.
- **Translated Source** > **Address** = ipv4_pool network object.



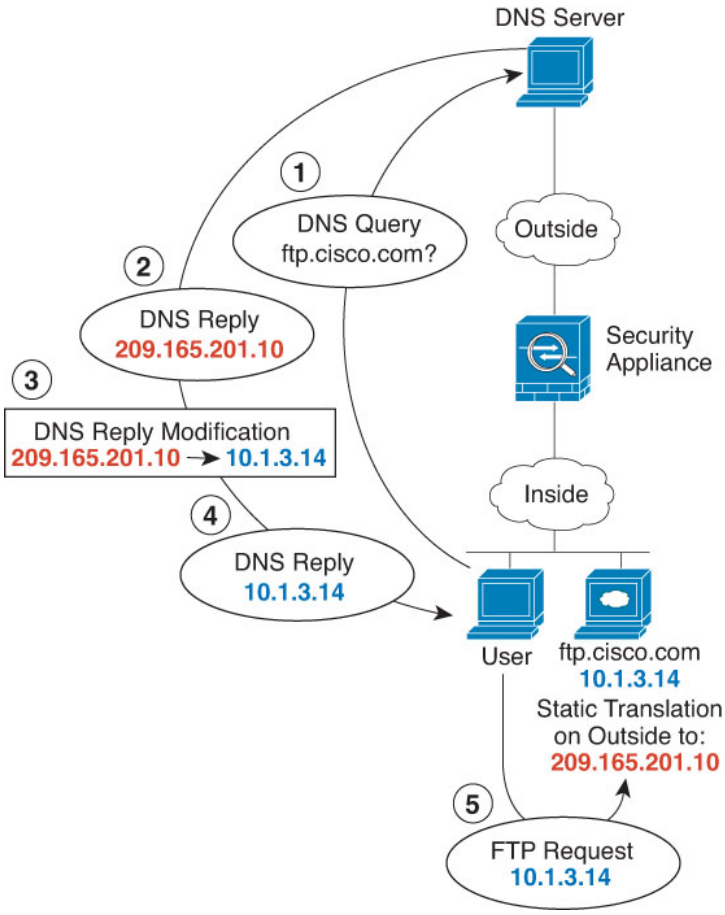
f) Click **OK**.

DNS Reply Modification, DNS Server on Outside

The following figure shows a DNS server that is accessible from the outside interface. A server, ftp.cisco.com, is on the inside interface. You configure NAT to statically translate the ftp.cisco.com real address (10.1.3.14) to a mapped address (209.165.201.10) that is visible on the outside network.

In this case, you want to enable DNS reply modification on this static rule so that inside users who have access to ftp.cisco.com using the real address receive the real address from the DNS server, and not the mapped address.

When an inside host sends a DNS request for the address of ftp.cisco.com, the DNS server replies with the mapped address (209.165.201.10). The system refers to the static rule for the inside server and translates the address inside the DNS reply to 10.1.3.14. If you do not enable DNS reply modification, then the inside host attempts to send traffic to 209.165.201.10 instead of accessing ftp.cisco.com directly.



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

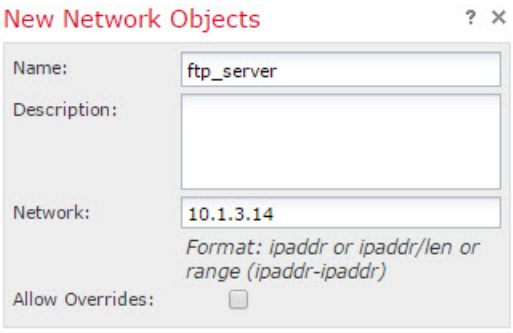
Step 1 Create the network objects for the FTP server.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the real FTP server address.

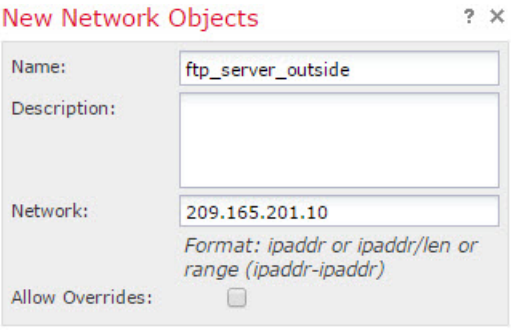
Name the network object (for example, ftp_server) and enter the host address, 10.1.3.14.



d) Click **Save**.

e) Click **Add Network** > **Add Object** and define the FTP server's translated address.

Name the network object (for example, ftp_server_outside) and enter the host address, 209.165.201.10.



f) Click **Save**.

Step 2 Configure the static NAT rule with DNS modification for the FTP server.

a) Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

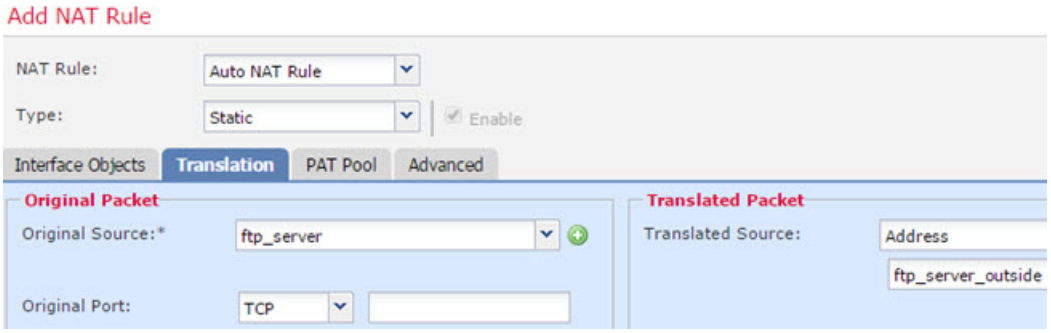
d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = ftp_server network object.
- **Translated Source** > **Address** = ftp_server_outside network object.

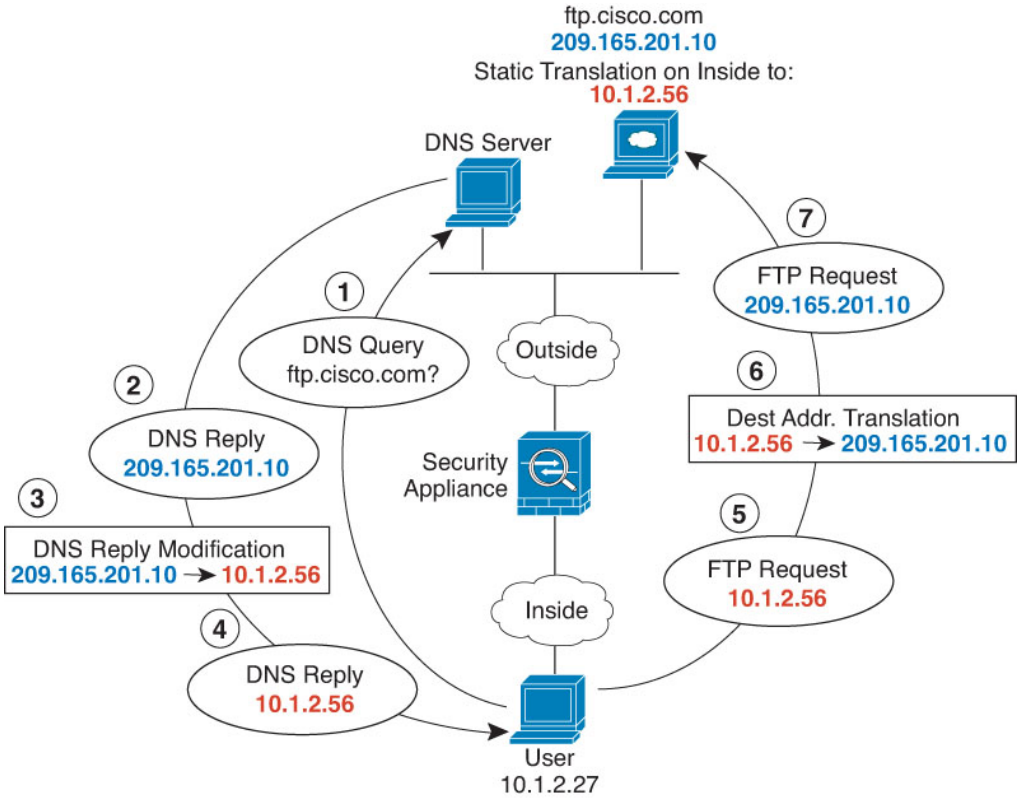
f) On **Advanced**, select **Translate DNS replies that match this rule**.



g) Click **OK**.

DNS Reply Modification, DNS Server on Host Network

The following figure shows an FTP server and DNS server on the outside. The system has a static translation for the outside server. In this case, when an inside user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.20.10. Because you want inside users to use the mapped address for ftp.cisco.com (10.1.2.56) you need to configure DNS reply modification for the static translation.



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

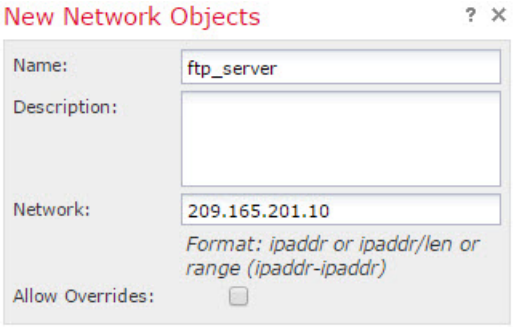
Step 1 Create the network objects for the FTP server.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the real FTP server address.

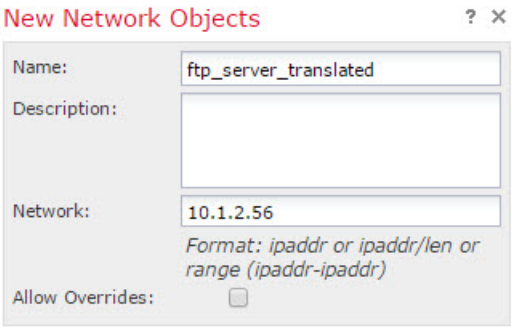
Name the network object (for example, ftp_server) and enter the host address, 209.165.201.10.



d) Click **Save**.

e) Click **Add Network** > **Add Object** and define the FTP server's translated address.

Name the network object (for example, ftp_server_translated) and enter the host address, 10.1.2.56.



f) Click **Save**.

Step 2 Configure the static NAT rule with DNS modification for the FTP server.

a) Select **Devices** > **NAT** and create or edit an Firepower Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

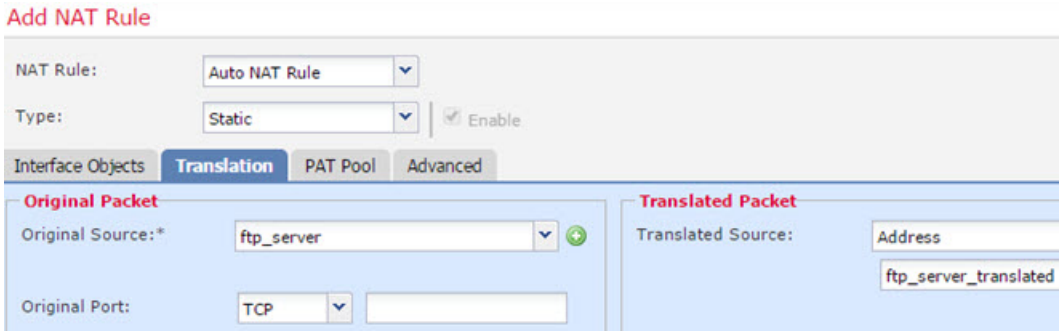
d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = outside.
- **Destination Interface Objects** = inside.

e) On **Translation**, configure the following:

- **Original Source** = ftp_server network object.
- **Translated Source** > **Address** = ftp_server_translated network object.

f) On **Advanced**, select **Translate DNS replies that match this rule**.



g) Click **OK**.

History for FTD NAT

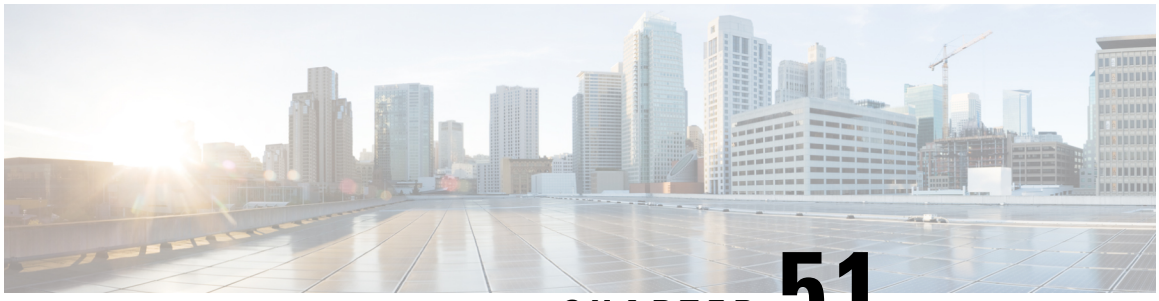
Feature	Version	Details
Network Address Translation (NAT) for Firepower Threat Defense.	6.0.1	The NAT policy for Firepower Threat Defense was added. New/modified screens: Threat Defense was added as a type of NAT policy to the **Devices** > **NAT** page. Supported platforms: Firepower Threat Defense
Support for network range objects in NAT for Firepower Threat Defense.	6.1.0	You can now use network range objects in Firepower Threat Defense NAT rules where appropriate.



PART XIV

7000 and 8000 Series Advanced Deployment Options

- Setting Up Virtual Switches, on page 1105
- Setting Up Virtual Routers, on page 1115
- Aggregate Interfaces and LACP, on page 1147
- Hybrid Interfaces, on page 1161
- Gateway VPNs, on page 1165



CHAPTER 51

Setting Up Virtual Switches

The following topics describe how to set up virtual switches in the Firepower System:

- Virtual Switches, on page 1105
- Switched Interface Configuration, on page 1105
- Virtual Switch Configuration, on page 1110

Virtual Switches

You can configure a 7000 or 8000 Series device in a Layer 2 deployment so that it provides packet switching between two or more networks. In a Layer 2 deployment, you can configure virtual switches to operate as standalone broadcast domains, dividing your network into logical segments. A virtual switch uses the media access control (MAC) address from a host to determine where to send packets.

When you configure a virtual switch, the switch initially broadcasts packets through every available port on the switch. Over time, the switch uses tagged return traffic to learn which hosts reside on the networks connected to each port.

A virtual switch must contain two or more switched interfaces to handle traffic. For each virtual switch, traffic becomes limited to the set of ports configured as switched interfaces. For example, if you configure a virtual switch with four switched interfaces, packets sent in through one port for broadcast can only be sent out of the remaining three ports on the switch.

When you configure a physical switched interface, you must assign it to a virtual switch. You can also define additional logical switched interfaces on a physical port as needed. You can group multiple physical interfaces into a single logical switched interface called a link aggregation group (LAG). This single aggregate logical link provides higher bandwidth, redundancy, and load-balancing between two endpoints.



Caution If a Layer 2 deployment fails for any reason, the device no longer passes traffic.

Switched Interface Configuration

You can set up switched interfaces to have either physical or logical configurations. You can configure physical switched interfaces for handling untagged VLAN traffic. You can also create logical switched interfaces for handling traffic with designated VLAN tags.

In a Layer 2 deployment, the system drops any traffic received on an external physical interface that does not have a switched interface waiting for it. If the system receives a packet with no VLAN tag and you have not configured a physical switched interface for that port, it drops the packet. If the system receives a VLAN-tagged packet and you have not configured a logical switched interface, it also drops the packet.

The system handles traffic that has been received with VLAN tags on switched interfaces by stripping the outermost VLAN tag on ingress before any rules evaluation or forwarding decisions. Packets leaving the device through a VLAN-tagged logical switched interface are encapsulated with the associated VLAN tag on egress.

Note that if you change the parent physical interface to inline or passive, the system deletes all the associated logical interfaces.

Switched Interface Configuration Notes

You can configure one or more physical ports on a managed device as switched interfaces. You must assign a physical switched interface to a virtual switch before it can handle traffic. You can configure link mode settings and MDI/MDIX settings only for copper interfaces.

Note Interfaces on 8000 Series appliances do not support half-duplex options.

For each physical switched interface, you can add multiple logical switched interfaces. You must associate each logical interface with a VLAN tag to handle traffic received by the physical interface with that specific tag. You must assign a logical switched interface to a virtual switch to handle traffic.

When configuring a switched interface, the range within which you can set the MTU can vary depending on the Firepower System device model and interface type.

The range of MTU values can vary depending on the model of the managed device and the interface type.

Caution Changing the highest MTU value among all non-management interfaces on the device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all non-management interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. See Snort® Restart Traffic Behavior, on page 350 for more information.

To edit an existing logical switched interface, click the edit icon ( ) next to the interface.

When you delete a logical switched interface, you remove it from the physical interface where it resides, as well as the virtual switch and security zone it is associated with.

Related Topics

MTU Ranges for 7000 and 8000 Series Devices and NGIPSv, on page 500
Snort® Restart Scenarios, on page 348

Configuring Physical Switched Interfaces

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to configure the switched interface, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Next to the interface you want to configure as a switched interface, click **Edit** ( ).

Step 4 Click the **Switched** tab.

Step 5 If you want to associate the switched interface with a security zone, do one of the following:

- Choose an existing security zone from the **Security Zone** drop-down list.
- Choose **New** to add a new security zone; see Creating Security Zone and Interface Group Objects, on page 411.

Step 6 If you want to associate the switched interface with a virtual switch, do one of the following:

- Choose an existing virtual switch from the **Virtual Switch** drop-down list.
- Choose **New** to add a new virtual switch; see Adding Virtual Switches, on page 1111.

Step 7 Check the **Enabled** check box to allow the switched interface to handle traffic.

Note If you clear the check box, the interface becomes disabled so that users cannot access it for security purposes.

Step 8 From the **Mode** drop-down list, choose an option to designate the link mode, or choose **Autonegotiation** to specify that the interface is configured to auto negotiate speed and duplex settings.

Mode settings are available only for copper interfaces.

Interfaces on 8000 Series appliances do not support half-duplex options.

Step 9 From the **MDI/MDIX** drop-down list, choose an option to designate whether the interface is configured for MDI (medium dependent interface), MDIX (medium dependent interface crossover), or Auto-MDIX.

By default, MDI/MDIX is set to Auto-MDIX, which automatically handles switching between MDI and MDIX to attain link.

Step 10 In the **MTU** field, enter a maximum transmission unit (MTU), which designates the largest size packet allowed.

The range of MTU values can vary depending on the model of the managed device and the interface type.

Caution Changing the highest MTU value among all non-management interfaces on the device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all non-management interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. See Snort® Restart Traffic Behavior, on page 350 for more information.

Step 11 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

MTU Ranges for 7000 and 8000 Series Devices and NGIPSv, on page 500
Snort® Restart Scenarios, on page 348

Adding Logical Switched Interfaces

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to add the switched interface, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Choose **Add Logical Interface** from the **Add** drop-down menu.

Step 4 Click **Switched**.

Step 5 From the **Interface** drop-down list, choose the physical interface that will receive the VLAN-tagged traffic.

Step 6 In the **VLAN Tag** field, enter a tag value that gets assigned to inbound and outbound traffic on this interface.

The tag value can be any integer from 1 to 4094.

Step 7 If you want to associate the switched interface with a security zone, do one of the following:

- Choose an existing security zone from the **Security Zone** drop-down list.
- Choose **New** to add a new security zone; see Creating Security Zone and Interface Group Objects, on page 411.

Step 8 If you want to associate the switched interface with a virtual switch, do one of the following:

- Choose an existing virtual switch from the **Virtual Switch** drop-down list.
- Choose **New** to add a new virtual switch; see Adding Virtual Switches, on page 1111.

Step 9 Check the **Enabled** check box to allow the switched interface to handle traffic.

If you clear the check box, the interface becomes disabled and administratively taken down. If you disable a physical interface, you also disable all of the logical interfaces associated with it.

Step 10 In the **MTU** field, enter a maximum transmission unit (MTU), which designates the largest size packet allowed.

The range of MTU values can vary depending on the model of the managed device and the interface type.

Caution Changing the highest MTU value among all non-management interfaces on the device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all non-management interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. See Snort® Restart Traffic Behavior, on page 350 for more information.

Step 11 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

MTU Ranges for 7000 and 8000 Series Devices and NGIPSv, on page 500
Snort® Restart Scenarios, on page 348

Deleting Logical Switched Interfaces

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the managed device that contains the switched interface you want to delete, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Next to the logical switched interface you want to delete, click the delete icon ( ).

Step 4 When prompted, confirm that you want to delete the interface.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Virtual Switch Configuration

Before you can use switched interfaces in a Layer 2 deployment, you must configure virtual switches and assign switched interfaces to them. A virtual switch is a group of switched interfaces that process inbound and outbound traffic through your network.

Virtual Switch Configuration Notes

You can add virtual switches from the Virtual Switches tab of the Device Management page. The Virtual Switches tab displays a list of all the virtual switches you have configured on a device. The page includes summary information about each switch.

Table 83: Virtual Switches Table View Fields

Field	Description
Name	The name of the virtual switch.
Interfaces	All switched interfaces that are assigned to the virtual switch. Interfaces that you have disabled from the Interfaces tab are not available.
Hybrid Interface	The optionally configured hybrid interface that ties the virtual switch to a virtual router.
Unicast Packets	Unicast packet statistics for the virtual switch, including: • Unicast packets received • Unicast packets forwarded (excludes drops by host) • Unicast packets unintentionally dropped
Broadcast Packets	Broadcast packet statistics for the virtual switch, including: • Broadcast packets received • Broadcast packets forwarded • Broadcast packets unintentionally dropped

You can also add switches as you configure switched interfaces. You can assign only switched interfaces to a virtual switch. If you want to create a virtual switch before you configure the switched interfaces on your managed devices, you can create an empty virtual switch and add interfaces to it later.

Tip To edit an existing virtual switch, click the edit icon ( ) next to the switch.

Adding Virtual Switches

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to add the virtual switch, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Virtual Switches** tab.

Step 4 Click **Add Virtual Switch**.

Step 5 Enter a name in the **Name** field.

Step 6 From the **Available** list, choose one or more switched interfaces to add to the virtual switch.

Tip Interfaces that you have disabled from the Interfaces tab are not available; disabling an interface after you add it removes it from the configuration.

Step 7 Click **Add**.

Step 8 If you want to tie the virtual switch to a virtual router, choose a hybrid interface from the **Hybrid Interface** drop-down list.

Step 9 Optionally, configure advanced settings for the switch; see Advanced Virtual Switch Settings, on page 1111

Step 10 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Logical Hybrid Interfaces, on page 1161

Advanced Virtual Switch Settings

Adding Static MAC Entries

Over time, a virtual switch learns MAC addresses by tagging return traffic from the network. You can manually add a static MAC entry, which designates that a MAC address resides on a specific port. Regardless of whether you ever receive traffic from that port, the MAC address remains static in the table. You can specify one or more static MAC addresses for each virtual switch.

Enabling Spanning Tree Protocol (STP) and Dropping Bridge Protocol Data Units (BPDU)

STP is a network protocol used to prevent network loops. BPDUs are exchanged through the network, carrying information about network bridges. The protocol uses BPDUs to identify and select the fastest network links, if there are redundant links in the network. If a network link fails, Spanning Tree fails over to an existing alternate link.

Note Cisco strongly recommends that you enable STP when configuring a virtual switch that you plan to deploy in a 7000 or 8000 Series device high-availability pair. Only enable STP if your virtual switch switches traffic between multiple network interfaces.

If your virtual switch routes traffic between VLANs, similar to a router on a stick, BPDUs enter and exit the device through different logical switched interfaces, but the same physical switched interface. As a result, STP identifies the device as a redundant network loop, which can cause issues in certain Layer 2 deployments. To prevent this, you can configure the virtual switch at the domain level to have the device drop BPDUs when monitoring traffic. You can only drop BPDUs if you disable STP.

Note Drop BPDUs only if your virtual swtich routes traffic between VLANs on a single physical interface.

Enabling Strict TCP Enforcement

To maximize TCP security, you can enable strict enforcement, which blocks connections where the three-way handshake was not completed. Strict enforcement also blocks:

- non-SYN TCP packets for connections where the three-way handshake was not completed
- non-SYN/RST packets from the initiator on a TCP connection before the responder sends the SYN-ACK
- non-SYN-ACK/RST packets from the responder on a TCP connection after the SYN but before the session is established
- SYN packets on an established TCP connection from either the initiator or the responder

Note that if you associate the virtual switch with a logical hybrid interface, the switch uses the same strict TCP enforcement setting as the virtual router associated with the logical hybrid interface. You cannot specify strict TCP enforcement on the switch in this case.

Configuring Advanced Virtual Switch Settings

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device that contains the virtual switch you want to edit, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Virtual Switches** tab.

Step 4 Next to the virtual switch that you want to edit, click the edit icon ( ).

Step 5 Click the **Advanced** tab.

Step 6 To add a static MAC entry, click **Add**.

Step 7 In the **MAC Address** field, enter the address using the standard format of six groups of two hexadecimal digits separated by colons (for example, 01:23:45:67:89:AB).

Note Broadcast addresses (00:00:00:00:00:00 and FF:FF:FF:FF:FF:FF) cannot be added as static MAC addresses.

Step 8 From the **Interface** drop-down list, choose the interface where you want to assign the MAC address.

Step 9 Click **OK**.

Step 10 If you want to enable the Spanning Tree Protocol, check the **Enable Spanning Tree Protocol** check box.

Step 11 If you want to enable strict TCP enforcement, check the **Strict TCP Enforcement** check box.

If you associate the virtual switch with a logical hybrid interface, this option does not appear and the switch uses the same setting as the virtual router associated with the logical hybrid interface.

Step 12 If you want to drop BPDUs at the domain level, check the **Drop BPDUs** check box.

Step 13 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Deleting Virtual Switches

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

When you delete a virtual switch, any switched interfaces assigned to the switch become available for inclusion in another switch.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the managed device that contains the virtual switch you want to delete, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

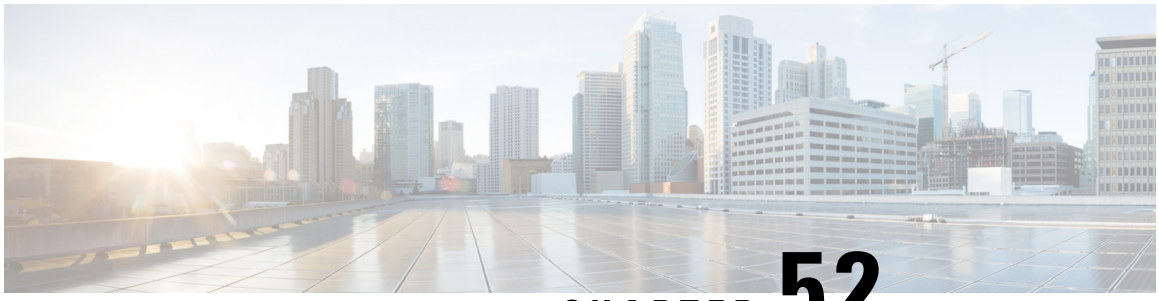
Step 3 Click the **Virtual Switches** tab.

Step 4 Next to the virtual switch that you want to delete, click the delete icon ( ).

Step 5 When prompted, confirm that you want to delete the virtual switch.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.



CHAPTER 52

Setting Up Virtual Routers

The following topics describe how to set up virtual routers in the Firepower System:

- Virtual Routers, on page 1115
- Routed Interfaces, on page 1116
- Configuring Physical Routed Interfaces, on page 1117
- Adding Logical Routed Interfaces, on page 1119
- Deleting Logical Routed Interfaces, on page 1121
- Configuring SFRP, on page 1122
- Virtual Router Configuration, on page 1123
- Adding Virtual Routers, on page 1124
- DHCP Relay, on page 1125
- Static Routes, on page 1127
- Dynamic Routing, on page 1129
- Virtual Router Filters, on page 1141
- Adding Virtual Router Authentication Profiles, on page 1143
- Viewing Virtual Router Statistics, on page 1144
- Deleting Virtual Routers, on page 1145

Virtual Routers

You can configure a managed device in a Layer 3 deployment so that it routes traffic between two or more interfaces. To route traffic, you must assign an IP address to each interface and assign the interfaces to the virtual router. The interfaces assigned to virtual routers can be physical, logical, or link aggregation group (LAG) interfaces.

You can configure the system to route packets by making packet forwarding decisions according to the destination address. Interfaces configured as routed interfaces receive and forward the Layer 3 traffic. Routers obtain the destination from the outgoing interface based on the forwarding criteria, and access control rules designate the security policies to be applied.

In Layer 3 deployments, you can define static routes. In addition, you can configure Routing Information Protocol (RIP) and Open Shortest Path First (OSPF) dynamic routing protocols. You can also configure a combination of static routes and RIP or static routes and OSPF.

Note that you can only configure virtual routers, physical routed interfaces, or logical routed interfaces on a 7000 or 8000 Series device.

Caution If a Layer 3 deployment fails for any reason, the device no longer passes traffic.

Related Topics

LAG Configuration, on page 1148

Routed Interfaces

You can set up routed interfaces with either physical or logical configurations. You can configure physical routed interfaces for handling untagged VLAN traffic. You can also create logical routed interfaces for handling traffic with designated VLAN tags.

In a Layer 3 deployment, the system drops any traffic received on an external physical interface that does not have a routed interface waiting for it. The system drops a packet if:

- It receives a packet with no VLAN tag, and you have not configured a physical routed interface for that port.
- It receives a VLAN-tagged packet, and you have not configured a logical routed interface for that port.

The system handles traffic that has been received with VLAN tags on switched interfaces by stripping the outermost VLAN tag on ingress prior to any rules evaluation or forwarding decisions. Packets leaving the device through a VLAN-tagged logical routed interface are encapsulated with the associated VLAN tag on egress. The system drops any traffic received with a VLAN tag after the stripping process completes.

You can add static Address Resolution Protocol (ARP) entries to a routed interface. If an external host needs to know the MAC address of the destination IP address it needs to send traffic to on your local network, it sends an ARP request. When you configure static ARP entries, the virtual router responds with an IP address and associated MAC address.

Note that disabling the **ICMP Enable Responses** option for logical routed interfaces does not prevent ICMP responses in all scenarios. You can add network-based rules to an access control policy to drop packets where the destination IP is the routed interface's IP and the protocol is ICMP.

If you have enabled the **Inspect Local Router Traffic** option on the managed device, the system drops the packets before they reach the host, thereby preventing any response.

The range of MTU values can vary depending on the model of the managed device and the interface type.

Caution Changing the highest MTU value among all non-management interfaces on the device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all non-management interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. See Snort® Restart Traffic Behavior, on page 350 for more information.

If you change the parent physical interface to inline or passive, the system deletes all the associated logical interfaces.

Related Topics

Advanced Settings, on page 251
MTU Ranges for 7000 and 8000 Series Devices and NGIPSv, on page 500

Snort® Restart Scenarios, on page 348

Configuring Physical Routed Interfaces

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

You can configure one or more physical ports on a managed device as routed interfaces. You must assign a physical routed interface to a virtual router before it can route traffic.

Caution Adding a routed interface pair on a 7000 or 8000 Series device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Next to the interface you want to modify, click **Edit** ( ).

Step 4 Click **Routed** to display the routed interface options.

Step 5 If you want to apply a security zone, do one of the following:

- Choose an existing security zone from the **Security Zone** drop-down list.
- Choose **New** to add a new security zone; see Creating Security Zone and Interface Group Objects, on page 411.

Step 6 If you want to specify a virtual router, do one of the following:

- Choose an existing virtual router from the **Virtual Router** drop-down list.
- Choose **New** to add a new virtual router; Adding Virtual Routers, on page 1124.

Step 7 Check the **Enabled** check box to allow the routed interface to handle traffic. If you clear the check box, the interface becomes disabled so that users cannot access it for security purposes.

Step 8 From the **Mode** drop-down list, choose an option to designate the link mode, or choose **Autonegotiation** to specify that the interface is configured to auto negotiate speed and duplex settings.

Mode settings are available only for copper interfaces.

Interfaces on 8000 Series appliances do not support half-duplex options.

Step 9 From the **MDI/MDIX** drop-down list, choose an option to designate whether the interface is configured for MDI (medium dependent interface), MDIX (medium dependent interface crossover), or Auto-MDIX.

Normally, MDI/MDIX is set to Auto-MDIX, which automatically handles switching between MDI and MDIX to attain link.

MDI/MDIX settings are available only for copper interfaces.

Step 10 In the **MTU** field, choose a maximum transmission unit (MTU), which designates the largest size packet allowed.

The MTU is the Layer 2 MTU/MRU and not the Layer 3 MTU.

The range of MTU values can vary depending on the model of the managed device and the interface type.

Caution Changing the highest MTU value among all non-management interfaces on the device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all non-management interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. See Snort® Restart Traffic Behavior, on page 350 for more information.

Step 11 Next to **ICMP**, check the **Enable Responses** check box to allow the interface to respond to ICMP traffic such as pings and traceroute.

Step 12 Next to **IPv6 NDP**, check the **Enable Router Advertisement** check box to enable the interface to broadcast router advertisements.

Step 13 To add an IP address, click **Add**.

Step 14 In the **Address** field, enter the routed interface's IP address and subnet mask using CIDR notation.

Note the following:

- You cannot add network and broadcast addresses, or the static MAC addresses 00:00:00:00:00:00 and FF:FF:FF:FF:FF:FF.
- You cannot add identical IP addresses, regardless of subnet mask, to interfaces in virtual routers.

Step 15 If your organization uses IPv6 addresses and you want to set the IP address of the interface automatically, check the **Address Autoconfiguration** check box next to the **IPv6** field.

Step 16 For **Type**, choose either **Normal** or **SFRP**.

For SFRP options, see Configuring SFRP, on page 1122 for more information.

Step 17 Click **OK**.

- To edit an IP address, click **Edit** ( ).
- To delete an IP address, click **Delete** ( ).

When adding an IP address to a routed interface of a 7000 or 8000 Series device in a high-availability pair, you must add a corresponding IP address to the routed interface on the high-availability pair peer.

Step 18 To add a static ARP entry, click **Add**.

Step 19 In the **IP Address** field, enter an IP address for the static ARP entry.

Step 20 In the **MAC Address** field, enter a MAC address to associate with the IP address. Use the standard address format of six groups of two hexadecimal digits separated by colons (for example, 01:23:45:67:89:AB).

Step 21 Click **OK**.

Step 22 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

MTU Ranges for 7000 and 8000 Series Devices and NGIPSv, on page 500
Snort® Restart Scenarios, on page 348

Adding Logical Routed Interfaces

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

For each physical routed interface, you can add multiple logical routed interfaces. You must associate each logical interface with a VLAN tag to handle traffic received by the physical interface with that specific tag. You must assign a logical routed interface to a virtual router to route traffic.

Caution Adding a routed interface pair on 7000 or 8000 Series devices restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Add Interface**.

Step 4 Click **Routed** to display the routed interface options.

Step 5 From the **Interface** drop-down list, choose the physical interface where you want to add the logical interface.

Step 6 In the **VLAN Tag** field, enter a tag value that gets assigned to inbound and outbound traffic on this interface. The value can be any integer from 1 to 4094.

Step 7 If you want to apply a security zone, do one of the following:

- Choose an existing security zone from the **Security Zone** drop-down list.

- Choose **New** to add a new security zone; see Creating Security Zone and Interface Group Objects, on page 411.

Step 8 If you want to specify a virtual router, do one of the following:

- Choose an existing virtual router from the **Virtual Router** drop-down list.
- Choose **New** to add a new virtual router; Adding Virtual Routers, on page 1124.

Step 9 Check the **Enabled** check box to allow the routed interface to handle traffic.

If you clear the check box, the interface becomes disabled and administratively taken down. If you disable a physical interface, you also disable all of the logical interfaces associated with it.

Step 10 In the **MTU** field, enter a maximum transmission unit (MTU), which designates the largest size packet allowed.

The MTU is the Layer 2 MTU/MRU and not the Layer 3 MTU.

The range of MTU values can vary depending on the model of the managed device and the interface type.

Caution Changing the highest MTU value among all non-management interfaces on the device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all non-management interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. See Snort® Restart Traffic Behavior, on page 350 for more information.

Step 11 Next to **ICMP**, check the **Enable Responses** check box to communicate updates or error information to other routers, intermediary devices, or hosts.

Step 12 Next to **IPv6 NDP**, check the **Enable Router Advertisement** check box to enable the interface to broadcast router advertisements.

Step 13 To add an IP address, click **Add**.

Step 14 In the **Address** field, enter the IP address in CIDR notation.

Note the following:

- You cannot add network and broadcast addresses, or the static MAC addresses 00:00:00:00:00:00 and FF:FF:FF:FF:FF:FF.
- You cannot add identical IP addresses, regardless of subnet mask, to interfaces in virtual routers.

Step 15 If your organization uses IPv6 addresses and you want to set the IP address of the interface automatically, choose the **Address Autoconfiguration** check box next to the **IPv6** field.

Step 16 For **Type**, choose either **Normal** or **SFRP**.

For SFRP options, see Configuring SFRP, on page 1122 for more information.

Step 17 Click **OK**.

- To edit an IP address, click the edit icon ( ).
- To delete an IP address, click the delete icon ( ).

When you add an IP address to a routed interface of a 7000 or 8000 Series device in a high-availability pair, you must add a corresponding IP address to the routed interface on the high-availability pair peer.

Step 18 To add a static ARP entry, click **Add**.

Step 19 In the **IP Address** field, enter an IP address for the static ARP entry.

Step 20 In the **MAC Address** field, enter a MAC address to associate with the IP address. Use the standard address format of six groups of two hexadecimal digits separated by colons (for example, 01:23:45:67:89:AB).

Step 21 Click **OK**. The static ARP entry is added.

Step 22 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

MTU Ranges for 7000 and 8000 Series Devices and NGIPSv, on page 500
Snort® Restart Scenarios, on page 348

Deleting Logical Routed Interfaces

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

When you delete a logical routed interface, you remove it from the physical interface where it resides, as well as its assigned virtual router and security zone.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click the edit icon **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Next to the logical routed interface you want to delete, click **Delete** ( ).

Step 4 When prompted, confirm that you want to delete the interface.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configuring SFRP

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

You can configure Cisco Redundancy Protocol (SFRP) to achieve network redundancy for high availability on either a 7000 or 8000 Series device high-availability pair or individual devices. SFRP provides gateway redundancy for both IPv4 and IPv6 addresses. You can configure SFRP on routed and hybrid interfaces.

If the interfaces are configured on individual devices, they must be in the same broadcast domain. You must designate at least one of the interfaces as primary and an equal number as backup. The system supports only one primary and one backup per IP address. If network connectivity is lost, the system automatically promotes the backup to primary to maintain connectivity.

The options you set for SFRP must be the same on all interfaces in a group of SFRP interfaces. Multiple IP addresses in a group must be in the same primary/backup state. Therefore, when you add or edit an IP address, the state you set for that address propagates to all the addresses in the group. For security purposes, you must enter values for **Group ID** and **Shared Secret** that are shared among the interfaces in the group.

To enable SFRP IP addresses on a virtual router, you must also configure one non-SFRP IP address. Note that only one non-SFRP address should be configured per interface.

As all SFRPs in a group failover together, all SFRPs on the same virtual router should be in the same SFRP group. In addition, you should also set up an HA link interface on each device in a high-availability pair when using NAT, HA state sharing, or VPN. For more information on HA link interfaces, see Configuring HA Link Interfaces, on page 498.

For 7000 or 8000 Series devices in a high-availability pair, you designate the shared secret and the system copies it to the high-availability pair peer along with the SFRP IP configuration. The shared secret authenticates peer data.

Note We do not recommend enabling more than one non-SFRP IP address on a 7000 or 8000 Series device high-availability pair's routed or hybrid interface where one SFRP IP address is already configured. The system does not perform NAT if a 7000 or 8000 Series device high-availability pair fails over while in standby mode.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Next to the interface where you want to configure SFRP, click **Edit** ( ).

Step 4 Choose the type of interface where you want to configure SFRP, either **Routed** or **Hybrid**.

Step 5 You can configure SFRP while adding or editing an IP address. Click **Add** to add an IP address. To edit an IP address, click **Edit** ( ).

Step 6 For **Type**, choose **SFRP** to display the SFRP options.

Step 7 In the **Group ID** field, enter a value that designates a group of primary or backup interfaces configured for SFRP.

Step 8 For **Priority**, choose either priamry or backup to designate the preferred interface:

- For individual devices, you must set one interface to primary on one device and the other to backup on a second device.
- For 7000 or 8000 Seriesdevice high-availability pairs, when you set one interface as primary, the other automatically becomes the backup.

Step 9 In the **Shared Secret** field, enter a shared secret.

The Shared Secret field populates automatically for a group in a 7000 or 8000 Series device high-availability pair.

Step 10 In the **Adv. Interval (seconds)** field, enter an interval for route advertisements for Layer 3 traffic.

Step 11 Click **OK**.

Step 12 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

About 7000 and 8000 Series Device High Availability, on page 517

Virtual Router Configuration

Caution Adding a virtual router on a 7000 or 8000 Series device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Before you can use routed interfaces in a Layer 3 deployment, you must configure virtual routers and assign routed interfaces to them. A virtual router is a group of routed interfaces that route Layer 3 traffic.

You can assign only routed and hybrid interfaces to a virtual router.

To maximize TCP security, you can enable strict enforcement, which blocks connections where the three-way handshake was not completed. Strict enforcement also blocks:

- non-SYN TCP packets for connections where the three-way handshake was not completed
- non-SYN/RST packets from the initiator on a TCP connection before the responder sends the SYN-ACK

- non-SYN-ACK/RST packets from the responder on a TCP connection after the SYN but before the session is established
- SYN packets on an established TCP connection from either the initiator or the responder

Note that if you change the configuration of a Layer 3 interface to a non-Layer 3 interface or remove a Layer 3 interface from the virtual router, the router may fall into an invalid state. For example, if it is used in DHCPv6, it may cause an upstream and downstream mismatch.

Adding Virtual Routers

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

You can add virtual routers from the **Virtual Routers** tab of the device management page. You can also add routers as you configure routed interfaces.

If you want to create a virtual router before you configure the interfaces on your managed devices, you can create an empty virtual router and add interfaces to it later.

Caution Adding a virtual router on a 7000 or 8000 Series device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Virtual Routers** tab.

Tip If your devices are in a stack in a high-availability pair, choose the stack you want to modify from the **Selected Device** drop-down list.

Step 4 Click **Add Virtual Router**.

Step 5 In the **Name** field, enter a name for the virtual router. You can use alphanumeric characters and spaces.

Step 6 Configure IPv6 static routing, OSPFv3, and RIPng on your virtual router by checking or clearing the **IPv6 Support** check box.

Step 7 If you do not want to enable strict TCP enforcement, clear the **Strict TCP Enforcement** check box. This option is enabled by default.

Step 8 Choose one or more interfaces from the **Available** list under **Interfaces**, and click **Add**.

The **Available** list contains all enabled Layer 3 interfaces, routed and hybrid, on the device that you can assign to the virtual router.

Tip To remove a routed or hybrid interface from the virtual router, click **Delete** ( ). Disabling a configured interface from the Interfaces tab also removes it.

Step 9 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

DHCP Relay

DHCP provides configuration parameters to Internet hosts. A DHCP client that has not yet acquired an IP address cannot communicate directly with a DHCP server outside its broadcast domain. To allow DHCP clients to communicate with DHCP servers, you can configure DHCP relay instances to handle cases where the client is not on the same broadcast domain as the server.

You can set up DHCP relay for each virtual router you configure. By default, this feature is disabled. You can enable either DHCPv4 relay or DHCPv6 relay.

Note You cannot run a DHCPv6 Relay chain through two or more virtual routers running on the same device.

Setting Up DHCPv4 Relay

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

The following procedure explains how to set up DHCPv4 relay on a virtual router.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Virtual Routers** tab.

Step 4 Next to the virtual router you want to modify, click the edit icon ( ).

Step 5 Check the **DHCPv4** check box.

Step 6 Under the **Servers** field, enter a server IP address.

Step 7 Click **Add**.

You can add up to four DHCP servers.

Step 8 In the **Max Hops** field, enter the maximum number of hops from 1 to 255.

Step 9 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Setting Up DHCPv6 Relay

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

You cannot run a DHCPv6 Relay chain through two or more virtual routers running on the same device.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Virtual Routers** tab.

Step 4 Next to the virtual router where you want to set up DHCP relay, click the edit icon ( ).

Step 5 Check the **DHCPv6** check box.

Step 6 In the **Interfaces** field, check the check boxes next to one or more interfaces that have been assigned to the virtual router.

Tip You cannot disable an interface from the **Interfaces** tab while it is configured for DHCPv6 Relay. You must first clear the DHCPv6 Relay interfaces check box and save the configuration.

Step 7 Next to a selected interface, click the drop-down icon and choose whether the interface relays DHCP requests **Upstream**, **Downstream**, or **Both**.

Note You must include at least one downstream interface and one upstream interface. Choosing both means that the interface is both downstream and upstream.

Step 8 In the **Max Hops** field, enter the maximum number of hops from 1 to 255

Step 9 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Static Routes

Static routing allows you to write rules about the IP addresses of traffic passing through a router. It is the simplest way of configuring path selection of a virtual router because there is no communication with other routers regarding the current topology of the network.

Do not configure routing to IP interfaces for DHCPv4 servers that the assigned virtual router cannot route packets to. Doing so will render previously specified routable DHCP4 servers unroutable.

The Static Routes table includes summary information about each route, as described in the following table.

Table 84: Static Routes Table View Fields

Field	Description
Enabled	Specifies whether this route is currently enabled or disabled.
Name	The name of the static route.
Destination	The destination network where traffic is routed.
Type	Specifies the action that is taken for this route, which will is one of the following: • IP — designates that the route forwards packets to the address of a neighboring router. • Interface — designates that the route forwards packets to an interface through which traffic is routed to hosts on a directly connected network. • Discard — designates that the static route drops packets.
Gateway	The target IP address if you selected IP as the static route type or the interface if you selected Interface as the static route type.
Preference	Determines the route selection. If you have multiple routes to the same destination, the system selects the route with the higher preference.

Viewing the Static Routes Table

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Any	Admin/Network Admin

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to view, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Virtual Routers** tab.

Step 4 Next to the virtual router where you want to view static routes, click the edit icon ( ).

If a view icon ( ) appears instead, the configuration belongs to a descendant domain, or you do not have permission to modify the configuration.

Step 5 Click the **Static** tab.

Adding Static Routes

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to add the static route, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Virtual Routers** tab.

Step 4 Next to the virtual router where you want to add the static route, click the edit icon ( ).

Step 5 Click **Static** to display the static route options.

Step 6 Click **Add Static Route**.

Step 7 In the **Route Name** field, enter a name for the static route. You can use alphanumeric characters and spaces.

Step 8 For **Enabled**, check the check box to specify that the route is currently enabled.

Step 9 In the **Preference** field, enter a numerical value between 1 and 65535 to determine the route selection.

Note If you have multiple routes to the same destination, the system uses the route with the higher preference.

Step 10 From the **Type** drop-down list, choose the type of static route you are configuring.

Step 11 In the **Destination** field, enter the IP address for the destination network where traffic should be routed.

Step 12 In the **Gateway** field, you have two options:

- If you chose **IP** as the selected static route type, choose an IP address.
- If you chose **Interface** as the selected static route type, choose an enabled interface from the drop-down list.

Tip Interfaces you have disabled from the **Interfaces** tab are not available; disabling an interface you have added removes it from the configuration.

Step 13 Click **OK**.

Step 14 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Dynamic Routing

Dynamic, or adaptive, routing uses a routing protocol to alter the path that a route takes in response to a change in network conditions. The adaptation is intended to allow as many routes as possible to remain valid, that is, have destinations that can be reached in response to the change. This allows the network to "route around" damage, such as loss of a node or a connection between nodes, so long as other path choices are available. You can configure a router with no dynamic routing, or you can configure the Routing Information Protocol (RIP) or the Open Shortest Path First (OSPF) routing protocol.

RIP Configuration

Routing Information Protocol (RIP) is a dynamic routing protocol, designed for small IP networks, that relies on hop count to determine routes. The best routes use the fewest number of hops. The maximum number of hops allowed for RIP is 15. This hop limit also limits the size of the network that RIP can support.

Adding Interfaces for RIP Configuration

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

While configuring RIP, you must choose interfaces from those already included in the virtual router, where you want to configure RIP. Disabled interfaces are not available.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Virtual Routers** tab.

Step 4 Next to the virtual router you want to modify, click the edit icon ( ).

Step 5 Click **Dynamic Routing** to display the dynamic routing options.

Step 6 Click **RIP** to display the RIP options.

Step 7 Under **Interfaces**, click the add icon ( ).

Step 8 From the **Name** drop-down list, choose the interface where you want to configure RIP.

Tip Interfaces you have disabled from the Interfaces tab are not available; disabling an interface you have added removes it from the configuration.

Step 9 In the **Metric** field, enter a metric for the interface. When routes from different RIP instances are available and all of them have the same preference, the route with the lowest metric becomes the preferred route.

Step 10 From the **Mode** drop-down list, choose one of the following options:

- **Multicast** — default mode where RIP multicasts the entire routing table to all adjacent routers at a specified address.
- **Broadcast** — forces RIP to use broadcast (for example, RIPv1) even though multicast mode is possible.
- **Quiet** — RIP will not transmit any periodic messages to this interface.
- **No Listen** — RIP will send to this interface but not listen to it.

Step 11 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configuring Authentication Settings for RIP Configuration

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

RIP authentication uses one of the authentication profiles you configured on the virtual router.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Virtual Routers** tab.

Step 4 Next to the virtual router where you want to add the RIP authentication profile, click the edit icon ( ).

Step 5 Click **Dynamic Routing** to display the dynamic routing options.

Step 6 Click **RIP** to display the RIP options.

Step 7 Under **Authentication**, choose an existing virtual router authentication profile from the **Profile** drop-down list, or choose **None**.

Step 8 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configuring Advanced Settings for RIP Configuration

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

You can configure several advanced RIP settings pertaining to various timeout values and other features that affect the behavior of the protocol.

Caution Changing any of the advanced RIP settings to incorrect values can prevent the router from communicating successfully with other RIP routers.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Virtual Routers** tab.

Step 4 Next to the virtual router you want to modify, click the edit icon ( ).

Step 5 Click **Dynamic Routing** to display the dynamic routing options.

Step 6 Click **RIP** to display the RIP options.

Step 7 In the **Preference** field, enter a numerical value (higher is better) for the preference of the routing protocol. The system prefers routes learned through RIP over static routes.

Step 8 In the **Period** field, enter the interval, in seconds, between periodic updates. A lower number determines faster convergence, but larger network load.

Step 9 In the **Timeout Time** field, enter a numerical value that specifies how old routes must be, in seconds, before being considered unreachable.

Step 10 In the **Garbage Time** field, enter a numerical value that specifies how old routes must be, in seconds, before being discarded.

Step 11 In the **Infinity** field, enter a numerical value that specifies a value for infinity distance in convergence calculations. Larger values will make protocol convergence slower.

Step 12 From the **Honor** drop-down list, choose one of the following options to designate when requests for dumping routing tables should be honored:

- **Always** — always honor requests
- **Neighbor** — only honor requests sent from a host on a directly connected network
- **Never** — never honor requests

Step 13 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Adding Import Filters for RIP Configuration

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

You can add an import filter to designate which routes are accepted or rejected from RIP into the route table. Import filters are applied in the order they appear in the table.

When adding an import filter, you use one of the filters you configured on the virtual router.

Tip To edit a RIP import filter, click the edit icon ( ). To delete a RIP import filter, click the delete icon ( ).

Before you begin

- Add a virtual router as described in Adding Virtual Routers, on page 1124.
- Configure a filter on the virtual router as described in Setting Up Virtual Router Filters, on page 1142.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Virtual Routers** tab.

Step 4 Next to the virtual router where you want to add the RIP virtual router filter, click the edit icon ( ).

Step 5 Click **Dynamic Routing** to display the dynamic routing options.

Step 6 Click **RIP** to display the RIP options.

Step 7 Under **Import Filters**, click the add icon ( ).

Step 8 From the **Name** drop-down list, choose the filter you want to add as an import filter.

Step 9 Next to **Action**, choose **Accept** or **Reject**.

Step 10 Click **OK**.

Tip To change the order of the import filters, click the move up ( ) and move down ( ) icons as needed. You can also drag the filters up or down in the list.

Step 11 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Adding Export Filters for RIP Configuration

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

You can add an export filter to define which routes will be accepted or rejected from the route table to RIP. Export filters are applied in the order they appear in the table.

When adding an export filter, you use one of the filters you configured on the virtual router.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Virtual Routers** tab.

Step 4 Next to the virtual router where you want to add the RIP virtual router filter, click the edit icon ( ).

Step 5 Click **Dynamic Routing** to display the dynamic routing options.

Step 6 Click **RIP** to display the RIP options.

Step 7 Under **Export Filters**, click the add icon ( ).

Step 8 From the **Name** drop-down list, choose the filter you want to add as an export filter.

Step 9 Next to **Action**, choose **Accept** or **Reject**.

Step 10 Click **OK**.

Tip To change the order of the export filters, click the move up ( ) and move down ( ) icons as needed. You can also drag the filters up or down in the list.

Step 11 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

OSPF Configuration

Open Shortest Path First (OSPF) is an adaptive routing protocol that defines routes dynamically by obtaining information from other routers and advertising routes to other routers using link state advertisements. The router keeps information about the links between it and the destination to make routing decisions. OSPF assigns a cost to each routed interface, and considers the best routes to have the lowest costs.

OSPF Routing Areas

An OSPF network may be structured, or subdivided, into routing areas to simplify administration and optimize traffic and resource use. Areas are identified by 32-bit numbers, expressed either simply in decimal or often in octet-based dot-decimal notation.

By convention, area zero or 0.0.0.0 represents the core or backbone region of an OSPF network. You may choose to identify other areas. Often, administrators select the IP address of a main router in an area as the area's identification. Each additional area must have a direct or virtual connection to the backbone OSPF area. Such connections are maintained by an interconnecting router, known as the area border router (ABR). An ABR maintains separate link state databases for each area it serves and maintains summarized routes for all areas in the network.

Adding OSPF Areas

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Virtual Routers** tab.

Step 4 Next to the virtual router you want to modify, click the edit icon ( ).

Step 5 Click **Dynamic Routing** to display the dynamic routing options.

Step 6 Click **OSPF** to display the OSPF options.

Step 7 Under **Areas**, click the add icon ( ).

Step 8 In the **Area Id** field, enter a numerical value for the area. This value can be either an integer or an IPv4 address.

Step 9 Optionally, check the **Stubnet** check box to designate that the area does not receive router advertisements external to the autonomous system and routing from within the area is based entirely on a default route. If you clear the check box, the area becomes a backbone area or otherwise non-stub area.

Step 10 In the **Default cost** field, enter a cost associated with the default route for the area.

Step 11 Under **Stubnets**, click the add icon ( ).

Step 12 In the **IP Address** field, enter an IP address in CIDR notation.

Step 13 Choose the **Hidden** check box to indicate that the stubnet is hidden.

Hidden stubnets are not propagated into other areas.

Step 14 Choose the **Summary** check box to designate that default stubnets that are subnetworks of this stubnet are suppressed.

Step 15 In the **Stub cost** field, enter a value that defines the cost associated with routing to this stub network.

Step 16 Click **OK**.

Step 17 If you want to add a network, click the add icon ( ) under **Networks**.

Step 18 In the **IP Address** field, enter an IP address in CIDR notation for the network.

Step 19 Check the **Hidden** check box to indicate that the network is hidden. Hidden networks are not propagated into other areas.

Step 20 Click **OK**.

Step 21 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

OSPF Area Interfaces

You can configure a subset of the interfaces assigned to the virtual router for OSPF. The following list describes the options you can specify on each interface.

Interfaces

Select the interface where you want to configure OSPF. Interfaces you have disabled from the Interfaces tab are not available.

Type

Select the type of OSPF interface from the following choices:

- Broadcast — On broadcast networks, flooding and hello messages are sent using multicasts, a single packet for all the neighbors. The option designates a router to be responsible for synchronizing the link state databases and originating network link state advertisements. This network type cannot be used on physically non-broadcast multiple-access (NBMP) networks and on unnumbered networks without proper IP prefixes.
- Point-to-Point (PtP) — Point-to-point networks connect just two routers together. No election is performed and no network link state advertisement is originated, which makes it simpler and faster to establish.

This network type is useful not only for physically PtP interfaces, but also for broadcast networks used as PtP links. This network type cannot be used on physically NBMP networks.

- Non-Broadcast — On NBMP networks, the packets are sent to each neighbor separately because of the lack of multicast capabilities. Similar to broadcast networks, the option designates a router, which plays a central role in the propagation of link state advertisements. This network type cannot be used on unnumbered networks.
- Autodetect — The system determines the correct type based on the specified interface.

Cost

Specify the output cost of the interface.

Stub

Specify whether the interface should listen for OSPF traffic and transmit its own traffic.

Priority

Enter a numerical value that specifies the priority value used in designated router election. On every multiple access network, the system designates a router and backup router. These routers have some special functions in the flooding process. Higher priority increases preferences in this election. You cannot configure a router with a priority of 0.

Nonbroadcast

Specify whether hello packets are sent to any undefined neighbors. This switch is ignored on any NBMA network.

Authentication

Select the OSPF authentication profile that this interface uses from one of the authentication profiles you configured on the virtual router or select **None**. For more information about configuring authentication profiles, see Adding Virtual Router Authentication Profiles, on page 1143.

Hello Interval

Type the interval, in seconds, between the sending of hello messages.

Poll

Type the interval, in seconds, between the sending of hello messages for some neighbors on NBMA networks.

Retrans Interval

Type the interval, in seconds, between retransmissions of unacknowledged updates.

Retrans Delay

Type the estimated number of seconds it takes to transmit a link state update packet over the interface.

Wait Time

Type the number of seconds that the router waits between starting election and building adjacency.

Dead Interval

Type the number of seconds that the router waits before declaring a neighbor down when not receiving messages from it. If this value is defined, it overrides the value calculated from dead count.

Dead Count

Type a numerical value that when multiplied by the hello interval specifies the number of seconds that the router waits before declaring a neighbor down when not receiving messages from it.

To edit an OSPF area interface, click the edit icon ( ). To delete an OSPF area interface, click the delete icon ( ). Disabling a configured interface from the Interfaces tab also deletes it.

Adding OSPF Area Interfaces

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Seriess	Leaf only	Admin/Network Admin

You can configure a subset of the interfaces assigned to the virtual router for OSPF.

You can choose only one interface for use in an OSPF area.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to add the OSPF interface, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Virtual Routers** tab.

Step 4 Next to the virtual router where you want to add the OSPF interface, click the edit icon ( ).

Step 5 Click **Dynamic Routing** to display the dynamic routing options.

Step 6 Click **OSPF** to display the OSPF options.

Step 7 Under **Areas**, click the add icon ( ).

Step 8 Click **Interfaces**.

Step 9 Click the add icon ( ).

Step 10 Take any of the actions as described in OSPF Area Interfaces, on page 1135.

Step 11 If you want to add a network, click the add icon ( ) under **Networks**.

Step 12 In the **IP address** field, enter an IP address for the neighbor receiving hello messages on non-broadcast networks from this interface.

Step 13 Check the **Eligible** check box to indicate that the neighbor is eligible to receive messages.

Step 14 Click **OK**.

Tip To edit a neighbor, click the edit icon ( ). To delete a neighbor, click the delete icon ( ).

Step 15 Click **OK**.

Step 16 Click **Save**.

Step 17 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Adding OSPF Area Vlinks

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

All areas in an OSPF autonomous system must be physically connected to the backbone area. In some cases where this physical connection is not possible, you can use a vlink to connect to the backbone through a non-backbone area. Vlinks can also be used to connect two parts of a partitioned backbone through a non-backbone area.

You must add a minimum of two OSPF areas before you can add a vlink.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Virtual Routers** tab.

Step 4 Next to the virtual router you want to modify, click the edit icon ( ).

Step 5 Click **Dynamic Routing** to display the dynamic routing options.

Step 6 Click **OSPF** to display the OSPF options.

Step 7 Under **Areas**, click the add icon ( ).

Step 8 Click **Vlinks**.

Step 9 Click the add icon ( ).

Step 10 In the **Router ID** field, enter an IP address for the router.

Step 11 From the **Authentication** drop-down list, choose the authentication profile the vlink will use.

Step 12 In the **Hello Interval** field, enter the interval, in seconds, between sending of hello messages.

Step 13 In the **Retrans Interval** field, enter the interval, in seconds, between retransmissions of unacknowledged updates.

Step 14 In the **Wait Time** field, enter the number of seconds that the router waits between starting election and building adjacency.

Step 15 In the **Dead Interval** field, enter the number of seconds that the router waits before declaring a neighbor down when not receiving messages from it. If this value is defined, it overrides the value calculated from dead count.

Step 16 In the **Dead Count** field, enter a numerical value that when multiplied by the hello interval, specifies the number of seconds that the router waits before declaring a neighbor down when not receiving messages from it.

Step 17 Click **OK**.

Step 18 Click **Save**.

Step 19 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Adding Import Filters for OSPF Configuration

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

You can add an import filter to define which routes are accepted or rejected from OSPF into the route table. Import filters are applied in the order they appear in the table.

When adding an import filter, you use one of the filters you configured on the virtual router.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Virtual Routers**.

Step 4 Next to the virtual router you want to modify, click the edit icon ( ).

Step 5 Click **Dynamic Routing** to display the dynamic routing options.

Step 6 Click **OSPF** to display the OSPF options.

Step 7 Under **Import Filters**, click the add icon ( ).

Step 8 From the **Name** drop-down list, choose the filter you want to add as an import filter.

Step 9 Next to **Action**, choose **Accept** or **Reject**.

Step 10 Click **OK**.

Tip To change the order of the import filters, click the move up ( ) and move down ( ) icons as needed. You can also drag the filters up or down in the list.

Step 11 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Adding Export Filters for OSPF Configuration

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

You can add an export filter to define which routes will be accepted or rejected from the route table to OSPF. Export filters are applied in the order they appear in the table.

When adding an export filter, you use one of the filters you configured on the virtual router.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Virtual Routers** tab.

Step 4 Next to the virtual router where you want to add the OSPF virtual router filter, click the edit icon ( ).

Step 5 Click the **Dynamic Routing** tab to display the dynamic routing options.

Step 6 Click **OSPF** to display the OSPF options.

Step 7 Under **Export Filters**, click the add icon ( ).

Step 8 From the **Name** drop-down list, choose the filter you want to add as an export filter.

Step 9 Next to **Action**, choose **Accept** or **Reject**.

Step 10 Click **OK**.

Tip To change the order of the export filters, click the move up ( ) and move down ( ) icons as needed. You can also drag the filters up or down in the list.

Step 11 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Virtual Router Filters

Filters provide a way to match routes for importing into the virtual router's route table and for exporting routes to dynamic protocols. You can create and manage a list of filters. Each filter defines specific criteria to look for in routes that are defined statically or received from a dynamic protocol.

The Virtual Routers Filters table includes summary information about each filter you have configured on a virtual router, as described in the following table.

Table 85: Virtual Router Filters Table View Fields

Field	Description
Name	The name of the filter.
Protocol	The protocol that the route originates from: • Static — The route originates as a local static route. • RIP — The route originates from a dynamic RIP configuration. • OSPF — The route originates from a dynamic OSPF configuration.
From Router	The router IP addresses that this filter attempts to match in a router. You must enter this valu and RIP filters.
Next Hop	The next hop where packets using this route are forwarded. You must enter this value for sta filters.
Destination Type	The type of destination where packets are sent: • Router • Device • Discard
Destination Network	The networks that this filter attempts to match in a route.
OSPF Path Type	Applies only to OSPF protocol. The path type can be one of the following: • Ext-1 • Ext-2 • Inter Area • Intra Area
OSPF Router ID	Applies only to OSPF protocol. The router ID of the router advertising that route/network.

Viewing Virtual Router Filters

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Any	Admin/Network Admin

The **Filter** tab of the virtual router editor displays a table listing of all the filters you have configured on a virtual router. The table includes summary information about each filter.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to view, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Virtual Routers** tab.

Step 4 Next to the virtual router where you want to view the filters, click the edit icon ( ).

Step 5 Click the **Filter** tab.

Setting Up Virtual Router Filters

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Virtual Routers** tab.

Step 4 Next to the virtual router you want to modify, click the edit icon ( ).

Step 5 Click the **Filter** tab.

Step 6 Click **Add Filter**.

Step 7 In the **Name** field, enter a name for the filter. You can use alphanumeric characters only.

Step 8 Under **Protocol**, choose **All** or choose the protocol that applies to the filter.

Step 9 If you chose All, Static, or RIP as the **Protocol**, under **From Router**, enter the router IP addresses that this filter will attempt to match in a route.

Note You can also enter a /32 CIDR block for IPv4 addresses and a /128 prefix length for IPv6 addresses. All other address blocks are invalid for this field.

Step 10 Click **Add**.

Step 11 If you chose All, Static, or RIP as the **Protocol**, under **Next Hop**, enter the IP addresses for the gateways that this filter will attempt to match in a route.

Note You can also enter a /32 CIDR block for IPv4 addresses and a /128 prefix length for IPv6 addresses. All other address blocks are invalid for this field.

Step 12 Click **Add**.

Step 13 Under **Destination Type**, choose the options that apply to the filter.

Step 14 Under **Destination Network**, enter the IP address of the network that this filter will attempt to match in a route.

Step 15 Click **Add**.

Step 16 If you chose All or OSPF as the **Protocol**, under **Path Type**, choose the options that apply to the filter. You must choose at least one path type.

Step 17 If you chose OSPF as the **Protocol**, under **Router ID**, enter the IP address that serves as the router ID of the router advertising the route/network.

Step 18 Click **Add**.

Step 19 Click **OK**.

Step 20 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Adding Virtual Router Authentication Profiles

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

You can set up Authentication Profiles for use in RIP and OSPF configurations. You can configure a simple password or specify a shared cryptographic key. Simple passwords allow for every packet to carry eight bytes of the password. The system ignores received packets lacking this password. Cryptographic keys allow for validation, a 16-byte long digest generated from a password to be appended to every packet.

Note that for OSPF, each area can have a different authentication method. Therefore, you create authentication profiles that can be shared among many areas. You cannot add authentication for OSPFv3.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Virtual Routers** tab.

Step 4 Next to the virtual router you want to modify, click the edit icon ( ).

Step 5 Click **Authentication Profile**.

Step 6 Click **Add Authentication Profile**.

Step 7 In the **Authentication Profile Name** field, enter a name for the authentication profile.

Step 8 From the **Authentication Type** drop down list, choose **simple** or **cryptographic**.

Step 9 In the **Password** field, enter a secure password.

Step 10 In the **Confirm Password** field, enter the password again to confirm it.

Step 11 Click **OK**.

Step 12 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Viewing Virtual Router Statistics

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Any	Admin/Network Admin

You can view runtime statistics for each virtual router. The statistics display unicast packets, packets dropped, and separate routing tables for IPv4 and IPv6 addresses.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to view statistics, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click the **Virtual Routers** tab.

Step 4 Next to the virtual router where you want to view the router statistics, click the view icon ( ).

Deleting Virtual Routers

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

When you delete a virtual router, any routed interfaces assigned to the router become available for inclusion in another router.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

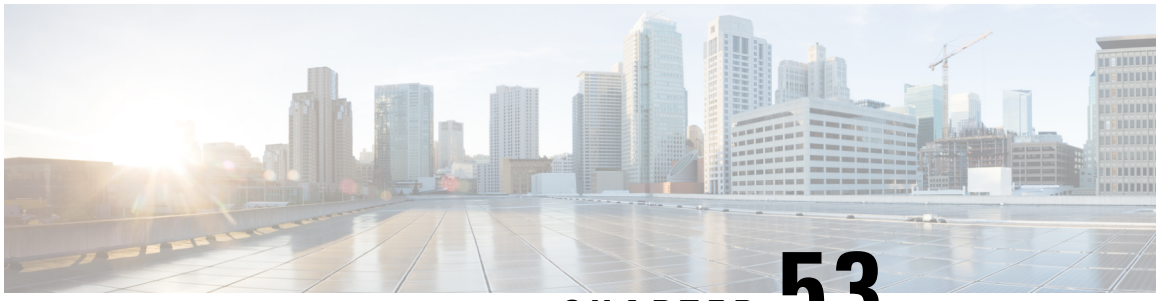
Step 3 Click the **Virtual Routers** tab.

Step 4 Next to the virtual router that you want to delete, click the delete icon ( ).

Step 5 When prompted, confirm that you want to delete the virtual router.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.



CHAPTER 53

Aggregate Interfaces and LACP

The following topics explain aggregate interface configuration and how LACP functions on managed devices:

- About Aggregate Interfaces, on page 1147
- LAG Configuration, on page 1148
- Link Aggregation Control Protocol (LACP), on page 1152
- Adding Aggregate Switched Interfaces, on page 1153
- Adding Aggregate Routed Interfaces, on page 1155
- Adding Logical Aggregate Interfaces, on page 1158
- Viewing Aggregate Interface Statistics, on page 1159
- Deleting Aggregate Interfaces, on page 1159

About Aggregate Interfaces

In the Firepower System, you can group multiple physical Ethernet interfaces into a single logical link on managed devices configured in either a Layer 2 deployment that provides packet switching between networks, or a Layer 3 deployment that routes traffic between interfaces. This single aggregate logical link provides higher bandwidth, redundancy, and load-balancing between two endpoints.

You create aggregate links by creating a switched or routed link aggregation group, or LAG. When you create an aggregation group, a logical interface called an aggregate interface is created. To an upper layer entity a LAG looks like a single logical link and data traffic is transmitted through the aggregate interface. The aggregate link provides increased bandwidth by adding the bandwidth of multiple links together. It also provides redundancy by load-balancing traffic across all available links. If one link fails, the system automatically load-balances traffic across all remaining links.



The endpoints in a LAG can be two 7000 or 8000 Series devices, as shown in the illustration above, or a 7000 or 8000 Series device connected to a third-party access switch or router. The two devices do not have to match, but they must have the same physical configuration and they must support the IEEE 802.ad link aggregation standard. A typical deployment for a LAG might be to aggregate access links between two managed devices, or to create a point-to-point connection between a managed device and an access switch or a router.

Note that you cannot configure aggregate interfaces on NGIPSv devices or ASA FirePOWER modules.

LAG Configuration

There are two types of aggregate interfaces:

- switched — Layer 2 aggregate interfaces
- routed — Layer 3 aggregate interfaces

You implement link aggregation through the use of link aggregation groups (LAGs). You configure a LAG by creating an aggregate switched or routed interface and then associating a set of physical interfaces with the link. All of the physical interfaces must be of the same speed and medium.

You create aggregate links either dynamically or statically. Dynamic link aggregation uses Link Aggregation Control Protocol (LACP), a component of the IEEE 802.ad link aggregation standard, while static link aggregation does not. LACP enables each device on either end of the LAG to exchange link and system information to determine which links will be actively used in the aggregation. A static LAG configuration requires you to manually maintain link aggregations and deploy load-balancing and link selection policies.

When you create a switched or routed aggregate interface, a link aggregation group of the same type is created and numbered automatically. For example, when you create your first LAG (switched or routed), the aggregate interface can be identified by the **lag0** label in the **Interfaces** tab for your managed device. When you associate physical and logical interfaces with this LAG, they appear nested below the primary LAG in a hierarchical tree menu. Note that a switched LAG can only contain switched physical interfaces, and a routed LAG can only contain routed physical interfaces.

Consider the following requirements when you configure a LAG:

- The Firepower System supports a maximum of 14 LAGs, and assigns a unique ID to each LAG interface in the range of 0 to 13. The LAG ID is not configurable.
- You must configure the LAG on both sides of the link, and you must set the interfaces on either side of the link to the same speed.
- You must associate at least two physical interfaces per LAG, up to a maximum of eight. A physical interface cannot belong to more than one LAG.
- Physical interfaces in a LAG cannot be used in any other mode of operation, either as inline or passive, or be used as part of another logical interface for tagged traffic.
- Physical interfaces in a LAG can span multiple NetMods, but cannot span multiple sensors (i.e. all physical interfaces must reside on the same device).
- A LAG cannot contain a stacking NetMod.

Aggregate Switched Interfaces

You can combine between two and eight physical ports on a managed device to create a switched LAG interface. You must assign a switched LAG interface to a virtual switch before it can handle traffic. A managed device can support up to 14 LAG interfaces.

The range of MTU values can vary depending on the model of the managed device and the interface type.

Caution Changing the highest MTU value among all non-management interfaces on the device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all non-management interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. See Snort® Restart Traffic Behavior, on page 350 for more information.

Related Topics

MTU Ranges for 7000 and 8000 Series Devices and NGIPSv, on page 500
Snort® Restart Scenarios, on page 348

Aggregate Routed Interfaces

You can combine between two and eight physical ports on a7000 or 8000 Series device to create a routed LAG interface. You must assign a routed LAG interface to a virtual router before it can route traffic. A managed device can support up to 14 LAG interfaces.

You can add static Address Resolution Protocol (ARP) entries to a routed LAG interface. If an external host needs to know the MAC address of the destination IP address it needs to send traffic to on your local network, it sends an ARP request. When you configure static ARP entries, the virtual router responds with an IP address and associated MAC address.

Disabling the **ICMP Enable Responses** option for routed LAG interfaces does not prevent ICMP responses in all scenarios. You can still use access control rules to handle connections where the destination IP is the routed interface's IP and the protocol is ICMP; see Port and ICMP Code Conditions, on page 373.

If you enable the **Inspect Local Router Traffic** option, the system blocks packets before they reach the host, thereby preventing any response. For more information about inspecting local router traffic, see Advanced Settings, on page 251.

The range of MTU values can vary depending on the model of the managed device and the interface type.

Caution Changing the highest MTU value among all non-management interfaces on the device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all non-management interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. See Snort® Restart Traffic Behavior, on page 350 for more information.

Related Topics

MTU Ranges for 7000 and 8000 Series Devices and NGIPSv, on page 500
Snort® Restart Scenarios, on page 348

Logical Aggregate Interfaces

For each switched or routed aggregate interface, you can add multiple logical interfaces. You must associate each logical LAG interface with a VLAN tag to handle traffic received by the LAG interface with that specific tag. You add logical interfaces to switched or routed aggregate interfaces in the same way you would add them to physical switched or routed interfaces.

Note When you create a LAG interface you also create an "untagged" logical interface by default, which is identified by the **lag*n*.0** label, where ***n*** is an integer from 0 to 13. To be operational, each LAG requires this one logical interface at a minimum. You can associate additional logical interfaces with any LAG to handle VLAN-tagged traffic. Each additional logical interface requires a unique VLAN tag. The Firepower System supports VLAN tags in the range of 1 through 4094.

You can also configure the Cisco Redundancy Protocol (SFRP) on a logical routed interface. SFRP allows devices to act as redundant gateways for specified IP addresses.

Note that disabling the **ICMP Enable Responses** option for logical routed interfaces does not prevent ICMP responses in all scenarios. You can add network-based rules to an access control policy to drop packets where the destination IP is the routed interface's IP and the protocol is ICMP.

If you have enabled the **Inspect Local Router Traffic** option, which is an advanced setting on the managed device, it drops the packets before they reach the host, thereby preventing any response.

The range of MTU values can vary depending on the model of the managed device and the interface type.

Caution Changing the highest MTU value among all non-management interfaces on the device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all non-management interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. See Snort® Restart Traffic Behavior, on page 350 for more information.

Related Topics

SFRP
Advanced Settings, on page 251
MTU Ranges for 7000 and 8000 Series Devices and NGIPSv, on page 500
Snort® Restart Scenarios, on page 348

Load-Balancing Algorithms

You assign an egress load-balancing algorithm to the LAG that determines how to distribute traffic to the LAG bundle's member links. The load-balancing algorithm makes hashing decisions based on values in various packet fields, such as Layer 2 MAC addresses, Layer 3 IP addresses, and Layer 4 port numbers (TCP/UDP traffic). The load-balancing algorithm you select applies to all of the LAG bundle's member links.

Choose the load-balancing algorithm that supports your deployment scenario from the following options when you configure a LAG:

- Destination IP
- Destination MAC
- Destination Port
- Source IP
- Source MAC

- Source Port
- Source and Destination IP
- Source and Destination MAC
- Source and Destination Port

Note You should configure both ends of the LAG to have the same load-balancing algorithm. Higher layer algorithms will back off to lower layer algorithms as necessary (such as a Layer 4 algorithm backing off to Layer 3 for ICMP traffic).

Link Selection Policies

Link aggregation requires the speed and medium of each link to be the same at both endpoints. Because link properties can change dynamically, the link selection policy helps determine how the system manages the link selection process. A link selection policy that maximizes the highest port count supports link redundancy, while a link selection policy that maximizes total bandwidth supports overall link speed. A stable link selection policy attempts to minimize excessive changes in link states.

Note You should configure both ends of the LAG to have the same link selection policy.

Choose the link selection policy that supports your deployment scenario from the following options:

- Highest Port Count — Choose this option for the highest total active port count to provide added redundancy.
- Highest Total Bandwidth — Choose this option to provide the highest total bandwidth for the aggregated link.
- Stable — Choose this option if your primary concern is link stability and reliability. Once you configure a LAG, the active links change only when absolutely necessary (such as link failure) rather than doing so for added port count or bandwidth.
- LACP Priority — Choose this option to use the LACP algorithm to determine which links are active in the LAG. This setting is appropriate if you have undefined deployment goals, or if the device at the other end of the LAG is not managed by the Firepower Management Center.

LACP is a key aspect of automating the link selection method that supports dynamic link aggregation. When LACP is enabled, a link selection policy based on LACP priority uses the following properties of LACP:

LACP system priority

You configure this value on each partnered device running LACP to determine which one is superior in link aggregation. The system with the lower value has the higher system priority. In dynamic link aggregation, the system with the higher LACP system priority sets the selected state of member links on its side first, then the system with the lower priority sets its member links accordingly. You can specify 0 to 65535. If you do not specify a value, the default priority is 32768.

LACP link priority

You configure this value on each link belonging to the aggregation group. The link priority determines the active and standby links in the LAG. Links with lower values have higher priority. If an active link goes down, the standby link with the highest priority is selected to replace the downed link. However, if two or more links have the same LACP link priority, the link with the lowest physical port number is selected as the standby link. You can specify 0 to 65535. If you do not specify a value, the default priority is 32768.

Link Aggregation Control Protocol (LACP)

Link Aggregation Control Protocol (LACP), a component of IEEE 802.3ad, is a method of exchanging system and port information to create and maintain LAG bundles. When you enable LACP, each device on either end of the LAG uses LACP to determine which links will be actively used in the aggregation. LACP provides availability and redundancy by exchanging LACP packets (or control messages) between links. It learns the capabilities of the links dynamically and informs the other links. Once LACP identifies correctly matched links, it facilitates grouping the links into the LAG. If a link fails, traffic continues on the remaining links. LACP must be enabled at both ends of the LAG for the link to be operational.

LACP

When you enable LACP, you need to specify a transmission mode for each end of the LAG that determines how LACP packets are exchanged between partnered devices. There are two options for LACP mode:

- Active — Choose this mode to place a device into an active negotiating state, in which the device initiates negotiations with remote links by sending LACP packets.
- Passive — Choose this mode to place a device into a passive negotiating state, in which the device responds to LACP packets it receives but does not initiate LACP negotiation.



Note Both modes allow LACP to negotiate between links to determine if they can form a link bundle based on criteria such as port speed. However, you should avoid a passive-passive configuration, which essentially places both ends of the LAG in listening mode.

LACP has a timer which defines how often LACP packets are sent between devices. LACP exchanges packets at these rates:

- Slow — 30 seconds
- Fast — 1 second

The device where this option is applied expects to receive LACP packets with this frequency from the partner device on the other side of the LAG.

Note When a LAG is configured on a managed device that is part of a device stack, only the primary device participates in LACP communication with the partner system. All secondary devices forward LACP messages to the primary device. The primary device relays any dynamic LAG modifications to the secondary devices.

Adding Aggregate Switched Interfaces

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

You can combine between two and eight physical ports on a managed device to create a switched LAG interface. You must assign a switched LAG interface to a virtual switch before it can handle traffic. A managed device can support up to 14 LAG interfaces.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **Edit** ( ) next to the device where you want to configure the switched LAG interface.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Choose **Add Aggregate Interface** from the **Add** drop-down menu.

Step 4 Click **Switched** to display the switched LAG interface options.

Step 5 If you want to apply a security zone, do one of the following:

- Choose an existing security zone from the **Security Zone** drop-down list.
- Choose **New** to add a new security zone; see Creating Security Zone and Interface Group Objects, on page 411.

Step 6 Specify a virtual switch:

- Choose an existing virtual switch from the **Virtual Switch** drop-down list.
- Choose **New** to add a new virtual switch; see Adding Virtual Switches, on page 1111.

Step 7 Check the **Enabled** check box to allow the switched LAG interface to handle traffic.

If you clear the check box, the interface becomes disabled so that users cannot access it for security purposes.

Step 8 From the **Mode**, choose an option to designate the link mode, or choose **Autonegotiation** to specify that the interface is configured to auto negotiate speed and duplex settings.

Mode settings are available only for copper interfaces.

Interfaces on 8000 Series appliances do not support half-duplex options. When links auto negotiate speed, all active links are selected for the LAG based on the same speed setting.

Step 9 From the **MDI/MDIX** drop-down list, choose an option to designate whether the interface is configured for MDI (medium dependent interface), MDIX (medium dependent interface crossover), or Auto-MDIX.

MDI/MDIX settings are available only for copper interfaces.

By default, MDI/MDIX is set to Auto-MDIX, which automatically handles switching between MDI and MDIX to attain link.

Step 10 Enter a maximum transmission unit (MTU) in the **MTU** field.

The range within which you can set the MTU can vary depending on the Firepower System device model and interface type. See MTU Ranges for 7000 and 8000 Series Devices and NGIPSv, on page 500 for more information.

Step 11 Under **Link Aggregation**, choose one or more physical interfaces from **Available Interfaces** to add to the LAG bundle.

Tip To remove physical interfaces from the LAG bundle, choose one or more physical interfaces and click the **Remove Selected icon**. To remove all physical interfaces from the LAG bundle, click the **Remove All icon**. Deleting the LAG interface from the Interfaces tab also removes the interfaces.

Step 12 Choose an option from the **Load-Balancing Algorithm** drop-down list.

Step 13 Choose a **Link Selection Policy** from the drop-down list.

Tip Choose **LACP Priority** if you are configuring an aggregate interface between a Firepower System device and a third-party network device.

Step 14 If you chose **LACP Priority** as the **Link Selection Policy**, assign a value for **System Priority** and click the **Configure Interface Priority** link to assign a priority value for each interface in the LAG.

Step 15 Choose either **Inner** or **Outer** from the **Tunnel Level** drop-down list.

Note The tunnel level only applies to IPv4 traffic when Layer 3 load balancing is configured. The outer tunnel is always used for Layer 2 and IPv6 traffic. If the **Tunnel Level** is not explicitly set, the default is **Outer**.

Step 16 Under **LACP**, check the **Enabled** check box to allow the switched LAG interface to handle traffic using the Link Aggregation Control Protocol.

If you clear the check box, the LAG interface becomes a static configuration and the Firepower System will use all of the physical interfaces selected for the aggregation.

Step 17 Click a **Rate** radio button to set the frequency that determines how often LACP control messages are received from the partner device:

- Click **Slow** to receive packets every 30 seconds.
- Click **Fast** to receive packets every 1 second.

Step 18 Click a **Mode** radio button to establish the listening mode of the device:

- Click **Active** to initiate negotiations with remote links by sending LACP packets to the partner device.
- Click **Passive** to respond to LACP packets received.

Step 19 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

MTU Ranges for 7000 and 8000 Series Devices and NGIPSv, on page 500
Snort® Restart Scenarios, on page 348

Adding Aggregate Routed Interfaces

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

You can combine between two and eight physical ports on a managed device to create a routed LAG interface. You must assign a routed LAG interface to a virtual router before it can route traffic. A managed device can support up to 14 LAG interfaces.

Caution Adding a routed interface pair on 7000 or 8000 Series devices restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **Edit** ( ) next to the device where you want to configure the routed LAG interface.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Choose **Add Aggregate Interface** from the **Add** drop-down menu.

Step 4 Click **Routed** to display the routed LAG interface options.

Step 5 If you want to apply a security zone, do one of the following:

- Choose an existing security zone from the **Security Zone** drop-down list.
- Choose **New** to add a new security zone; see Creating Security Zone and Interface Group Objects, on page 411.

Step 6 Specify a virtual router:

- Choose an existing virtual router from the **Virtual Router** drop-down list.
- Choose **New** to add a new virtual router; Adding Virtual Routers, on page 1124.

Step 7 Check the **Enabled** check box to allow the routed LAG interface to handle traffic.

If you clear the check box, the interface becomes disabled so that users cannot access it for security purposes.

Step 8 From the **Mode** drop-down list, choose an option to designate the link mode, or choose **Autonegotiation** to specify that the LAG interface is configured to auto negotiate speed and duplex settings.

Mode settings are available only for copper interfaces.

Interfaces on 8000 Series appliances do not support half-duplex options. When links auto negotiate speed, all active links are selected for the LAG based on the same speed setting.

Step 9 Choose an option from the **MDI/MDIX** drop-down list to designate whether the interface is configured for MDI (medium dependent interface), MDIX (medium dependent interface crossover), or Auto-MDIX.

MDI/MDIX settings are available only for copper interfaces.

By default, MDI/MDIX is set to Auto-MDIX, which automatically handles switching between MDI and MDIX to attain link.

Step 10 Enter a maximum transmission unit (MTU) in the **MTU** field.

The range of MTU values can vary depending on the model of the managed device and the interface type.

Caution Changing the highest MTU value among all non-management interfaces on the device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all non-management interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. See Snort® Restart Traffic Behavior, on page 350 for more information.

Step 11 If you want to allow the LAG interface to respond to ICMP traffic such as pings and traceroute, check the **Enable Responses** check box next to **ICMP**.

Step 12 If you want to enable the LAG interface to broadcast router advertisements, check the **Enable Router Advertisement** check box next to **IPv6 NDP**.

Step 13 Click **Add** to add an IP address.

Step 14 In the **Address** field, enter the routed LAG interface's IP address and subnet mask using CIDR notation.

Note the following:

- You cannot add network and broadcast addresses, or the static MAC addresses 00:00:00:00:00:00 and FF:FF:FF:FF:FF:FF.
- You cannot add identical IP addresses, regardless of subnet mask, to interfaces in virtual routers.

Step 15 If your organization uses IPv6 addresses and you want to set the IP address of the LAG interface automatically, check the **Address Autoconfiguration** check box next to the **IPv6** field.

Step 16 For **Type**, choose either Normal or SFRP.

Step 17 If you chose SFRP for **Type**, set options as described in SFRP.

Step 18 Click **OK**.

Note When adding an IP address to a routed interface of a 7000 or 8000 Series device in a high-availability pair, you must add a corresponding IP address to the routed interface on the high-availability peer.

Step 19 Click **Add** to add a static ARP entry.

Step 20 Enter an IP address the **IP Address** field.

Step 21 Enter a MAC address to associate with the IP address in the **MAC Address** field. Use the standard format (for example, 01:23:45:67:89:AB).

Step 22 Click **OK**.

Step 23 Under **Link Aggregation**, choose one or more physical interfaces from **Available Interfaces** to add to the LAG bundle.

Tip To remove physical interfaces from the LAG bundle, choose one or more physical interfaces and click the **Remove Selected icon** . To remove all physical interfaces from the LAG bundle, click the **Remove All icon**. Deleting the LAG interface from the **Interfaces** tab also removes the interfaces.

Step 24 Choose a **Load-Balancing Algorithm** from the drop-down list.

Step 25 Choose a **Link Selection Policy** from the drop-down list.

Tip Choose **LACP Priority** if you are configuring an aggregate interface between a Firepower System device and a third-party network device.

Step 26 If you chose **LACP Priority** as the **Link Selection Policy**, assign a value for **System Priority** and click the **Configure Interface Priority** link to assign a priority value for each interface in the LAG.

Step 27 Choose either **Inner** or **Outer** from the **Tunnel Level** drop-down list.

Note The tunnel level only applies to IPv4 traffic when Layer 3 load balancing is configured. The outer tunnel is always used for Layer 2 and IPv6 traffic. If the **Tunnel Level** is not explicitly set, the default is **Outer**.

Step 28 Under **LACP**, check the **Enabled** check box to allow the routed LAG interface to handle traffic using the Link Aggregation Control Protocol.

If you clear the check box, the LAG interface becomes a static configuration and the Firepower System will use all of the physical interfaces for the aggregation.

Step 29 Click a **Rate** radio button to set the frequency that determines how often LACP control messages are received from the partner device.

- Click **Slow** to receive packets every 30 seconds.
- Click **Fast** to receive packets every 1 second.

Step 30 Click a **Mode** radio button to establish the listening mode of the device.

- Click **Active** to initiate negotiations with remote links by sending LACP packets to the partner device.
- Click **Passive** to respond to LACP packets received.

Step 31 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Advanced Settings, on page 251

Adding Logical Aggregate Interfaces

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

For each switched or routed aggregate interface, you can add multiple logical interfaces. You must associate each logical LAG interface with a VLAN tag to handle traffic received by the LAG interface with that specific tag. You add logical interfaces to switched or routed aggregate interfaces in the same way you would add them to physical switched or routed interfaces.

Note When you create a LAG interface you also create an "untagged" logical interface by default, which is identified by the **lag*n*.0** label, where ***n*** is an integer from 0 to 13. To be operational, each LAG requires this one logical interface at a minimum. You can associate additional logical interfaces with any LAG to handle VLAN-tagged traffic. Each additional logical interface requires a unique VLAN tag. The Firepower System supports VLAN tags in the range of 1 through 4094.

Caution Adding a routed interface pair on 7000 or 8000 Series devices restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to add the logical LAG interface, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 From the **Add** drop-down menu, choose **Add Logical Interface**.

Step 4 Click **Switched** to display the switched interface options, or click **Routed** to display the routed interface options.

Step 5 Choose an available LAG from the **Interface** drop-down list. The aggregate interface is identified by the **lag*n*** label, where ***n*** is an integer from 0 to 13.

Step 6 Configure the remaining settings appropriate to the interface type you chose:

- Switched — See Adding Logical Switched Interfaces, on page 1108 for more information on adding a logical interface to a switched interface.
- Routed — See Adding Logical Routed Interfaces, on page 1119 for more information on adding a logical interface to a routed interface.

Related Topics

SFRP
Advanced Settings, on page 251
MTU Ranges for 7000 and 8000 Series Devices and NGIPSv, on page 500
Snort® Restart Scenarios, on page 348

Viewing Aggregate Interface Statistics

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

You can view protocol and traffic statistics for each aggregate interface. The statistics show LACP protocol information such as LACP key and partner information, packets received, packets transmitter, and packets dropped. Statistics are further refined per member interface to show traffic and link information on a per-port basis.

Aggregate interface information is also presented to the dashboard via predefined dashboard widgets. The Current Interface Status widget shows the status of all interfaces on the appliance, enabled or unused. The Interface Traffic widget shows the rate of traffic received (Rx) and transmitted (Tx) on the appliance's interfaces over the dashboard time range. See Predefined Dashboard Widgets, on page 258.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to view the logical aggregate interface statistics, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Next to the interface where you want to view the interface statistics, click the view icon ( ).

Deleting Aggregate Interfaces

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

The aggregate interface can be identified by the **lag*n*** label, where ***n*** can be an integer from 0 to 13.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to delete the aggregate interface, click the edit icon ( ).

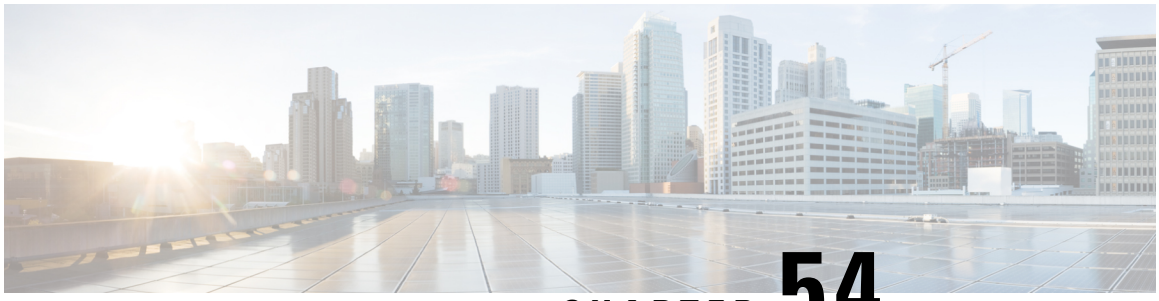
In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Next to the aggregate interface you want to delete, click the delete icon ( ).

Step 4 When prompted, confirm that you want to delete the aggregate interface.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.



CHAPTER 54

Hybrid Interfaces

The following topics describe how to configure local hybrid interfaces:

- About Hybrid Interfaces, on page 1161
- Logical Hybrid Interfaces, on page 1161
- Adding Logical Hybrid Interfaces, on page 1162
- Deleting Logical Hybrid Interfaces, on page 1164

About Hybrid Interfaces

You can configure logical hybrid interfaces on managed devices that allow the Firepower System to bridge traffic between virtual routers and virtual switches. If IP traffic received on interfaces in a virtual switch is addressed to the MAC address of an associated hybrid logical interface, the system handles it as Layer 3 traffic and either routes or responds to the traffic depending on the destination IP address. If the system receives any other traffic, it handles it as Layer 2 traffic and switches it appropriately. You cannot configure logical hybrid interfaces on an NGIPSv device.

Note that hybrid interfaces that are not associated with both a virtual switch and a virtual router are not available for routing, and do not generate or respond to traffic.

Logical Hybrid Interfaces

You must associate a logical hybrid interface with a virtual router and virtual switch to bridge traffic between Layer 2 and Layer 3. You can only associate a single hybrid interface with a virtual switch. However, you can associate multiple hybrid interfaces with a virtual router.

You can also configure the Cisco Redundancy Protocol (SFRP) on a logical hybrid interface. SFRP allows devices to act as redundant gateways for specified IP addresses.

Note that disabling the **ICMP Enable Responses** option for hybrid interfaces does not prevent ICMP responses in all scenarios. You can add network-based rules to an access control policy to drop packets where the destination IP is the hybrid interface's IP and the protocol is ICMP.

If you have enabled the **Inspect Local Router Traffic** option on the managed device, it drops the packets before they reach the host, thereby preventing any response.

The range of MTU values can vary depending on the model of the managed device and the interface type.

Caution Changing the highest MTU value among all non-management interfaces on the device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all non-management interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. See Snort® Restart Traffic Behavior, on page 350 for more information.

Related Topics

Configuring SFRP, on page 1122
Advanced Settings, on page 251
MTU Ranges for 7000 and 8000 Series Devices and NGIPSv, on page 500
Snort® Restart Scenarios, on page 348

Adding Logical Hybrid Interfaces

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

Caution Adding a routed interface pair on 7000 or 8000 Series devices restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to add the hybrid interface, click the edit icon ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 From the **Add** drop-down menu, choose **Add Logical Interface**.

Step 4 Click **Hybrid** to display the hybrid interface options.

Step 5 In the **Name** field, enter a name for the interface.

Step 6 From the **Virtual Router** drop-down list, choose an existing virtual router, choose **None**, or choose **New** to add a new virtual router.

Note If you add a new virtual router, you must configure it on the Device Management page after you finish setting up the hybrid interface. See Adding Virtual Routers, on page 1124.

Step 7 From the **Virtual Switch** drop-down list, choose an existing virtual switch, choose **None**, or choose **New** to add a new virtual switch.

Note If you add a new virtual switch, you must configure it on the Device Management page after you finish setting up the hybrid interface. See Adding Virtual Switches, on page 1111.

Step 8 Check the **Enabled** check box to allow the hybrid interface to handle traffic.

Note If you clear the check box, the interface becomes disabled and administratively taken down.

Step 9 In the **MTU** field, enter a maximum transmission unit (MTU), which designates the largest size packet allowed. The range of MTU values can vary depending on the model of the managed device and the interface type.

Caution Changing the highest MTU value among all non-management interfaces on the device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all non-management interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. See Snort® Restart Traffic Behavior, on page 350 for more information.

Step 10 Next to **ICMP**, check the **Enable Responses** check box to allow the interface to respond to ICMP traffic such as pings and traceroute.

Step 11 Next to **IPv6 NDP**, check the **Enable Router Advertisement** check box to enable the interface to broadcast router advertisements. You can only enable this option if you added IPv6 addresses.

Step 12 To add an IP address, click **Add**.

Step 13 In the **Address** field, enter the IP address and subnet mask. Note the following:

- You cannot add network and broadcast addresses, or the static MAC addresses 00:00:00:00:00:00 and FF:FF:FF:FF:FF:FF.
- You cannot add identical IP addresses, regardless of subnet mask, to interfaces in virtual routers.

Step 14 Optionally if you have IPv6 addresses, next to the **IPv6** field, check the **Address Autoconfiguration** check box to set the IP address of the interface automatically.

Step 15 For **Type**, choose either Normal or SFRP.

Step 16 If you chose SFRP for **Type**, set options as described in SFRP.

Step 17 Click **OK**.

Step 18 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

MTU Ranges for 7000 and 8000 Series Devices and NGIPSv, on page 500
Snort® Restart Scenarios, on page 348

Deleting Logical Hybrid Interfaces

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to delete the logical hybrid interface, click the edit icon ( ).

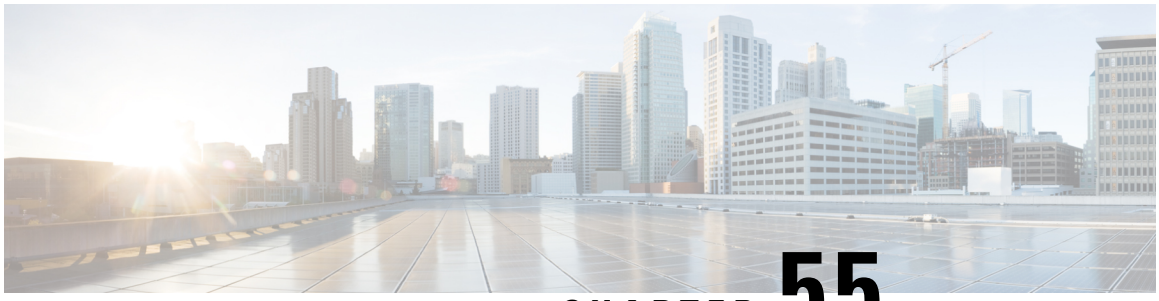
In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Next to the logical hybrid interface you want to delete, click the delete icon ( ).

Step 4 When prompted, confirm that you want to delete the interface.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.



CHAPTER 55

Gateway VPNs

The following topics describe how to manage your VPN deployment:

- Gateway VPN Basics, on page 1165
- VPN Deployments, on page 1166
- VPN Deployment Management, on page 1168
- VPN Deployment Status, on page 1179
- VPN Statistics and Logs, on page 1179

Gateway VPN Basics

A virtual private network (VPN) is a network connection that establishes a secure tunnel between endpoints via a public source, such as the internet or other network. You can configure the Firepower System to build secure VPN tunnels between the virtual routers of Firepower managed devices. The system builds tunnels using the Internet Protocol Security (IPsec) protocol suite.

After the VPN connection is established, the hosts behind the local gateway can connect to the hosts behind the remote gateway through the secure VPN tunnel. A connection consists of the IP addresses and host names of the two gateways, the subnets behind them, and the shared secrets for the two gateways to authenticate to each other.

The VPN endpoints authenticate to each other with either the Internet Key Exchange (IKE) version 1 or version 2 protocol to create a security association for the tunnel. The system uses either the IPsec authentication header (AH) protocol or the IPsec encapsulating security payload (ESP) protocol to authenticate the data entering the tunnel. The ESP protocol encrypts the data as well as providing the same functionality as AH.

If you have access control policies in your deployment, the system does not send VPN traffic until it has passed through access control. In addition, the system does not send tunnel traffic to the public source when the tunnel is down.

To configure and deploy VPN for Firepower, you must have a VPN license enabled on each of your target managed devices. Additionally, VPN features are only available on 7000 and 8000 Series devices.

IPsec

The IPsec protocol suite defines how IP packets across a VPN tunnel are hashed, encrypted, and encapsulated in the ESP or AH security protocol. The Firepower System uses the hash algorithm and encryption key of the Security Association (SA), which becomes established between the two gateways by the Internet Key Exchange (IKE) protocol.

Security associations (SA) establish shared security attributes between two devices and allow VPN endpoints to support secure communication. An SA allows two VPN endpoints to handle the parameters for how the VPN tunnel is secured between them.

The system uses the Internet Security Association and Key Management Protocol (ISAKMP) during the initial phase of negotiating the IPsec connection to establish the VPN between endpoints and the authenticated key exchange. The IKE protocol resides within ISAKMP.

The AH security protocol provides protection for packet headers and data, but it cannot encrypt them. ESP provides encryption and protection for packets, but it cannot secure the outermost IP header. In many cases, this protection is not required, and most VPN deployments use ESP more frequently than AH because of its encryption capabilities. Since VPN only operates in tunnel mode, the system encrypts and authenticates the entire packet from Layer 3 and up in the ESP protocol. ESP in tunnel mode encrypts the data as well as providing the latter's encryption capabilities.

IKE

The Firepower System uses the IKE protocol to mutually authenticate the two gateways against each other as well as to negotiate the SA for the tunnel. The process consists of two phases.

IKE phase 1 establishes a secure authenticated communication channel by using the Diffie-Hellman key exchange to generate a pre-shared key to encrypt further IKE communications. This negotiation results in a bidirectional ISAKMP security association. The system allows you to perform the authentication using a pre-shared key. Phase 1 operates in main mode, which seeks to protect all data during the negotiation, while also protecting the identity of the peers.

During IKE phase 2, the IKE peers use the secure channel established in phase 1 to negotiate security associations on behalf of IPsec. The negotiation results in a minimum of two unidirectional security associations, one inbound and one outbound.

VPN Deployments

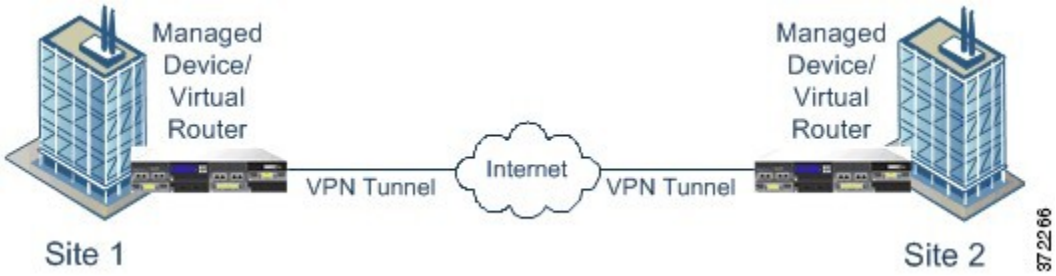
A VPN deployment specifies the endpoints and networks that are included in a VPN and how they connect to each other. After you configure a VPN deployment on the Firepower Management Center, you can then deploy it to your managed devices or devices managed by another Firepower Management Center.

The system supports three types of VPN deployments: point-to-point, star, and mesh.

Point-to-Point VPN Deployments

In a point-to-point VPN deployment, two endpoints communicate directly with each other. You configure the two endpoints as peer devices, and either device can start the secured connection. Each of the devices in this configuration must be a VPN-enabled managed device.

The following diagram displays a typical point-to-point VPN deployment.



Star VPN Deployments

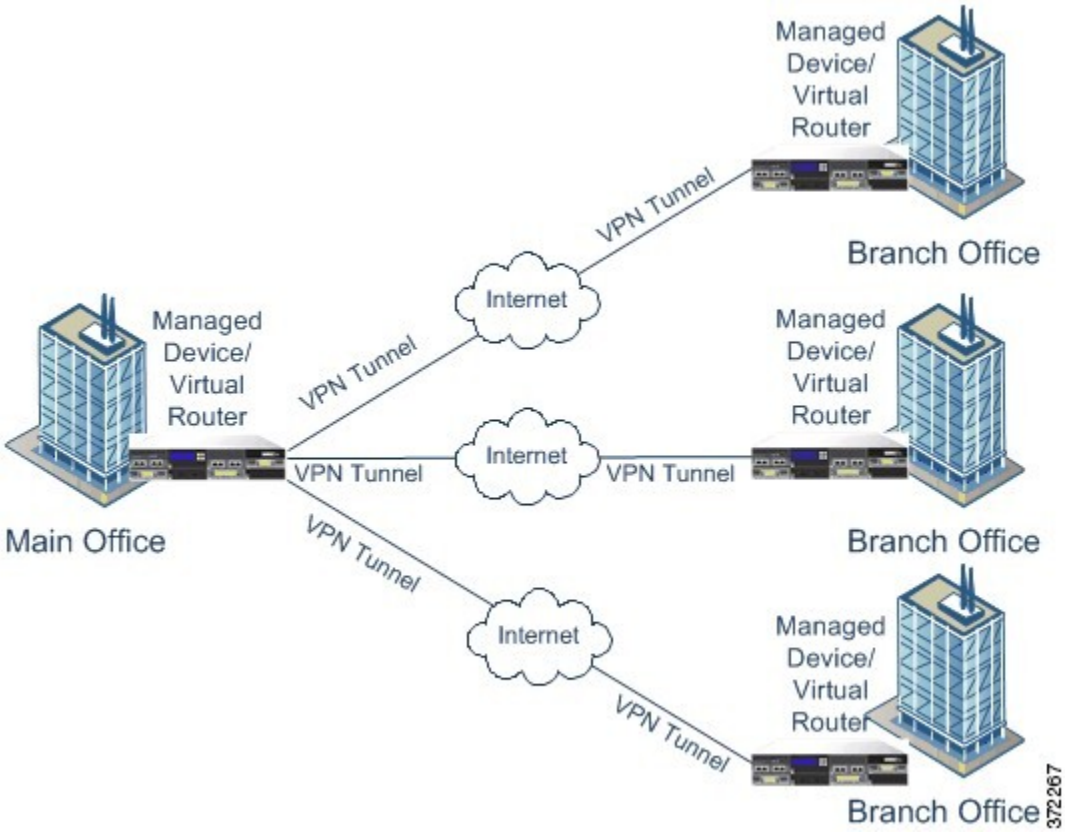
In a star VPN deployment, a central endpoint (hub node) establishes a secure connection with multiple remote endpoints (leaf nodes). Each connection between the hub node and an individual leaf node is a separate VPN tunnel. The hosts behind any of the leaf nodes can communicate with each other through the hub node.

Star deployments commonly represent a VPN that connects an organization's main and branch office locations using secure connections over the Internet or other third-party network. Star VPN deployments provide all employees with controlled access to the organization's network.

In a typical star deployment, the hub node is located at the main office. Leaf nodes are located at branch offices and start most of the traffic. Each of the nodes must be a VPN-enabled managed device.

Star deployments only support IKE version 2.

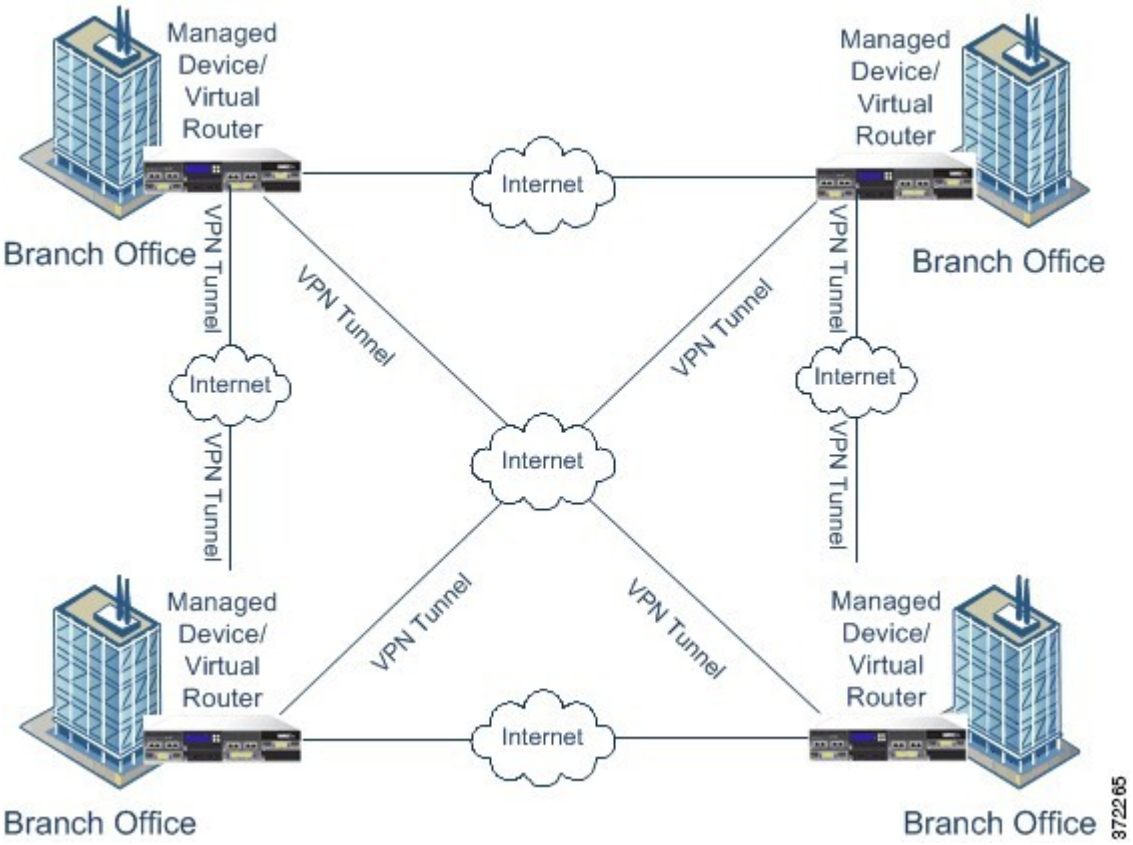
The following diagram displays a typical star VPN deployment.



Mesh VPN Deployments

In a mesh VPN deployment, all endpoints can communicate with every other endpoint by an individual VPN tunnel. The mesh deployment offers redundancy so that when one endpoint fails, the remaining endpoints can still communicate with each other. This type of deployment commonly represents a VPN that connects a group of decentralized branch office locations. The number of VPN-enabled managed devices you deploy in this configuration depends on the level of redundancy you require. Each of the endpoints must be a VPN-enabled managed device.

The following diagram displays a typical mesh VPN deployment.



VPN Deployment Management

On the VPN page (**Devices** > **VPN** > **Site to Site**), you can view all of your current VPN deployments by name and the endpoints contained in the deployment. Options on this page allow you to view the status of a VPN deployment, create a new deployment, deploy to managed devices, and edit or delete a deployment.

Note that when you register a device to a Firepower Management Center, deployed VPN deployments sync to the Firepower Management Center during registration.

Related Topics

Managing VPN Deployments, on page 1174

VPN Deployment Options

When you create a new VPN deployment you must, at minimum, give it a unique name, specify a deployment type, and designate a preshared key. You can select from three types of deployment, each containing a group of VPN tunnels:

- Point-to-point (PTP) deployments establish a VPN tunnel between two endpoints.
- Star deployments establish a group of VPN tunnels connecting a hub endpoint to a group of leaf endpoints.
- Mesh deployments establish a group of VPN tunnels among a set of endpoints.

Only Cisco managed devices can be used as endpoints in VPN deployments. Third-party endpoints are not supported.

You must define a pre-shared key for VPN authentication. You can specify a default key to use in all of the VPN connections you generate in a deployment. For point-to-point deployments, you can specify a preshared key for each endpoint pair.

In a multidomain deployment, you can configure a VPN deployment across domains; that is, you can assign endpoints to devices that belong to different domains. In such cases, you can view but not modify the ancestor deployment in the related descendant domains. When you drill down for deployment details, the system displays information for devices that belong to the current domain only.

Point-to-Point VPN Deployment Options

When configuring a point-to-point VPN deployment, you define a group of endpoint pairs and then create a VPN between the two nodes in each pair.

The following list describes the options you can specify in your deployment.

Name

Specify a unique name for the deployment.

Type

Click **PTP** to specify that you are configuring a point-to-point deployment.

Pre-shared Key

Define a unique pre-shared key for authentication. The system uses this key for all the VPNs in your deployment, unless you specify a pre-shared key for each endpoint pair.

Device

You can choose a managed device, including a device stack or device high-availability pair, as an endpoint for your deployment. For Cisco-managed devices not managed by the Firepower Management Center you are using, choose **Other** and then specify an IP address for the endpoint.

Virtual Router

If you chose a managed device as your endpoint, choose a virtual router that is currently applied to the selected device. You cannot choose the same virtual router for more than one endpoint.

Interface

If you chose a managed device as your endpoint, choose a routed interface that is assigned to the virtual router you specified.

IP Address

- If you chose a managed device as an endpoint, choose an IP address that is assigned to the specified routed interface.
- If the managed device is a device high-availability pair, you can choose only from a list of SFRP IP addresses.
- If you choose a managed device **not** managed by the Firepower Management Center, specify an IP address for the endpoint.

Protected Networks

Specify the networks in your deployment that are encrypted. Enter a subnet with CIDR block for each network. IKE version 1 only supports a single protected network.

Note that VPN endpoints cannot have the same IP address and that protected networks in a VPN endpoint pair cannot overlap. If a list of protected networks for an endpoint contains one or more IPv4 or IPv6 entry, the other endpoint's protected network must have at least one entry of the same type (i.e., IPv4 or IPv6). If it does not, then the other endpoint's IP address must be of the same type and must not overlap with the entries in the protected network. (Use /32 CIDR address blocks for IPv4 and /128 CIDR address blocks for IPv6). If both of these checks fail, the endpoint pair is invalid.

Internal IP

Check the check box if the endpoint resides behind a firewall with network address translation.

Public IP

If you checked the **Internal IP** check box, specify a public IP address for the firewall. If the endpoint is a responder, you must specify this value.

Public IKE Port

If you checked the **Internal IP** check box, specify a single numerical value from 1 to 65535 for the UDP port on the firewall that is being port-forwarded to the internal endpoint. If the endpoint is a responder and the port on the firewall being forwarded is not 500 or 4500, you must specify this value.

Use Deployment Key

Check the check box to use the pre-shared key defined for the deployment. Clear the check box to specify a pre-shared key for VPN authentication for this endpoint pair.

Pre-shared Key

If you cleared the **Use Deployment Key** check box, specify a pre-shared key in this field.

Related Topics

Configuring Point-to-Point VPN Deployments, on page 1175

Star VPN Deployment Options

When configuring a star VPN deployment, you define a single hub node endpoint and a group of leaf node endpoints. You must define the hub node endpoint and at least one leaf node endpoint to configure the deployment.

The following list describes the options you can specify in your deployment.

Name

Specify a unique name for the deployment.

Type

Click **Star** to specify that you are configuring a star deployment.

Pre-shared Key

Define a unique pre-shared key for authentication.

Device

You can choose a managed device, including a device stack or device high-availability pair, as an endpoint for your deployment. For Cisco-managed devices not managed by the Firepower Management Center you are using, choose **Other** and then specify an IP address for the endpoint.

Virtual Router

If you chose a managed device as your endpoint, choose a virtual router that is currently applied to the selected device. You cannot choose the same virtual router for more than one endpoint.

Interface

If you chose a managed device as your endpoint, choose a routed interface that is assigned to the selected virtual router.

IP Address

- If you chose a managed device as an endpoint, choose an IP address that is assigned to the specified routed interface.
- If the managed device is a device high-availability pair, you can choose only from a list of SFRP IP addresses.
- If you chose a managed device **not** managed by the Firepower Management Center, specify an IP address for the endpoint.

Protected Networks

Specify the networks in your deployment that are encrypted. Enter a subnet with CIDR block for each network.

Note that VPN endpoints cannot have the same IP address and that protected networks in a VPN endpoint pair cannot overlap. If a list of protected networks for an endpoint contains one or more IPv4 or IPv6 entry, the other endpoint's protected network must have at least one entry of the same type (i.e., IPv4 or IPv6). If it does not, then the other endpoint's IP address must be of the same type and must not overlap with the entries in the protected network. (Use /32 CIDR address blocks for IPv4 and /128 CIDR address blocks for IPv6). If both of these checks fail, the endpoint pair is invalid.

Internal IP

Check the check box if the endpoint resides behind a firewall with network address translation.

Public IP

If you checked the **Internal IP** check box, specify a public IP address for the firewall. If the endpoint is a responder, you must specify this value.

Public IKE Port

If you checked the **Internal IP** check box, specify a single numerical value from 1 to 65535 for the UDP port on the firewall that is being port-forwarded to the internal endpoint. If the endpoint is a responder and the port on the firewall being forwarded is not 500 or 4500, you must specify this value.

Related Topics

Configuring Star VPN Deployments, on page 1175

Mesh VPN Deployment Options

When configuring a mesh VPN deployment, you define a group of VPNs to link any two points for a given set of endpoints.

The following list describes the options you can specify in your deployment.

Name

Specify a unique name for the deployment.

Type

Click **Mesh** to specify that you are configuring a mesh deployment.

Pre-shared Key

Define a unique pre-shared key for authentication.

Device

You can choose a managed device, including a device stack or device high-availability pair, as an endpoint for your deployment. For Cisco-managed devices not managed by the Firepower Management Center you are using, choose **Other** and then specify an IP address for the endpoint.

Virtual Router

If you chose a managed device as your endpoint, choose a virtual router that is currently applied to the specified device. You cannot choose the same virtual router for more than one endpoint.

Interface

If you chose a managed device as your endpoint, choose a routed interface that is assigned to the specified virtual router.

IP Address

- If you chose a managed device as an endpoint, choose an IP address that is assigned to the selected routed interface.
- If the managed device is a device high-availability pair, you can choose only from a list of SFRP IP addresses.
- If you chose a managed device **not** managed by the Firepower Management Center, specify an IP address for the endpoint.

Protected Networks

Specify the networks in your deployment that are encrypted. Enter a subnet with CIDR block for each network. IKE version 1 only supports a single protected network.

Note that VPN endpoints cannot have the same IP address and that protected networks in a VPN endpoint pair cannot overlap. If a list of protected networks for an endpoint contains one or more IPv4 or IPv6 entry, the other endpoint's protected network must have at least one entry of the same type (i.e., IPv4 or IPv6). If it does not, then the other endpoint's IP address must be of the same type and must not overlap with the entries in the protected network. (Use /32 CIDR address blocks for IPv4 and /128 CIDR address blocks for IPv6). If both of these checks fail, the endpoint pair is invalid.

Internal IP

Check the check box if the endpoint resides behind a firewall with network address translation.

Public IP

If you checked the **Internal IP** check box, specify a public IP address for the firewall. If the endpoint is a responder, you must specify this value.

Public IKE Port

If you checked the **Internal IP** check box, specify a single numerical value from 1 to 65535 for the UDP port on the firewall that is being port-forwarded to the internal endpoint. If the endpoint is a responder and the port on the firewall being forwarded is not 500 or 4500, you must specify this value.

Related Topics

Configuring Mesh VPN Deployments, on page 1176

Advanced VPN Deployment Options

VPN deployments contain some common settings that can be shared among the VPNs in a deployment. Each VPN can use the default settings or you can override the default settings. Advanced settings typically require little or no modification and are not common to every deployment.

The following list describes the advanced options you can specify in your deployment.

Other Algorithm Allowed

Check the check box to enable auto negotiation to an algorithm not listed in the Algorithm list, but proposed by the remote peer.

Algorithm

Specify the phase one and phase two algorithm proposals to secure data in your deployment. Choose **Cipher**, **Hash**, and Diffie-Hellman (**DH**) group authentication messages for both phases.

IKE Life Time

Specify a numerical value and choose a time unit for the maximum IKE SA renegotiation interval. You can specify a minimum of 15 minutes and a maximum of 30 days.

IKE v2

Check the check box to specify that the system uses IKE version 2. This version supports the star deployment and multiple protected networks.

Life Time

Specify a numerical value and select a time unit for the maximum SA renegotiation interval. You can specify a minimum of 5 minutes and a maximum of 24 hours.

Life Packets

Specify the number of packets that can be transmitted over an IPsec SA before it expires. You can use any integer between 0 and 18446744073709551615.

Life Bytes

Specify the number of bytes that can be transmitted over an IPsec SA before it expires. You can use any integer between 0 and 18446744073709551615.

AH

Check the check box to specify that the system uses the authentication header security protocol for the data to be protected. Clear the check box to use encryption service payload (ESP) protocol.

Related Topics

Configuring Advanced VPN Deployment Settings, on page 1177

Managing VPN Deployments

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	VPN	7000 & 8000 Series	Any	Admin/Network Admin

Caution Adding or removing a VPN on a 7000 or 8000 Series device restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Procedure

Step 1 Choose **Devices** > **VPN** > **Site to Site**.

Step 2 Manage your VPN deployments:

- Add — To create a new VPN deployment, click **Add VPN** > **Firepower Device**, and continue as follows depending on deployment type:
  - Configuring Mesh VPN Deployments, on page 1176
  - Configuring Point-to-Point VPN Deployments, on page 1175
  - Configuring Star VPN Deployments, on page 1175
- Edit — To modify the settings in an existing VPN deployment, click the edit icon ( ); see Editing VPN Deployments, on page 1178.
- Delete — To delete a VPN deployment, click the delete icon ( ).
- Deploy—Click **Deploy**; see Deploy Configuration Changes, on page 346.
- View VPN status — To view the status of an existing VPN deployment, click the status icon; see Viewing VPN Status, on page 1179.

Related Topics

Snort® Restart Scenarios, on page 348

Configuring Point-to-Point VPN Deployments

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	VPN	7000 & 8000 Series	Any	Admin/Network Admin

Before you begin

If you are using managed devices as endpoints, create a virtual router and apply it to the appropriate device.

Note You cannot use the same virtual router for more than one endpoint. For more information, see Setting Up Virtual Routers, on page 1115

Procedure

Step 1 Choose **Devices** > **VPN** > **Site to Site**.

Step 2 Click **Add VPN** > **Firepower Device**.

Step 3 Enter a unique **Name**.

Step 4 Verify that **PTP** is chosen as the **Type**.

Step 5 Enter a unique **Pre-shared Key**.

Step 6 Next to **Node Pairs**, click the add icon ( ).

Step 7 Configure the VPN deployment options described in Point-to-Point VPN Deployment Options, on page 1169.

Step 8 Under **Node A**, next to **Protected Networks**, click the add icon ( ).

Step 9 Enter a CIDR block for the protected network.

Step 10 Click **OK**.

Step 11 Repeat step 8 through step 10 for **Node B**.

Step 12 Click **Save**.
The endpoint pair is added to your deployment.

Step 13 Click **Save** to finish configuring your deployment.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configuring Star VPN Deployments

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	VPN	7000 & 8000 Series	Any	Admin/Network Admin

Before you begin

If you are using managed devices as endpoints, create a virtual router and apply it to the appropriate device.

Note You cannot use the same virtual router for more than one endpoint. For more information, see Setting Up Virtual Routers, on page 1115

Procedure

Step 1 Choose **Devices** > **VPN** > **Site to Site**.
Step 2 Click **Add VPN** > **Firepower Device**.
Step 3 Enter a unique **Name**.
Step 4 Click **Star** to specify the **Type**.
Step 5 Enter a unique **Pre-shared Key**.
Step 6 Next to **Hub Node**, click the edit icon ( ).
Step 7 Configure the VPN deployment options described in Star VPN Deployment Options, on page 1170.
Step 8 Next to **Protected Networks**, click the add icon ( ).
Step 9 Enter an IP address for the protected network.
Step 10 Click **OK**.
Step 11 Click **Save**. The hub node is added to your deployment.
Step 12 Next to **Leaf Nodes**, click the add icon ( ).
Step 13 Repeat step 7 through step 10 to complete the leaf node, which has the same options as the hub node.
Step 14 Click **Save**.
The leaf node is added to your deployment.
Step 15 Click **Save** to finish configuring your deployment.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configuring Mesh VPN Deployments

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	VPN	7000 & 8000 Series	Any	Admin/Network Admin

Before you begin

If you are using managed devices as endpoints, create a virtual router and apply it to the appropriate device.

Note You cannot use the same virtual router for more than one endpoint. For more information, see Setting Up Virtual Routers, on page 1115

Procedure

Step 1 Choose **Devices** > **VPN** > **Site to Site**.
Step 2 Click **Add VPN** > **Firepower Device**.
Step 3 Enter a unique **Name**.
Step 4 Click **Mesh** to specify the **Type**.
Step 5 Enter a unique **Pre-shared Key**.
Step 6 Next to **Nodes**, click the add icon ( ).
Step 7 Configure the VPN deployment options described in Mesh VPN Deployment Options, on page 1172.
Step 8 Next to **Protected Networks**, click the add icon ( ).
Step 9 Enter a CIDR block for the protected network.
Step 10 Click **OK**.
The protected network is added.
Step 11 Click **Save**.
The endpoint is added to your deployment.
Step 12 Repeat step 6 through step 11 to add more endpoints.
Step 13 Click **Save** to complete your deployment.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configuring Advanced VPN Deployment Settings

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	VPN	7000 & 8000 Series	Any	Admin/Network Admin

In a multidomain deployment, the system displays VPN deployments created in the current domain, which you can edit. It also displays VPN deployments created in ancestor domains if one of the endpoint devices belongs to your domain. You cannot edit VPN deployments created in ancestor domains. To view and edit VPN deployments created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Devices** > **VPN** > **Site to Site**.

Step 2 Click the edit icon ( ).

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click the **Advanced** tab.

Step 4 Configure the advanced settings, as described in Advanced VPN Deployment Options, on page 1173.

Step 5 Next to **Algorithms**, click the add icon ( ).

Step 6 Chose **Cipher**, **Hash**, and Diffie-Hellman (**DH**) group authentication messages for both phases.

Step 7 Click **OK**.

Step 8 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Editing VPN Deployments

Caution Two users should **not** edit the same deployment simultaneously; however, note that the web interface does not prevent simultaneous editing.

In a multidomain deployment, the system displays VPN deployments created in the current domain, which you can edit. It also displays VPN deployments created in ancestor domains if one of the endpoint devices belongs to your domain. You cannot edit VPN deployments created in ancestor domains. To view and edit VPN deployments created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Devices** > **VPN** > **Site to Site**.

Step 2 Click the edit icon ( ).

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Modify the desired settings:

- Advanced settings; see Configuring Advanced VPN Deployment Settings, on page 1177.
- Mesh deployment settings; see Configuring Mesh VPN Deployments, on page 1176.
- Point-to-point deployment settings; see Configuring Point-to-Point VPN Deployments, on page 1175.
- Star deployment settings; see Configuring Star VPN Deployments, on page 1175.

Tip You cannot edit the deployment type after you initially save the deployment. To change the deployment type, you must delete the deployment and create a new one.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

VPN Deployment Status

After you configure a VPN deployment, you can view the status of your configured VPN tunnels. The VPN page displays a status icon for each VPN deployment once it has been deployed:

- The () icon designates that all VPN endpoints are up.
- The () icon designates that all VPN endpoints are down.
- The () icon designates that some endpoints are up, while others are down.

You can click a status icon to view the deployment status along with basic information about the endpoints in the deployment, such as endpoint name and IP address. The VPN status updates every minute or when a status change occurs, such as an endpoint going down or coming up.

Related Topics

Viewing VPN Status, on page 1179

Viewing VPN Status

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	VPN	7000 & 8000 Series	Any	Admin/Network Admin

In a multidomain deployment, the system displays VPN deployments created in the current domain. It also displays VPN deployments created in ancestor domains if one of the endpoint devices belongs to your domain. To view VPN deployments created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Devices** > **VPN** > **Site to Site**.

Step 2 Click the VPN status icon next to the deployment where you want to view the status.

Step 3 Click **OK**.

VPN Statistics and Logs

After you configure a VPN deployment, you can view statistics about the data traversing your configured VPN tunnels. In addition, you can view the latest VPN system and IKE logs for each endpoint.

The system displays the following statistics:

Endpoint

The device path to the routed interface and IP address designated as the VPN endpoint.

Status

Whether the VPN connection is up or down.

Protocol

The protocol used for encryption, either ESP or AH.

Packets Received

The number of packets per interface the VPN tunnel receives during an IPsec SA negotiation.

Packets Forwarded

The number of packets per interface the VPN tunnel transmits during an IPsec SA negotiation.

Bytes Received

The number of bytes per interface the VPN tunnel receives during an IPsec SA negotiation.

Bytes Forwarded

The number of bytes per interface the VPN tunnel transmits during an IPsec SA negotiation.

Time Created

The date and time the VPN connection was created.

Time Last Used

The last time a user initiated a VPN connection.

NAT Traversal

If "Yes" is displayed, at least one of the VPN endpoints resides behind a device with network address translation.

IKE State

The state of the IKE SA: connecting, established, deleting, or destroying.

IKE Event

The IKE SA event: reauthentication or rekeying.

IKE Event Time

The time in seconds the next event should occur.

IKE Algorithm

The IKE algorithm being used by the VPN deployment.

IPsec State

The state of the IPsec SA: installing, installed, updating, rekeying, deleting, and destroying.

IPsec Event

Notification of when the IPsec SA event is rekeying.

IPsec Event Time

The time in seconds until the next event should occur.

IPsec Algorithm

IPsec algorithm being used by the VPN deployment.

Related Topics

Viewing VPN Statistics and Logs, on page 1181

Viewing VPN Statistics and Logs

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	VPN	7000 & 8000 Series	Any	Admin/Network Admin

In a multidomain deployment, the system displays VPN deployments created in the current domain. It also displays VPN deployments created in ancestor domains if one of the endpoint devices belongs to your domain. To view VPN deployments created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Devices** > **VPN** > **Site to Site**.

Step 2 Click the VPN status icon next to the deployment for which you want to view statistics.

Step 3 Click the view statistics icon ( ).

Step 4 Optionally, click **Refresh** to update the VPN statistics.

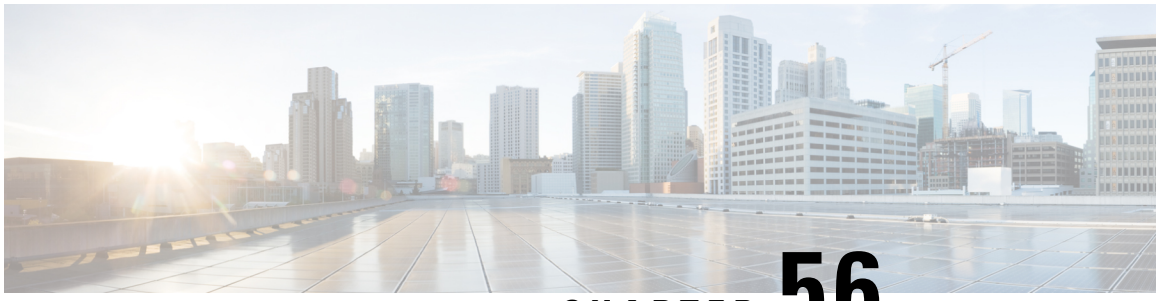
Step 5 Optionally, click **View Recent Log** to view the latest data log for each endpoint. To view the log for 7000 or 8000 Series devices in high-availability pairs and stacked devices, you can click the link for either the active/primary or backup/secondary device.



PART XV

Access Control

- Understanding Access Control, on page 1185
- Best Practices for Access Control, on page 1193
- Access Control Policies, on page 1201
- Access Control Rules, on page 1215
- URL Filtering, on page 1229
- HTTP Response Pages and Interactive Blocking, on page 1243
- Blocking Traffic with Security Intelligence, on page 1249
- DNS Policies, on page 1261
- Prefiltering and Prefilter Policies , on page 1273
- Intelligent Application Bypass, on page 1289
- Access Control Using Content Restriction, on page 1297



CHAPTER 56

Understanding Access Control

- Introduction to Access Control, on page 1185
- Access Control Policy Default Action, on page 1185
- Deep Inspection Using File and Intrusion Policies, on page 1187
- Access Control Policy Inheritance, on page 1191

Introduction to Access Control

Access control is a hierarchical policy-based feature that allows you to specify, inspect, and log (non-fast-pathed) network traffic.

Each managed device can be targeted by one access control policy. The data that the policy's *target devices* collect about your network traffic can be used to filter and control that traffic based on:

- simple, easily determined transport and network layer characteristics: source and destination, port, protocol, and so on
- the latest contextual information on the traffic, including characteristics such as reputation, risk, business relevance, application used, or URL visited
- realm, user, user group, or ISE attribute
- custom Security Group Tag (SGT)
- characteristics of encrypted traffic; you can also decrypt this traffic for further analysis
- whether unencrypted or decrypted traffic contains a prohibited file, detected malware, or intrusion attempt

Each type of traffic inspection and control occurs where it makes the most sense for maximum flexibility and performance. For example, reputation-based blocking uses simple source and destination data, so it can block prohibited traffic early in the process. In contrast, detecting and blocking intrusions and exploits is a last-line defense.

Access Control Policy Default Action

A newly created access control policy directs its target devices to handle all traffic using its *default action*.

In a simple access control policy, the default action specifies how target devices handle all traffic. In a more complex policy, the default action handles traffic that:

- is not trusted by Intelligent Application Bypass
- is not on a Security Intelligence Block list
- is not blocked by SSL inspection (encrypted traffic only)
- matches none of the rules in the policy (except Monitor rules, which match and log—but do not handle or inspect—traffic)

The access control policy default action can block or trust traffic without further inspection, or inspect traffic for intrusions and discovery data.

Note You **cannot** perform file or malware inspection on traffic handled by the default action. Logging for connections handled by the default action is initially disabled, though you can enable it.

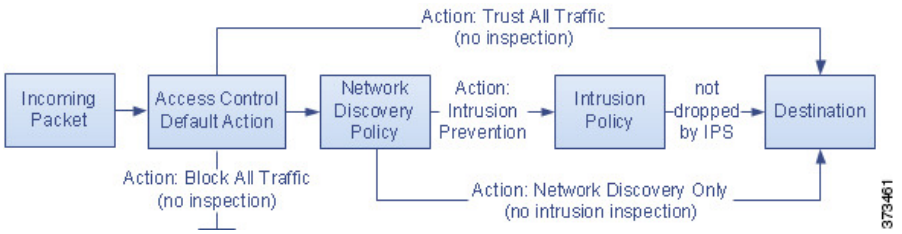
If you are using policy inheritance, the default action for the lowest-level descendant determines final traffic handling. Although an access control policy can inherit its default action from its base policy, you cannot enforce this inheritance.

The following table describes the types of inspection you can perform on traffic handled by each default action.

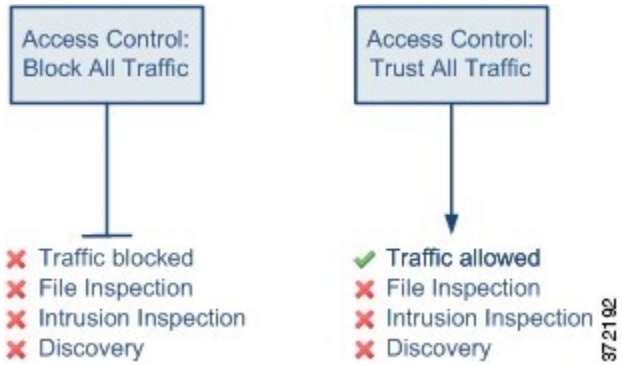
Table 86: Access Control Policy Default Actions

Default Action	Effect on Traffic	Inspection Type and Policy
Access Control: Block All Traffic	block without further inspection	none
Access Control: Trust All Traffic	trust (allow to its final destination without further inspection)	none
Intrusion Prevention	allow, as long as it is passed by the intrusion policy you specify	intrusion, using the specified intrusion policy and associated variable set, and discovery, using the network discovery policy
Network Discovery Only	allow	discovery only, using the network discovery policy
Inherit from base policy	defined in base policy	defined in base policy

The following diagram illustrates the table.



The following diagrams illustrate the **Block All Traffic** and **Trust All Traffic** default actions.



The following diagrams illustrate the **Intrusion Prevention** and **Network Discovery Only** default actions.



Tip The purpose of **Network Discovery Only** is to improve performance in a discovery-only deployment. Different configurations can disable discovery if you are only interested in intrusion detection and prevention.

Related Topics

Performance Considerations for Limited Deployments, on page 358
Logging Connections with a Policy Default Action, on page 2221

Deep Inspection Using File and Intrusion Policies

Deep inspection uses intrusion and file policies as the last line of defense before traffic is allowed to its destination.

- *Intrusion policies* govern the system's intrusion prevention capabilities.

  For complete information, see Intrusion Detection and Prevention, on page 1445.

- *File policies* govern the system's file control and AMP for Networks capabilities.

  For complete information, see File Policies and Malware Protection, on page 1401.

Access control occurs before deep inspection; access control rules and the access control default action determine which traffic is inspected by intrusion and file policies.

By associating an intrusion or file policy with an access control rule, you are telling the system that before it passes traffic that matches the access control rule's conditions, you first want to inspect the traffic with an intrusion policy, a file policy, or both.

In an access control policy, you can associate one intrusion policy with each Allow and Interactive Block rule, as well as with the default action. Every unique **pair** of intrusion policy and variable set counts as one policy.

To associate intrusion and file policies with an access control rule, see:

- Access Control Rule Configuration to Perform Intrusion Prevention, on page 1484
- Configuring an Access Control Rule to Perform Malware Protection, on page 1409

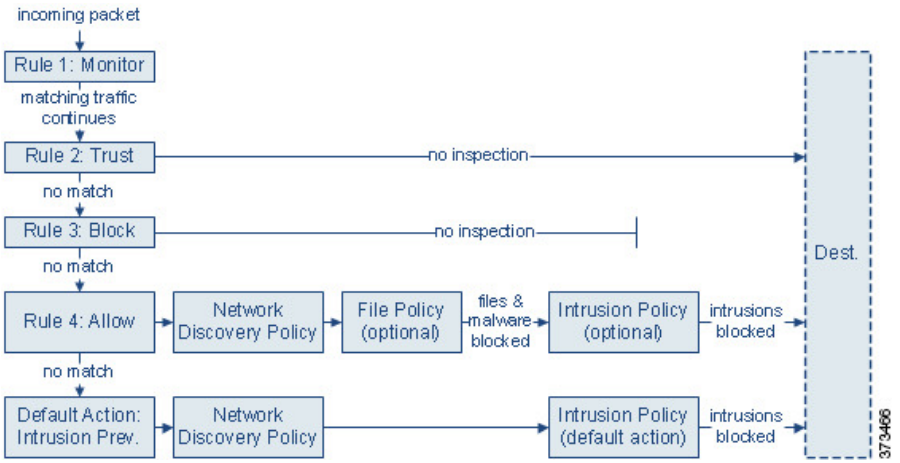
Note By default, the system disables intrusion and file inspection of encrypted payloads. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has intrusion and file inspection configured.

Related Topics

How Policies Examine Traffic For Intrusions, on page 1448
File Policies, on page 1401

Access Control Traffic Handling with Intrusion and File Policies

The following diagram shows the flow of traffic in an inline intrusion prevention and AMP for Networks deployment, as governed by an access control policy that contains four different types of access control rules and a default action.



In the scenario above, the first three access control rules in the policy—Monitor, Trust, and Block—cannot inspect matching traffic. Monitor rules track and log but do not inspect network traffic, so the system continues to match traffic against additional rules to determine whether to permit or deny it. (However, see an important exception and caveat at Access Control Rule Monitor Action, on page 1223.) Trust and Block rules handle matching traffic without further inspection of any kind, while traffic that does not match continues to the next access control rule.

The fourth and final rule in the policy, an Allow rule, invokes various other policies to inspect and handle matching traffic, in the following order:

- **Discovery: Network Discovery Policy**—First, the network discovery policy inspects traffic for discovery data. Discovery is passive analysis and does not affect the flow of traffic. Although you do not explicitly enable discovery, you can enhance or disable it. However, allowing traffic does not automatically guarantee discovery data collection. The system performs discovery only for connections involving IP addresses that are explicitly monitored by your network discovery policy.
- **AMP for Networks and File Control: File Policy**—After traffic is inspected by discovery, the system can inspect it for prohibited files and malware. AMP for Networks detects and optionally blocks malware in many types of files, including PDFs, Microsoft Office documents, and others. If your organization wants to block not only the transmission of malware files, but all files of a specific type (regardless of whether the files contain malware), *file control* allows you to monitor network traffic for transmissions of specific file types, then either block or allow the file.
- **Intrusion Prevention: Intrusion Policy**—After file inspection, the system can inspect traffic for intrusions and exploits. An intrusion policy examines decoded packets for attacks based on patterns, and can block or alter malicious traffic. Intrusion policies are paired with *variable sets*, which allow you to use named values to accurately reflect your network environment.
- **Destination**—Traffic that passes all the checks described above passes to its destination.

An Interactive Block rule (not shown in the diagram) has the same inspection options as an Allow rule. This is so you can inspect traffic for malicious content when a user bypasses a blocked website by clicking through a warning page.

Traffic that does not match any access control rules in the policy with an action other than Monitor is handled by the default action. In this scenario, the default action is an Intrusion Prevention action, which allows traffic to its final destination as long as it is passed by the intrusion policy you specify. In a different deployment, you might have a default action that trusts or blocks all traffic without further inspection. Note that the system can inspect traffic allowed by the default action for discovery data and intrusions, but not prohibited files or malware. You **cannot** associate a file policy with the access control default action.



Note Sometimes, when a connection is analyzed by an access control policy, the system must process the first few packets in that connection, **allowing them to pass**, before it can decide which access control rule (if any) will handle the traffic. However, so these packets do not reach their destination uninspected, you can specify an intrusion policy (in the Advanced settings for the access control policy) to inspect these packets and generate intrusion events.

File and Intrusion Inspection Order

In your access control policy, you can associate multiple Allow and Interactive Block rules with different intrusion and file policies to match inspection profiles to various types of traffic.



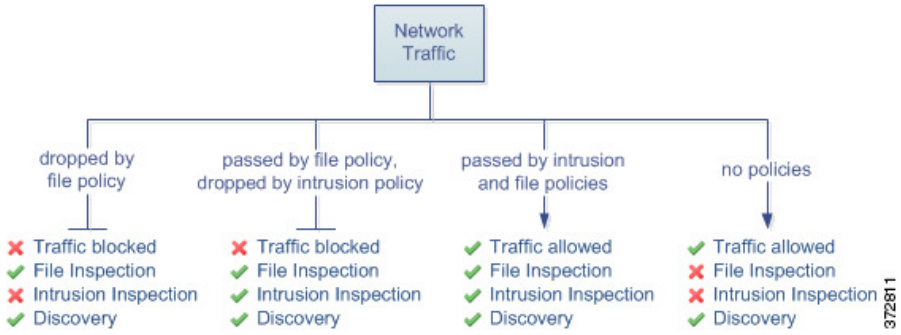
Note Traffic allowed by an Intrusion Prevention or Network Discovery Only default action can be inspected for discovery data and intrusions, but cannot be inspected for prohibited files or malware. You **cannot** associate a file policy with the access control default action.

You do not have to perform both file and intrusion inspection in the same rule. For a connection matching an Allow or Interactive Block rule:

- without a file policy, traffic flow is determined by the intrusion policy
- without an intrusion policy, traffic flow is determined by the file policy
- without either, allowed traffic is inspected by network discovery only

Tip The system does not perform any kind of inspection on trusted traffic. Although configuring an Allow rule with neither an intrusion nor file policy passes traffic like a Trust rule, Allow rules let you perform discovery on matching traffic.

The diagram below illustrates the types of inspection you can perform on traffic that meets the conditions of either an Allow or user-bypassed Interactive Block access control rule. For simplicity, the diagram displays traffic flow for situations where both (or neither) an intrusion and a file policy are associated with a single access control rule.



For any single connection handled by an access control rule, file inspection occurs before intrusion inspection. That is, the system does not inspect files blocked by a file policy for intrusions. Within file inspection, simple blocking by type takes precedence over malware inspection and blocking.

For example, consider a scenario where you normally want to allow certain network traffic as defined in an access control rule. However, as a precaution, you want to block the download of executable files, examine downloaded PDFs for malware and block any instances you find, and perform intrusion inspection on the traffic.

You create an access control policy with a rule that matches the characteristics of the traffic you want to provisionally allow, and associate it with both an intrusion policy and a file policy. The file policy blocks the download of all executables, and also inspects and blocks PDFs containing malware:

- First, the system blocks the download of all executables, based on simple type matching specified in the file policy. Because they are immediately blocked, these files are subject to neither malware nor intrusion inspection.
- Next, the system performs malware cloud lookups for PDFs downloaded to a host on your network. Any PDFs with a malware disposition are blocked, and are not subject to intrusion inspection.
- Finally, the system uses the intrusion policy associated with the access control rule to inspect any remaining traffic, including files not blocked by the file policy.

Note Until a file is detected and blocked in a session, packets from the session may be subject to intrusion inspection.

Access Control Policy Inheritance

Especially useful in multidomain deployments, you can nest access control policies, where each policy inherits the rules and settings from an ancestor (or *base*) policy. You can enforce this inheritance, or allow lower-level policies to override their ancestors.

Access control uses a hierarchical policy-based implementation. Just as you create a domain hierarchy, you can create a corresponding hierarchy of access control policies. A *descendant*, or *child*, access control policy inherits rules and settings from its direct *parent*, or base, policy. That base policy may have its own parent policy from which it inherits rules and settings, and so on.

An access control policy's rules are nested between its parent policy's Mandatory and Default rule sections. This implementation enforces Mandatory rules from ancestor policies, but allows the current policy to write rules that preempt Default rules from ancestor policies.

You can lock the following settings to enforce them in all descendant policies. Descendant policies can override unlocked settings.

- Security Intelligence — connections that are allowed or blocked based on the latest reputation intelligence for IP addresses, URLs, and domain names.
- HTTP Response pages — Displaying a custom or system-provided response page when you block a user's website request.
- Advanced settings — Specifying associated subpolicies, network analysis settings, performance settings, and other general options.

When using policy inheritance, the default action for the lowest-level descendant determines final traffic handling. Although an access control policy can inherit its default action from an ancestor policy, you cannot enforce this inheritance.

Policy Inheritance and Multitenancy

Access control's hierarchical policy-based implementation complements multitenancy.

In a typical multidomain deployment, access control policy hierarchy corresponds to domain structure, and you apply the lowest-level access control policy to managed devices. This implementation allows selective access control enforcement at a higher domain level, while lower-level domain administrators can tailor deployment-specific settings. (You must use roles, not policy inheritance and enforcement alone, to restrict administrators in descendant domains.)

For example, as a Global domain administrator for your organization, you can create an access control policy at the Global level. You can then require that all your devices, which are divided into subdomain by function, use that Global-level policy as a base policy.

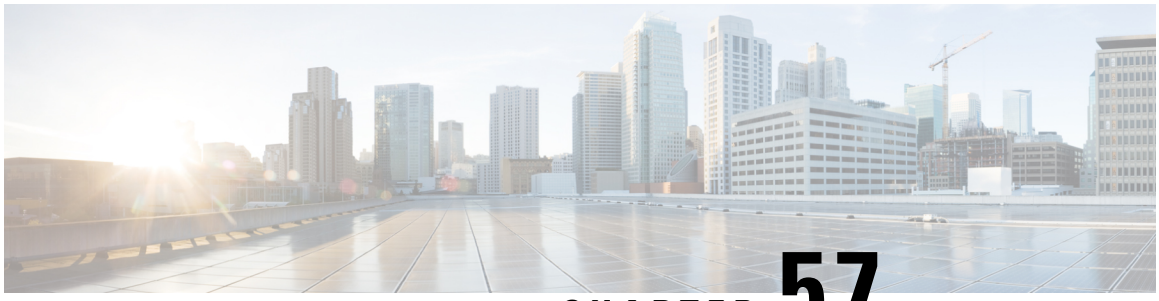
When subdomain administrators log into the Firepower Management Center to configure access control, they can deploy the Global-level policy as-is. Or, they can create and deploy a descendant access control policy within the boundaries of the Global-level policy.



Note Although the most useful implementation of access control inheritance and enforcement complements multitenancy, you can create a hierarchy of access control policies within a single domain. You can also assign and deploy access control policies at any level.

Related Topics

Managing Access Control Policy Inheritance, on page 1207
Blocking Traffic with Security Intelligence, on page 1249
HTTP Response Pages and Interactive Blocking, on page 1243
Access Control Policy Advanced Settings, on page 1211



CHAPTER 57

Best Practices for Access Control

- General Best Practices for Access Control, on page 1193
- Best Practices for Access Control Rules, on page 1194

General Best Practices for Access Control

Review the following requirements and general best practices:

- Use a prefilter policy to provide early blocking for unwanted traffic, and to fastpath traffic that does not benefit from access control inspection. For more information, see Best Practices for Prefiltering, on page 1276.
- Although you can configure the system without licensing your deployment, many features require that you enable the appropriate licenses before you deploy.
- For the system to affect traffic, you must deploy relevant configurations to managed devices using routed, switched, or transparent interfaces, or inline interface pairs.

  Sometimes, the system prevents you from deploying inline configurations to passively deployed devices, including inline devices in tap mode.

  In other cases, the policy may deploy successfully, but attempting to block or alter traffic using passively deployed devices can have unexpected results. For example, the system may report multiple beginning-of-connection events for each blocked connection, because blocked connections are not blocked in passive deployments.
- Certain features, including URL filtering, application detection, rate limiting, and Intelligent Application Bypass, must allow some packets to pass in order for the system to identify the traffic.

  To prevent these packets from reaching their destination uninspected, see Best Practices for Handling Packets That Pass Before Traffic Identification, on page 1670 and Specify a Policy to Handle Packets That Pass Before Traffic Identification, on page 1670.
- You cannot perform file or malware inspection on traffic handled by the access control policy's default action.
- Some features are only available on certain device models. Warning icons and confirmation dialog boxes designate unsupported features.

- If you will use syslog or store events externally, avoid special characters in object names such as policy and rule names. Object names should not contain special characters, such as commas, that the receiving application may use as separators.
- Logging for connections handled by the default action is initially disabled, though you can enable it.
- Best practices for creating, ordering, and implementing access control rules are detailed in Best Practices for Access Control Rules, on page 1194 and subtopics.

Best Practices for Access Control Rules

Properly configuring and ordering rules is essential to building an effective deployment. The following topics summarize rule performance guidelines.



Note When you deploy configuration changes, the system evaluates all rules together and creates an expanded set of criteria that target devices use to evaluate network traffic. If these criteria exceed the resources (physical memory, processors, and so on) of a target device, you cannot deploy to that device.

Related Topics

Best Practices for Application Control, on page 379
Best Practices for URL Filtering, on page 1230

Best Practices for Ordering Rules

General guidelines:

- In general, place top-priority rules that must apply to all traffic near the top of the policy.
- Specific rules should come before general rules, especially when the specific rules are exceptions to general rules.

  Otherwise, traffic will match the general rule first and never hit the applicable specific rule.
- Whenever possible, put specific drop rules near the top of the policy. This ensures the earliest possible decision on undesirable traffic.
- Rules that drop traffic based on layer-3/4 criteria only (such as IP address, security zone, and port number) should come as early as possible.
- URL filtering rules and application rules and others that require inspection should come after rules that drop traffic based on layer-3/4 criteria only (such as IP address, security zone, and port number), but before rules that specify file and intrusion policies.
- Put URL filtering rules above application rules, and follow application rules with micro-application rules and Common Industrial Protocol (CIP) sub-classification application filtering rules.
- Rules that specify file policies and intrusion policies should come at the bottom of the rule order. These rules require resource-intensive deep inspection, and you should eliminate as many threats as possible using less-intensive methods first, for performance reasons, in order to minimize the number of potential threats that require deep inspection.

- Always order rules to suit your organization's needs.

Exceptions and additions to the above guidelines are noted in the sections below.

Rule Preemption

Rule preemption occurs when a rule will never match traffic because a rule earlier in the evaluation order matches the traffic first. A rule's conditions govern whether it preempts other rules. In the following example, the second rule cannot block Admin traffic because the first rule allows it:

Access Control Rule 1: allow Admin users
Access Control Rule 2: block Admin users

Any type of rule condition can preempt a subsequent rule. The VLAN range in the first SSL rule includes the VLAN in the second rule, so the first rule preempts the second:

SSL Rule 1: do not decrypt VLAN 22-33
SSL Rule 2: block VLAN 27

In the following example, Rule 1 matches any VLAN because no VLANs are configured, so Rule 1 preempts Rule 2, which attempts to match VLAN 2:

Access Control Rule 1: allow Source Network 10.4.0.0/16
Access Control Rule 2: allow Source Network 10.4.0.0/16, VLAN 2

A rule also preempts an identical subsequent rule where all configured conditions are the same:

QoS Rule 1: rate limit VLAN 1 URL www.netflix.com
QoS Rule 2: rate limit VLAN 1 URL www.netflix.com

A subsequent rule would not be preempted if any condition is different:

QoS Rule 1: rate limit VLAN 1 URL www.netflix.com
QoS Rule 2: rate limit VLAN 2 URL www.netflix.com

Example: Ordering SSL Rules to Avoid Preemption

Consider a scenario where a trusted CA (Good CA) mistakenly issued a CA certificate to a malicious entity (Bad CA), but has not yet revoked that certificate. You want to use an SSL policy to block traffic encrypted with certificates issued by the untrusted CA, but otherwise allow traffic within the trusted CA's chain of trust. After you upload the CA certificates and all intermediate CA certificates, configure an SSL policy with rules in the following order:

SSL Rule 1: Block issuer CN=www.badca.com
SSL Rule 2: Do not decrypt issuer CN=www.goodca.com

If you reverse the rules, you first match all traffic trusted by Good CA, including traffic trusted by Bad CA. Because no traffic ever matches the subsequent Bad CA rule, malicious traffic may be allowed instead of blocked.

Rule Actions and Rule Order

A rule's action determines how the system handles matching traffic. Improve performance by placing rules that do not perform or ensure further traffic handling before the resource-intensive rules that do. Then, the system can divert traffic that it might otherwise have inspected.

The following examples show how you might order rules in various policies, given a set of rules where none is more critical and preemption is not an issue.

If your rules include application conditions, also see Best Practices for Configuring Application Control, on page 381.

Optimum Order: SSL Rules

Not only does decryption require resources, but so does further analysis of the decrypted traffic. Place SSL rules that decrypt traffic last.

1. Monitor—Rules that log matching connections, but take no other action on traffic.
2. Block, Block with reset—Rules that block traffic without further inspection.
3. Do not decrypt—Rules that do not decrypt encrypted traffic, passing the encrypted session to access control rules. The payloads of these sessions are not subject to deep inspection.
4. Decrypt - Known Key—Rules that decrypt incoming traffic with a known private key.
5. Decrypt - Resign—Rules that decrypt outgoing traffic by re-signing the server certificate.

Optimum Order: Access Control Rules

Intrusion, file, and malware inspection requires resources, especially if you use multiple custom intrusion policies and variable sets. Place access control rules that invoke deep inspection last.

1. Monitor—Rules that log matching connections, but take no other action on traffic. (However, see the important exception and caveat at Access Control Rule Monitor Action, on page 1223.)
2. Trust, Block, Block with reset—Rules that handle traffic without further inspection. Note that trusted traffic is subject to authentication requirements imposed by an identity policy, and to rate limiting.
3. Allow, Interactive Block (no deep inspection)—Rules that do not inspect traffic further, but allow discovery. Note that allowed traffic is subject to authentication requirements imposed by an identity policy, and to rate limiting.
4. Allow, Interactive Block (deep inspection)—Rules associated with file or intrusion policies that perform deep inspection for prohibited files, malware, and exploits.

Content Restriction Rule Order

To avoid rule preemption in both SSL and access control policies, position rules governing YouTube restriction above rules governing Safe Search restriction.

When you enable Safe Search for an access control rule, the system adds the `search engine` category to the **Selected Applications and Filters** list. This application category includes YouTube. As a result, YouTube traffic matches to the Safe Search rule unless YouTube EDU is enabled in a rule with a higher evaluation priority.

A similar rule preemption occurs if you position an SSL rule with the `safesearch supported` filter higher in the evaluation order than an SSL rule with specific YouTube application conditions.

Related Topics

About Content Restriction, on page 1297

Application Rule Order

Rules with application conditions are more likely to match traffic if you move them to a lower order in your list of rules.

Access control rules that use *specific* conditions (such as networks and IP addresses) should be ordered *before* rules that use general conditions (such as applications). If you're familiar with the Open Systems Interconnect (OSI) model, use similar numbering in concept. Rules with conditions for layers 1, 2, and 3 (physical, data link, and network) should be ordered first in your access control rules. Conditions for layers 5, 6, and 7 (session, presentation, and application) should be ordered later in your access control rules. For more information about the OSI model, see this Wikipedia article.

For more information and an example, see Best Practices for Configuring Application Control, on page 381 and Best Practices for Application Control, on page 379.

SSL Rule Order

In general, order your rules with specific conditions (such as IP addresses and networks) *before* rules with general conditions (such as applications).

Allow Traffic from Certificate Pinned Sites

Some applications use a technique referred to as *TLS/SSL pinning* or *certificate pinning*, which embeds the fingerprint of the original server certificate in the application itself. As a result, if you configured a TLS/SSL rule with a **Decrypt - Resign** action, when the application receives a resigned certificate from a managed device, validation fails and the connection is aborted.

To confirm that TLS/SSL pinning is occurring, attempt to log in to a mobile application like Facebook. If a network connection error is displayed, log in using a web browser. (For example, you *cannot* log in to a Facebook mobile application but *can* log in to Facebook using Safari or Chrome.) You can use Firepower Management Center connection events as further proof of TLS/SSL pinning



Note TLS/SSL pinning is not limited to mobile applications.

To allow this traffic, configure an SSL rule with the **Do Not Decrypt** action to match the server certificate common name or distinguished name. In the SSL policy, order this rule before all **Decrypt - Resign** rules that also match the traffic. You can retrieve the pinned certificate from the client's browser after a successful connection to the website. You can also view the certificate from the logged connection event, regardless of whether the connection succeeded or failed.

Prioritize ClientHello Modifications

To prioritize ClientHello modifications, place rules that match on conditions that are available in the ClientHello message before rules that match on ServerHello or server Certificate conditions.

When a managed device processes an SSL handshake, it can modify the ClientHello message to increase the likelihood of decryption. For example, it may remove compression methods because the Firepower System cannot decrypt compressed sessions.

The system modifies ClientHello messages only if it can conclusively match them to an SSL rule with a **Decrypt - Resign** action. The first time the system detects an encrypted session to a new server, server Certificate data is not available for ClientHello processing, which can result in an undecrypted first session.

For subsequent connections from the same client, the system can match the ClientHello message conclusively to rules with server Certificate conditions and process the message to maximize decryption potential.

If you place rules that match on ServerHello or server Certificate conditions (certificate, distinguished names, certificate status, cipher suites, version) before rules that match on ClientHello conditions (zones, networks, VLAN tags, ports, users, applications, URL categories), you can preempt ClientHello modification and increase the number of undecrypted sessions.

Situation Where SSL Policy is Bypassed

The SSL policy is bypassed for any connections that match access control rules with actions of **Trust**, **Block**, or **Block with reset** if those rules:

- Use security zone, network, geolocation, and port only as the traffic matching criteria.
- Precede other rules that require inspection, such as rules that match connections based on application or URL, or allow rules that apply intrusion or file inspection.

URL Rule Order

For the most effective URL matching, place rules that include URL conditions before other rules, particularly if the URL rules are block rules and the other rules meet both of the following criteria:

- They include application conditions.
- The traffic to be inspected is encrypted.

If you configure exceptions to a rule, put the exception above the other rule.

Best Practices for Simplifying and Focusing Rules

Simplify: Do Not Overconfigure

If one condition is enough to match the traffic you want to handle, do not use two.

Minimize individual rule criteria. Use as few individual elements in rule conditions as possible. For example, in network conditions use IP address blocks rather than individual IP addresses.

Combining elements into objects does **not** improve performance. For example, using a network object that contains 50 individual IP addresses gives you only an organizational—not a performance—benefit over including those IP addresses in the condition individually.

For recommendations related to application detection, see Best Practices for Configuring Application Control, on page 381.

Focus: Narrowly Constrain Resource-Intensive Rules, Especially by Interface

As much as possible, use rule conditions to narrowly define the traffic handled by resource-intensive rules. Focused rules are also important because rules with broad conditions can match many different types of traffic, and can preempt later, more specific rules. Examples of resource-intensive rules include:

- SSL rules that decrypt traffic—Not only the decryption, but further analysis of the decrypted traffic, requires resources. Narrow focus, and where possible, block or choose not to decrypt encrypted traffic.

- Access control rules that invoke deep inspection—Intrusion, file, and malware inspection requires resources, especially if you use multiple custom intrusion policies and variable sets. Make sure you only invoke deep inspection where required.

For maximum performance benefit, constrain rules by interface. If a rule excludes all of a device's interfaces, that rule does not affect that device's performance.

Maximum Number of Access Control Rules and Intrusion Policies

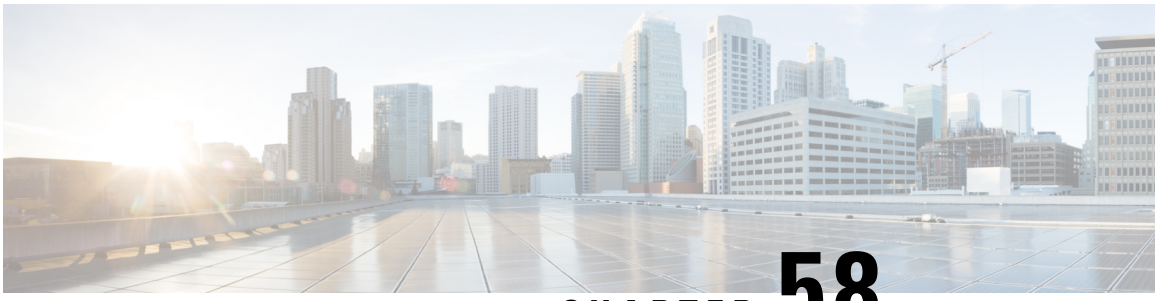
The maximum number of access control rules or intrusion policies that are supported by a target device depends on many factors, including policy complexity, physical memory, and the number of processors on the device.

If you exceed the maximum supported by your device, you cannot deploy your access control policy and must reevaluate.

Guidelines for intrusion policies:

In an access control policy, you can associate one intrusion policy with each Allow and Interactive Block rule, as well as with the default action. Every unique **pair** of intrusion policy and variable set counts as one policy.

You may want to consolidate intrusion policies or variable sets so you can associate a single intrusion policy-variable set pair with multiple access control rules. On some devices you may find you can use only a single variable set for all your intrusion policies, or even a single intrusion policy-variable set pair for the whole device.



CHAPTER 58

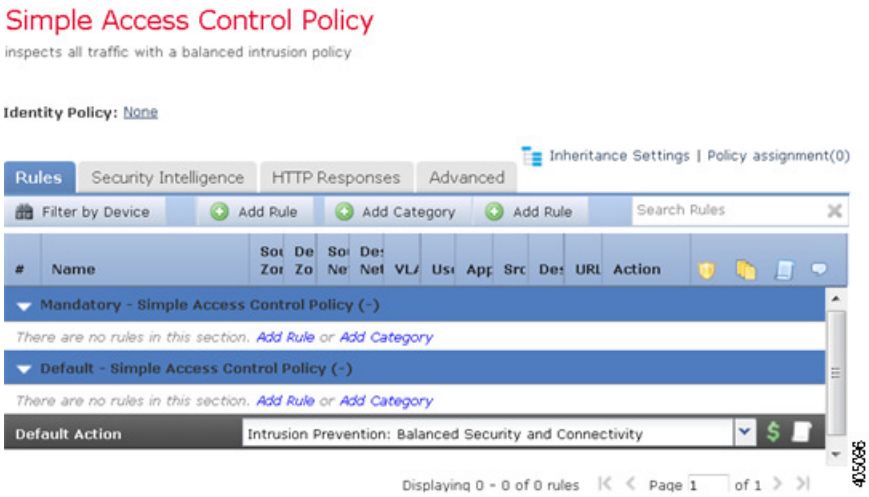
Access Control Policies

The following topics describe how to work with access control policies:

- Access Control Policy Components, on page 1201
- Requirements and Prerequisites for Access Control Policies, on page 1203
- Managing Access Control Policies, on page 1203
- System-Created Access Control Policies, on page 1204
- Creating a Basic Access Control Policy, on page 1204
- Editing an Access Control Policy, on page 1205
- Managing Access Control Policy Inheritance, on page 1207
- Setting Target Devices for an Access Control Policy, on page 1210
- Access Control Policy Advanced Settings, on page 1211
- History for Access Control Policies, on page 1214

Access Control Policy Components

In the following graphic, the default action uses the Balanced Security and Connectivity intrusion policy to inspect traffic before allowing it to its final destination.



Name and Description

Each access control policy must have a unique name. A description is optional.

Inheritance Settings

Policy inheritance allows you to create a hierarchy of access control policies. A parent (or *base*) policy defines and enforces default settings for its descendants, which is especially useful in multidomain deployments.

A policy's inheritance settings allow you to select its base policy. You can also lock settings in the current policy to force any descendants to inherit them. Descendant policies can override unlocked settings.

Policy Assignment

Each access control policy identifies the devices that use it. Each device can be targeted by only one access control policy. In a multidomain deployment, you can require that all the devices in a domain use the same base policy.

Rules

Access control rules provide a granular method of handling network traffic. Rules in an access control policy are numbered, starting at 1, including rules inherited from ancestor policies. The system matches traffic to access control rules in top-down order by ascending rule number.

Usually, the system handles network traffic according to the *first* access control rule where *all* the rule's conditions match the traffic. Conditions can be simple or complex, and their use often depends on certain licenses.

Default Action

The default action determines how the system handles and logs traffic that is not handled by any other access control configuration. The default action can block or trust all traffic without further inspection, or inspect traffic for intrusions and discovery data.

Although an access control policy can inherit its default action from an ancestor policy, you cannot enforce this inheritance.

Security Intelligence

Security Intelligence is a first line of defense against malicious internet content. This feature allows you to block connections based on the latest IP address, URL, and domain name reputation intelligence. To ensure continual access to vital resources, you can override Block list entries with custom Do Not Block list entries.

HTTP Responses

When the system blocks a user's website request, you can either display a generic system-provided response page, or a custom page. You can also display a page that warns users, but also allows them to continue to the originally requested site.

Advanced Access Control Options

Advanced access control policy settings typically require little or no modification. Often, the default settings are appropriate. Advanced settings you can modify include traffic preprocessing, SSL inspection, identity, and various performance options.

Related Topics

Rule Management: Common Characteristics, on page 361

Requirements and Prerequisites for Access Control Policies

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Managing Access Control Policies

The Firepower System allows you to edit system-provided access control policies and create custom access control policies.

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Policies** > **Access Control** .

Step 2 Manage access control policies:

- Copy—Click **Copy** ( ).
- Create—Click **New Policy**; see Creating a Basic Access Control Policy, on page 1204.
- Delete—Click **Delete** ( ).
- Deploy—Click **Deploy**; see Deploy Configuration Changes, on page 346.
- Edit—Click **Edit** ( ); see Editing an Access Control Policy, on page 1205
- Inheritance—Click **Plus** next to a policy with desdendants to expand your view of the policy's hierarchy.
- Import/Export—Click **Import/Export**; see Configuration Import and Export, on page 173.

- Report—Click **Report** (); see Generating Current Policy Reports, on page 356.

Related Topics

Out-of-Date Policies, on page 357

System-Created Access Control Policies

Depending on your devices' initial configurations, system-provided policies can include:

- Default Access Control—Blocks all traffic without further inspection.
- Default Intrusion Prevention—Allows all traffic, but also inspects with the Balanced Security and Connectivity intrusion policy and default intrusion variable set.
- Default Network Discovery—Allows all traffic while inspecting it for discovery data but not intrusions or exploits.

Creating a Basic Access Control Policy

When you create a new access control policy, you must, at minimum, choose a default action.

In most cases, logging of connections handled by a default action is initially disabled. An exception occurs if you create a subpolicy in a multidomain deployment. In that case, the system enables connection logging according to the logging configuration of the inherited default action.

Before you begin

Make sure you've addressed the steps up to this point in Setting Up Basic Policies and Configurations, on page 4.

Procedure

Step 1 Choose **Policies** > **Access Control** .

Step 2 Click **New Policy**.

Step 3 Enter a unique **Name** and, optionally, a **Description**.

Step 4 Optionally, choose a base policy from the **Select Base Policy** drop-down list.

If an access control policy is enforced on your domain, this step is not optional. You must choose the enforced policy or one of its descendants as the base policy.

Step 5 Specify the initial **Default Action**:

- If you chose a base policy, your new policy inherits its default action. You cannot change it here.
- **Block all traffic** creates a policy with the **Access Control: Block All Traffic** default action.
- **Intrusion Prevention** creates a policy with the **Intrusion Prevention: Balanced Security and Connectivity** default action, associated with the default intrusion variable set.
- **Network Discovery** creates a policy with the **Network Discovery Only** default action.

Tip If you want to trust all traffic by default, or if you chose a base policy and do not want to inherit the default action, you can change the default action later.

Caution Changing the total number of intrusion policies used by an access control policy restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information. You change the the total number of intrusion policies by adding an intrusion policy that is not currently used, or by removing the last instance of an intrusion policy. You can use an intrusion policy in an access control rule, as the default action, or as the default intrusion policy.

Step 6 Optionally, choose the **Available Devices** where you want to deploy the policy, then click **Add to Policy** (or drag and drop) to add the selected devices. To narrow the devices that appear, type a search string in the **Search** field.

If you want to deploy this policy immediately, you must perform this step.

Step 7 Click **Save**.

What to do next

- Optionally, further configure the new policy as described in Editing an Access Control Policy, on page 1205.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Note When you deploy an access control policy, its rules are not applied on the existing tunnel sessions. Hence, traffic on an existing connection is not bound by the new policy that is deployed. In addition, the policy hit count is incremented only for the first packet of a connection that matches a policy. Thus, the traffic of an existing connection that could match a policy is omitted from the hit count. To have the policy rules effectively applied, clear the existing tunnel sessions, and then deploy the policy.

Related Topics

Access Control Policy Default Action, on page 1185
Setting Target Devices for an Access Control Policy, on page 1210

Editing an Access Control Policy

Only one person should edit a policy at a time, using a single browser window. If multiple users save the same policy, the last saved changes are retained. For your convenience, the system displays information on who (if anyone) is currently editing each policy. To protect the privacy of your session, a warning appears after 30 minutes of inactivity on the policy editor. After 60 minutes, the system discards your changes.

Note You can only edit access control policies that were created in the current domain. Also, you cannot edit settings that are locked by an ancestor access control policy.

Procedure

Step 1 Choose **Policies** > **Access Control** .

Step 2 Click **Edit** ( ) next to the access control policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Edit your access control policy.

Settings:

- Name and Description—Click either field and enter new information.
- Default Action—Choose a value from the **Default Action** drop-down list.

Caution Changing the total number of intrusion policies used by an access control policy restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information. You change the the total number of intrusion policies by adding an intrusion policy that is not currently used, or by removing the last instance of an intrusion policy. You can use an intrusion policy in an access control rule, as the default action, or as the default intrusion policy.

- Default Action Variable Set—To change the variable set associated with an **Intrusion Prevention** default action, click **Variables** ( ). In the popup window that appears, select a new variable set and click **OK**.

  You can also click **Edit** ( ) to edit the selected variable set in a new window. For more information, see Managing Variables, on page 424.
- Default Action Logging—To configure logging for connections handled by the default action, click **Logging** ( ); see Logging Connections with a Policy Default Action, on page 2221.
- HTTP Responses—To specify what the user sees in a browser when the system blocks a website request, click **HTTP Responses**; see Choosing HTTP Response Pages, on page 1245.
- Inheritance: Change Base Policy—To change the base access control policy for this policy, click **Inheritance Settings**; see Choosing a Base Access Control Policy, on page 1208.
- Inheritance: Lock Settings in Descendants—To enforce this policy's settings in its descendant policies, click **Inheritance Settings**; see Locking Settings in Descendant Access Control Policies, on page 1209.
- Policy Assignment: Targets—To identify the managed devices targeted by this policy, click **Policy Assignment**; see Setting Target Devices for an Access Control Policy, on page 1210.
- Policy Assignment: Required in Domains—To enforce this policy in a subdomain, click **Policy Assignment**; see Requiring an Access Control Policy in a Domain, on page 1209.
- Rules—To manage access control rules, and to inspect and block malicious traffic using intrusion and file policies, click **Rules**; see Create and Edit Access Control Rules, on page 1220.
- Security Intelligence—To immediately block connections based on the latest reputation intelligence using a Block list, click **Security Intelligence**; see Configure Security Intelligence, on page 1252.

- Advanced Options—To set preprocessing, SSL inspection, identity, performance, and other advanced options, click **Advanced**; see Access Control Policy Advanced Settings, on page 1211.
- Warnings—To view a list of warnings or errors in your access control policy (and its descendant and associated policies), click **Show Warnings**. Warnings and errors mark configurations that could adversely affect traffic analysis and flow or prevent the policy from deploying. If there are no warnings, show warnings does not appear.

Step 4 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Rule and Other Policy Warnings, on page 392
Deep Inspection Using File and Intrusion Policies, on page 1187

Managing Access Control Policy Inheritance

Before you begin

Understand how inheritance works. See Access Control Policy Inheritance, on page 1191 and subtopics.

Procedure

Step 1 Edit the access control policy whose inheritance settings you want to change; see Editing an Access Control Policy, on page 1205.

Step 2 Manage policy inheritance:

- Change Base Policy — To change the base access control policy for this policy, click **Inheritance Settings** and proceed as described in Choosing a Base Access Control Policy, on page 1208.
- Lock Settings in Descendants — To enforce this policy's settings in its descendant policies, click **Inheritance Settings** and proceed as described in Locking Settings in Descendant Access Control Policies, on page 1209 .
- Required in Domains — To enforce this policy in a subdomain, click **Policy Assignment** and proceed as described in Requiring an Access Control Policy in a Domain, on page 1209.
- Inherit Settings from Base Policy — To inherit settings from a base access control policy, click **Security Intelligence**, **HTTP Responses**, or **Advanced** and proceed as directed in Inheriting Access Control Policy Settings from the Base Policy, on page 1208.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Choosing a Base Access Control Policy

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Any	Any	Any	Admin/Access Admin/Network Admin

You can use one access control policy as the base (parent) for another. By default, a child policy inherits its settings from its base policy, though you can change unlocked settings.

When you change the base policy for the current access control policy, the system updates the current policy with any locked settings from the new base policy.

Procedure

Step 1 In the access control policy editor, click **Inheritance Settings**.

Step 2 Choose a policy from the **Select Base Policy** drop-down list.

In a multidomain deployment, an access control policy may be required in the current domain. You can choose only the enforced policy or one of its descendants as the base policy.

Step 3 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Inheriting Access Control Policy Settings from the Base Policy

A new child policy inherits many settings from its base policy. If these settings are unlocked in the base policy, you can override them.

If you later reinherit the settings from the base policy, the system displays the base policy's settings and dims the controls. However, the system saves the overrides you made, and restores them if you disable inheritance again.

Procedure

Step 1 In the access control policy editor, click **Security Intelligence**, **HTTP Responses**, or **Advanced**.

Step 2 Check the **Inherit from base policy** check box for each setting you want to inherit.

If the controls are dimmed, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration.

Step 3 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Locking Settings in Descendant Access Control Policies

Lock a setting in an access control policy to enforce the setting in all descendant policies. Descendant policies can override unlocked settings.

When you lock settings, the system saves overrides already made in descendant polices so that the overrides can be restored if you unlock settings again.

Procedure

Step 1 In the access control policy editor, click **Inheritance Settings**.

Step 2 In the Child Policy Inheritance Settings area, check the settings you want to lock.

If the controls are dimmed, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration.

Step 3 Click **OK** to save the inheritance settings.

Step 4 Click **Save** to save the access control policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Requiring an Access Control Policy in a Domain

You can require that every device in a domain use the same base access control policy or one of its descendant policies.

Before you begin

- Configure at least one domain other than the Global domain.

Procedure

Step 1 In the access control policy editor, click **Policy Assignments**.

Step 2 Click **Required on Domains**.

Step 3 Build your domain list:

- Add — Select the domains where you want to enforce the current access control policy, then click **Add** or drag and drop into the list of selected domains.

- Delete — Click **Delete** ( ) next to a leaf domain, or right-click an ancestor domain and choose **Delete Selected**.
- Search — Type a search string in the search field. Click **Clear** ( ) to clear the search.

Step 4 Click **OK** to save the domain enforcement settings.

Step 5 Click **Save** to save the access control policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Setting Target Devices for an Access Control Policy

An access control policy specifies the devices that use it. Each device can be targeted by only one access control policy. In multidomain deployments, you can require that all the devices in a domain use the same base policy.

Procedure

Step 1 In the access control policy editor, click **Policy Assignments**.

Step 2 On **Targeted Devices**, build your target list:

- Add — Select one or more **Available Devices**, then click **Add to Policy** or drag and drop into the list of **Selected Devices**.
- Delete — Click **Delete** ( ) next to a single device, or select multiple devices, right-click, then choose **Delete Selected**.
- Search — Type a search string in the search field. Click **Clear** ( ) to clear the search.

Under **Impacted Devices**, the system lists the devices whose assigned access control policies are children of the current policy. Any change to the current policy affects these devices.

Step 3 Optionally, click **Required on Domains** to require that all the devices in the subdomains you choose use the same base policy. See Requiring an Access Control Policy in a Domain, on page 1209.

Step 4 Click **OK** to save your targeted device settings.

Step 5 Click **Save** to save the access control policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Access Control Policy Advanced Settings

Advanced access control policy settings typically require little or no modification. The default settings are appropriate for most deployments. Note that many of the advanced preprocessing and performance options in access control policies may be modified by rule updates as described in Update Intrusion Rules, on page 143.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Caution See Configurations that Restart the Snort Process When Deployed or Activated, on page 352 for a list of advanced setting modifications that restart the Snort process, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

General Settings

Option	Description
Maximum URL characters to store in connection events	To customize the number of characters you store for each URL requested by your users, see Limiting Logging of Long URLs, on page 2222. To customize the length of time before you re-block a website after a user bypasses an initial block, see Setting the User Bypass Timeout for a Blocked Website, on page 1246.
Allow an Interactive Block to bypass blocking for (seconds)	See Setting the User Bypass Timeout for a Blocked Website, on page 1246.
Retry URL cache miss lookup	The first time the system encounters a URL that does not have a locally stored category and reputation, it looks up that URL in the cloud and adds the result to the local data store, for faster processing of that URL in future. This setting determines what the system does when it needs to look up a URL's category and reputation in the cloud. By default, this setting is enabled: The system momentarily delays the traffic while it checks the cloud for the URL's reputation and category, and uses the cloud verdict to handle the traffic. If you disable this setting: When the system encounters a URL that is not in its local cache, the traffic is immediately passed and handled according to the rules configured for Uncategorized traffic. In passive deployments, the system does not retry the lookup, as it cannot hold packets.

Option	Description
Inspect traffic during policy apply	To inspect traffic when you deploy configuration changes unless specific configurations require restarting the Snort process, ensure that **Inspect traffic during policy apply** is set to its default value (enabled). When this option is enabled, resource demands could result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Scenarios, on page 348 for more information.

Associated Policies

Use advanced settings to associate subpolicies (SSL, identity, prefilter) with access control; see Associating Other Policies with Access Control, on page 1213.

Network Analysis and Intrusion Policies

Advanced network analysis and intrusion policy settings allow you to:

- Specify the intrusion policy and associated variable set that are used to inspect packets that must pass before the system can determine exactly how to inspect that traffic.
- Change the access control policy's default network analysis policy, which governs many preprocessing options.
- Use custom network analysis rules and network analysis policies to tailor preprocessing options to specific security zones, networks, and VLANs.

For more information, see Advanced Access Control Settings for Network Analysis and Intrusion Policies, on page 1669.

File and Malware Settings

File and Malware Inspection Performance and Storage Tuning, on page 1441 provides information on performance options for file control and AMP for Networks.

Intelligent Application Bypass Settings

Intelligent Application Bypass (IAB) is an expert-level configuration that specifies applications to bypass or test for bypass if traffic exceeds a combination of inspection performance and flow thresholds. For more information, see Intelligent Application Bypass, on page 1289.

Transport/Network Layer Preprocessor Settings

Advanced transport and network preprocessor settings apply globally to all networks, zones, and VLANs where you deploy your access control policy. You configure these advanced settings in an access control policy rather than in a network analysis policy. For more information, see Advanced Transport/Network Preprocessor Settings, on page 1758.

Detection Enhancement Settings

Advanced detection enhancement settings allow you to configure adaptive profiles so you can:

- Use file policies and applications in access control rules.
- Use service metadata in intrusion rules.
- In passive deployments, improve reassembly of packet fragments and TCP streams based on your network's host operating systems.

For more information, see Adaptive Profiles, on page 1811.

Performance Settings and Latency-Based Performance Settings

About Intrusion Prevention Performance Tuning, on page 1655 provides information on improving the performance of your system as it analyzes traffic for attempted intrusions.

For information specific to latency-based performance settings, see Packet and Intrusion Rule Latency Threshold Configuration, on page 1659.

Associating Other Policies with Access Control

Use an access control policy's advanced settings to associate one of each of the following subpolicies with the access control policy:

- SSL policy—Monitors, decrypts, blocks, or allows application layer protocol traffic encrypted with Secure Socket Layer (SSL) or Transport Layer Security (TLS).

Caution Adding or removing an SSL policy restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

- Identity policy—Performs user authentication based on the realm and authentication method associated with the traffic.
- Prefilter policy—Performs early traffic handling using limited network (layer 4) outer-header criteria.

Procedure

Step 1 In the access control policy editor, click **Advanced Settings**.

Step 2 Click **Edit** ( ) in the appropriate Policy Settings area.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Choose a policy from the drop-down list.

If you choose a user-created policy, you can click edit that appears to edit the policy.

Step 4 Click **OK**.

Step 5 Click **Save** to save the access control policy.

What to do next

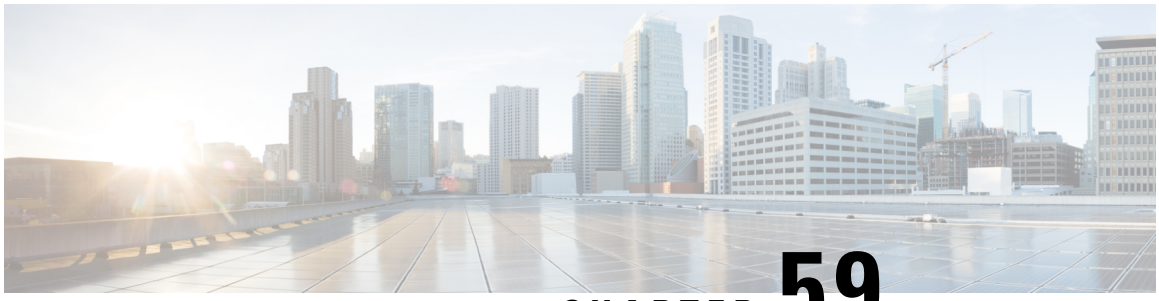
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Snort® Restart Scenarios, on page 348

History for Access Control Policies

Feature	Version	Details
New Security Intelligence categories	--	The following categories were introduced at about the time of the 6.6 release, but are not specific to 6.6: • banking_fraud • high_risk • ioc • link_sharing • malicious • newly_seen • spyware New/modified pages: Access control policy > Security Intelligence tab. Supported platforms: FMC



CHAPTER 59

Access Control Rules

The following topics describe how to configure access control rules:

- Introduction to Access Control Rules, on page 1215
- Requirements and Prerequisites for Access Control Rules, on page 1219
- Adding an Access Control Rule Category, on page 1220
- Create and Edit Access Control Rules, on page 1220
- Enabling and Disabling Access Control Rules, on page 1222
- Positioning an Access Control Rule, on page 1223
- Access Control Rule Actions, on page 1223
- Access Control Rule Comments, on page 1226

Introduction to Access Control Rules

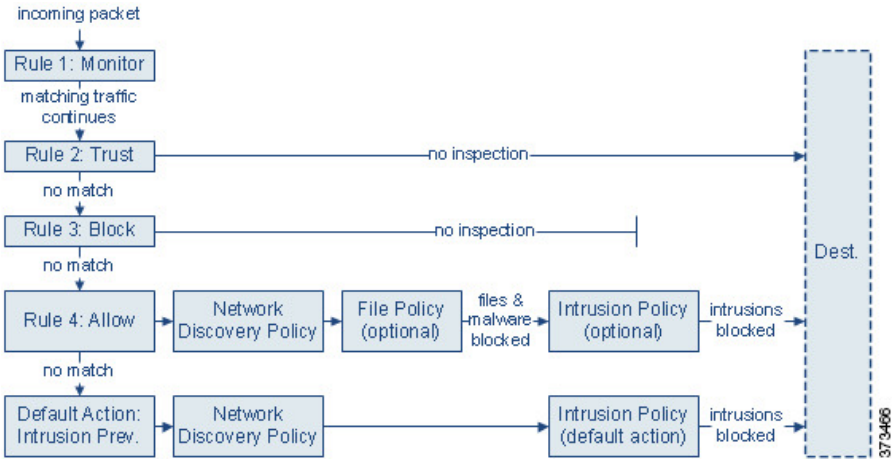
Within an access control policy, *access control rules* provide a granular method of handling network traffic across multiple managed devices.

Note Prefilter evaluation/8000 series fastpathing, Security Intelligence filtering, SSL inspection, user identification, and some decoding and preprocessing occur before access control rules evaluate network traffic.

The system matches traffic to access control rules in the order you specify. In most cases, the system handles network traffic according to the *first* access control rule where *all* the rule's conditions match the traffic.

Each rule also has an *action*, which determines whether you monitor, trust, block, or allow matching traffic. When you allow traffic, you can specify that the system first inspect it with intrusion or file policies to block any exploits, malware, or prohibited files before they reach your assets or exit your network.

The following scenario summarizes the ways that traffic can be evaluated by access control rules in an inline, intrusion prevention deployment.



In this scenario, traffic is evaluated as follows:

- **Rule 1: Monitor** evaluates traffic first. Monitor rules track and log network traffic. The system continues to match traffic against additional rules to determine whether to permit or deny it. (However, see an important exception and caveat at Access Control Rule Monitor Action, on page 1223.)
- **Rule 2: Trust** evaluates traffic next. Matching traffic is allowed to pass to its destination without further inspection, though it is still subject to identity requirements and rate limiting. Traffic that does not match continues to the next rule.
- **Rule 3: Block** evaluates traffic third. Matching traffic is blocked without further inspection. Traffic that does not match continues to the final rule.
- **Rule 4: Allow** is the final rule. For this rule, matching traffic is allowed; however, prohibited files, malware, intrusions, and exploits within that traffic are detected and blocked. Remaining non-prohibited, non-malicious traffic is allowed to its destination, though it is still subject to identity requirements and rate limiting. You can configure Allow rules that perform only file inspection, or only intrusion inspection, or neither.
- **Default Action** handles all traffic that does not match any of the rules. In this scenario, the default action performs intrusion prevention before allowing non-malicious traffic to pass. In a different deployment, you might have a default action that trusts or blocks all traffic, without further inspection. (You cannot perform file or malware inspection on traffic handled by the default action.)

Traffic you allow, whether with an access control rule or the default action, is automatically eligible for inspection for host, application, and user data by the network discovery policy. You do not explicitly enable discovery, although you can enhance or disable it. However, allowing traffic does not automatically guarantee discovery data collection. The system performs discovery only for connections involving IP addresses that are explicitly monitored by your network discovery policy; additionally, application discovery is limited for encrypted sessions.

Note that access control rules handle encrypted traffic when your SSL inspection configuration allows it to pass, or if you do not configure SSL inspection. However, some access control rule conditions require unencrypted traffic, so encrypted traffic may match fewer rules. Also, by default, the system disables intrusion and file inspection of encrypted payloads. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has intrusion and file inspection configured.

Access Control Rule Management

The **Rules** tab of the access control policy editor allows you to add, edit, categorize, search, move, enable, disable, delete, and otherwise manage access control rules in the current policy.

For each access control rule, the policy editor displays its name, a summary of its conditions, the rule action, and icons that communicate the rule's inspection options or status. These icons represent:

- **Intrusion policy** ( )
- **File policy** ( )
- **Safe search** ( )
- **YouTube EDU** ( )
- **Logging** ( )
- **Original Client option**
- **Comment** ( )
- **Warning** ( )
- **Errors** ( )
- important **Information** ( )

Disabled rules are dimmed and marked `(disabled)` beneath the rule name.

To create or edit a rule, use the access control rule editor. You can:

- Configure basic properties such as the rule's name, state, position, and action in the upper portion of the editor.
- Add conditions using the tabs on the left side of the lower portion of the editor.
- Use the tabs on the right side of the lower portion to configure inspection and logging options, and also to add comments to the rule. For your convenience, the editor lists the rule's inspection and logging options regardless of which tab you are viewing.

Note Properly creating and ordering access control rules is a complex task, but one that is essential to building an effective deployment. If you do not plan your policy carefully, rules can preempt other rules, require additional licenses, or contain invalid configurations. To help ensure that the system handles traffic as you expect, the access control policy interface has a robust warning and error feedback system for rules.

Related Topics

Access Control Rule Components, on page 1218
Example: Custom User Roles and Access Control, on page 44
Best Practices for Access Control Rules, on page 1194

Access Control Rule Components

In addition to its unique name, each access control rule has the following basic components:

State

By default, rules are enabled. If you disable a rule, the system does not use it and stops generating warnings and errors for that rule.

Position

Rules in an access control policy are numbered, starting at 1. If you are using policy inheritance, rule 1 is the first rule in the outermost policy. The system matches traffic to rules in top-down order by ascending rule number. With the exception of Monitor rules, the first rule that traffic matches is the rule that handles that traffic.

Rules can also belong to a section and a category, which are organizational only and do not affect rule position. Rule position goes across sections and categories.

Section and Category

To help you organize access control rules, every access control policy has two system-provided rule sections, Mandatory and Default. To further organize access control rules, you can create custom rule categories inside the Mandatory and Default sections.

If you are using policy inheritance, the current policy's rules are nested between its parent policy's Mandatory and Default sections.

Conditions

Conditions specify the specific traffic the rule handles. Conditions can be simple or complex; their use often depends on license.

Traffic must meet all of the conditions specified in all of the tabs in a rule. For example, if the Applications tab specifies HTTP but not HTTPS, the URL category and reputation conditions in the URLs tab will not apply to HTTPS traffic.

Action

A rule’s action determines how the system handles matching traffic. You can monitor, trust, block, or allow (with or without further inspection) matching traffic. The system does **not** perform deep inspection on trusted, blocked, or encrypted traffic.

Inspection

Deep inspection options govern how the system inspects and blocks malicious traffic you would otherwise allow. When you allow traffic with a rule, you can specify that the system first inspect it with intrusion or file policies to block any exploits, malware, or prohibited files before they reach your assets or exit your network.

Logging

A rule’s logging settings govern the records the system keeps of the traffic it handles. You can keep a record of traffic that matches a rule. In general, you can log sessions at the beginning or end of a connection, or both. You can log connections to the database, as well as to the system log (syslog) or to an SNMP trap server.

Comments

Each time you save changes to an access control rule, you can add comments.

Related Topics

Best Practices for Access Control Rules, on page 1194
Access Control Rule Management, on page 1217
Create and Edit Access Control Rules, on page 1220
Rule Condition Types, on page 363
Access Control Rule Actions, on page 1223
Deep Inspection Using File and Intrusion Policies, on page 1187
Best Practices for Connection Logging, on page 2216
Access Control Rule Comments, on page 1226

Access Control Rule Order

Rules in an access control policy are numbered, starting at 1. The system matches traffic to access control rules in top-down order by ascending rule number.

In most cases, the system handles network traffic according to the *first* access control rule where *all* the rule's conditions match the traffic. Except for Monitor rules, the system does not continue to evaluate traffic against additional, lower-priority rules after that traffic matches a rule.

To help you organize access control rules, every access control policy has two system-provided rule sections, Mandatory and Default. To further organize, you can create custom rule categories inside the Mandatory or Default sections. After you create a category, you cannot move it, although you can delete it, rename it, and move rules into, out of, within, and around it. The system assigns rule numbers across sections and categories.

If you use policy inheritance, the current policy's rules are nested between its parent policy's Mandatory and Default rule sections. Rule 1 is the first rule in the outermost policy, not the current policy, and the system assigns rule numbers across policies, sections, and categories.

Any predefined user role that allows you to modify access control policies also allows you to move and modify access control rules within and among rules categories. You can, however, create custom roles that restrict users from moving and modifying rules. Any user who is allowed to modify access control policies can add rules to custom categories and modify rules in them without restriction.

Tip Proper access control rule order reduces the resources required to process network traffic, and prevents rule preemption. Although the rules you create are unique to every organization and deployment, there are a few general guidelines to follow when ordering rules that can optimize performance while still addressing your needs.

Related Topics

Best Practices for Ordering Rules, on page 1194

Requirements and Prerequisites for Access Control Rules

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Adding an Access Control Rule Category

You can divide an access control policy's Mandatory and Default rule sections into custom categories. After you create a category, you cannot move it, although you can delete it, rename it, and move rules into, out of, within, and around it. The system assigns rule numbers across sections and categories.

Procedure

Step 1 In the access control policy editor, click **Add Category**.

Tip If your policy already contains rules, you can click a blank area in the row for an existing rule to set the position of the new category before you add it. You can also right-click an existing rule and select **Insert new category**.

Step 2 Enter a **Name**.

Step 3 From the **Insert** drop-down list, choose where you want to add the category:

- To insert a category below all existing categories in a section, choose **into Mandatory** or **into Default**.
- To insert a category above an existing category, choose **above category**, then choose a category.
- To insert a category above or below an access control rule, choose **above rule** or **below rule**, then enter an existing rule number.

Step 4 Click **OK**.

Step 5 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Create and Edit Access Control Rules

If you edit an access control rule that is actively in use, the changes do not apply to established connections at deploy-time. The updated rule is used to match against future connections. However, if the system is actively

inspecting a connection (for example, with an intrusion policy), it *will* apply changed matching or action criteria to existing connections.

For Firepower Threat Defense, you can ensure that your changes apply to all current connections by using the FTD **clear conn** CLI command to end established connections. Note that you should only do this if it is OK to end those connections, on the assumption that the sources for the connections will then attempt to reestablish the connection and thus be matched appropriately against the new rule.

Caution Changing the total number of intrusion policies used by an access control policy restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.You change the total number of intrusion policies by adding an intrusion policy that is not currently used, or by removing the last instance of an intrusion policy. You can use an intrusion policy in an access control rule, as the default action, or as the default intrusion policy.

Procedure

Step 1 In the access control policy editor, you have the following options:

- To add a new rule, click **Add Rule**.
- To edit an existing rule, click **Edit** ( ).

If **View** ( ) appears next to a rule instead, the rule belongs to an ancestor policy, or you do not have permission to modify the rule.

Step 2 If this is a new rule, enter a **Name**.

Step 3 Configure the rule components, or accept the defaults.

- Enabled—Specify whether the rule is **Enabled**.
- Position—Specify the rule position; see Access Control Rule Order, on page 1219.
- Action—Choose a rule **Action**; see Access Control Rule Actions, on page 1223.

  Note For FTD, VLAN tags in access rules only apply to inline sets; they cannot be used in access rules applied to firewall interfaces.

- Conditions—Click the corresponding condition you want to add. See Rule Condition Types, on page 363 for more information.
- Deep Inspection—For Allow and Interactive Block rules, click **Intrusion policy** ( ) or **File policy** ( ) to configure the rule's **Inspection** options. If the option is dimmed, no policy of that type is selected for the rule. See Understanding Access Control, on page 1185 for more information.
- Content Restriction—Click **Safe search** ( ) or **YouTube EDU** ( ) to configure content restriction settings on **Applications** of the rule editor. If the option are dimmed, content restriction is disabled for the rule. See About Content Restriction, on page 1297 for more information.

- Logging—Click **Logging** ( ) to specify **Logging** options. If the option is dimmed, connection logging is disabled for the rule. See Best Practices for Connection Logging, on page 2216 for more information.
- Comments—Click the number in the comment column to add **Comments**. The number indicates how many comments the rule already contains. See Access Control Rule Comments, on page 1226 for more information.

Step 4 Save the rule.

Step 5 Click **Save** to save the policy.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Best Practices for Access Control Rules, on page 1194

Enabling and Disabling Access Control Rules

When you create an access control rule, it is enabled by default. If you disable a rule, the system does not use it to evaluate network traffic and stops generating warnings and errors for that rule. When viewing the list of rules in an access control policy, disabled rules are grayed out, although you can still modify them.

Tip You can also enable or disable an access control rule using the rule editor.

Procedure

Step 1 In the access control policy editor, right-click the rule and choose a rule state.

If **View** ( ) appears next to a rule instead, the rule belongs to an ancestor policy, or you do not have permission to modify the rule.

Step 2 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Access Control Rule Components, on page 1218

Positioning an Access Control Rule

You can move an existing rule within an access control policy. When you add or move a rule to a category, the system places it last in the category.



Tip You can move multiple rules at once by selecting the rules then cutting and pasting using the right-click menu.

Before you begin

Review rule order guidelines in Best Practices for Access Control Rules, on page 1194.

Procedure

Step 1 In the access control rule editor, you have the following options:

- If you are adding a new rule, use the **Insert** drop-down list.
- If you are editing an existing rule, click **Move**.

Step 2 Choose where you want to move or insert the rule:

- Choose **into Mandatory** or **into Default**.
- Choose a **into Category**, then choose the user-defined category.
- Choose **above rule** or **below rule**, then type the appropriate rule number.

Step 3 Click **Save**.

Step 4 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Access Control Rule Actions

Every access control rule has an *action* that determines how the system handles and logs matching traffic. You can monitor, trust, block, or allow (with or without further inspection).

The access control policy's *default action* handles traffic that does not meet the conditions of any access control rule with an action other than Monitor.

Access Control Rule Monitor Action

The **Monitor** action is not designed to permit or deny traffic. Rather, its primary purpose is to force connection logging, regardless of how matching traffic is eventually handled.

If a connection matches a Monitor rule, the next non-Monitor rule that the connection matches should determine traffic handling and any further inspection. If there are no additional matching rules, the system should use the default action.

There is an exception, however. If a Monitor rule contains layer 7 conditions—such as an application condition—the system *allows early packets to pass* and the connection to be established (or the SSL handshake to complete). This occurs even if the connection should be blocked by a subsequent rule; this is because these early packets *are not evaluated against subsequent rules*. So that these packets do not reach their destination completely uninspected, you can specify an intrusion policy for this purpose in the access control policy's Advanced settings; see Inspection of Packets That Pass Before Traffic Is Identified, on page 1670. After the system completes its layer 7 identification, it applies the appropriate action to the remaining session traffic.

Caution As a best practice, *avoid placing layer 7 conditions on broadly-defined monitor rules high in your rule priority order*, to prevent inadvertently allowing traffic into your network. Also, if locally bound traffic matches a Monitor rule in a Layer 3 deployment, that traffic may bypass inspection. To ensure inspection of the traffic, enable **Inspect Local Router Traffic** in the advanced device settings for the managed device routing the traffic.

Related Topics

Logging for Monitored Connections, on page 2210

Access Control Rule Trust Action

The **Trust** action allows traffic to pass without deep inspection or network discovery. Trusted traffic is still subject to identity requirements and rate limiting.



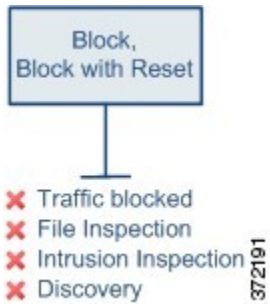
Note Some protocols, such as FTP and SIP, use secondary channels, which the system opens through the process of inspection. In some cases, trusted traffic can bypass all inspection, and these secondary channels cannot be opened properly. If you run into this problem, change the trust rule to **Allow**.

Related Topics

Logging for Trusted Connections, on page 2210

Access Control Rule Blocking Actions

The **Block** and **Block with reset** actions deny traffic without further inspection of any kind.



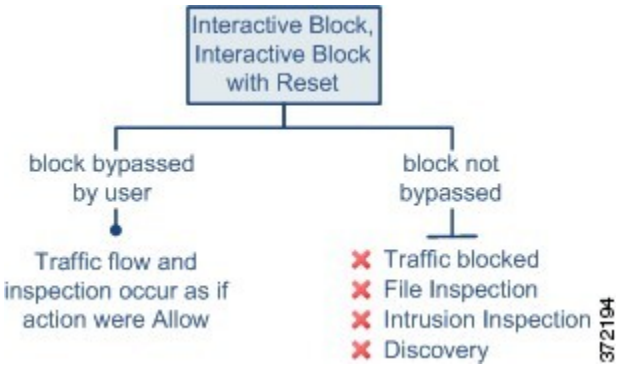
Block with reset rules reset the connection, with the exception of web requests met with an *HTTP response page*. This is because the response page, which you configure to appear when the system blocks web requests, cannot display if the connection is immediately reset. For more information, see HTTP Response Pages and Interactive Blocking, on page 1243.

Related Topics

Logging for Blocked Connections, on page 2211
About HTTP Response Pages, on page 1243

Access Control Rule Interactive Blocking Actions

For more information, see HTTP Response Pages and Interactive Blocking, on page 1243.



If a user bypasses the block, the rule mimics an allow rule. Therefore, you can associate interactive block rules with file and intrusion policies, and matching traffic is also eligible for network discovery.

If a user does not (or cannot) bypass the block, the rule mimics a block rule. Matching traffic is denied without further inspection.

Note that if you enable interactive blocking, you cannot reset *all* blocked connections. This is because the response page cannot display if the connection is immediately reset. Use the **Interactive Block with reset** action to (non-interactively) block-with-reset all non-web traffic, while still enabling interactive blocking for web requests.

For more information, see HTTP Response Pages and Interactive Blocking, on page 1243.

Related Topics

Logging for Allowed Connections, on page 2212
TLS/SSL Rule Blocking Actions, on page 1360

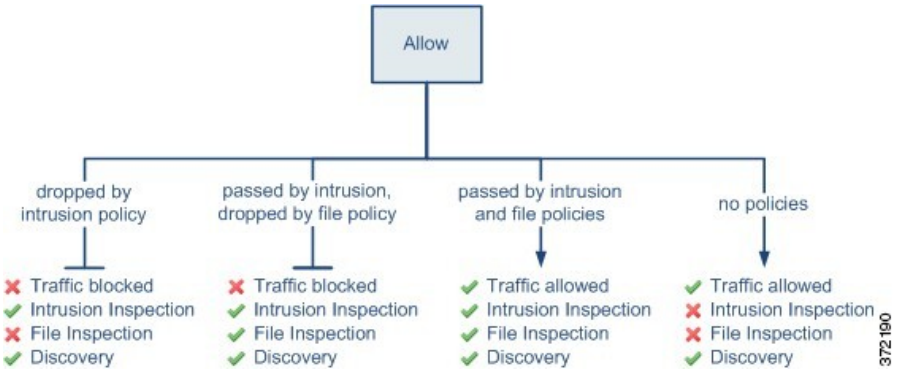
Access Control Rule Allow Action

The **Allow** action allows matching traffic to pass, though it is still subject to identity requirements and rate limiting.

Optionally, you can use deep inspection to further inspect and block unencrypted or decrypted traffic before it reaches its destination:

- You can use an intrusion policy to analyze network traffic according to intrusion detection and prevention configurations, and drop offending packets depending on the configuration.
- You can perform file control using a file policy. File control allows you to detect and block your users from uploading (sending) or downloading (receiving) files of specific types over specific application protocols.
- You can perform network-based advanced malware protection (AMP), also using a file policy. AMP for Networks can inspect files for malware, and block detected malware depending on the configuration.

The following diagram illustrates the types of inspection performed on traffic that meets the conditions of an Allow rule (or a user-bypassed Interactive Block rule. Notice that file inspection occurs before intrusion inspection; blocked files are not inspected for intrusion-related exploits.



For simplicity, the diagram displays traffic flow for situations where both (or neither) an intrusion and a file policy are associated with an access control rule. You can, however, configure one without the other. Without a file policy, traffic flow is determined by the intrusion policy; without an intrusion policy, traffic flow is determined by the file policy.

Regardless of whether the traffic is inspected or dropped by an intrusion or file policy, the system can inspect it using network discovery. However, allowing traffic does not automatically guarantee discovery inspection. The system performs discovery only for connections involving IP addresses that are explicitly monitored by your network discovery policy; additionally, application discovery is limited for encrypted sessions.

Related Topics

Logging for Allowed Connections, on page 2212

Access Control Rule Comments

When you create or edit an access control rule, you can add a comment. For example, you might summarize the overall configuration for the benefit of other users, or note when you change a rule and the reason for the

change. You can display a list of all comments for a rule along with the user who added each comment and the date the comment was added.

When you save a rule, all comments made since the last save become read-only.

Related Topics

Configuring Access Control Policy Preferences

Adding Comments to an Access Control Rule

Procedure

Step 1 In the access control rule editor, click **Comments**.

Step 2 Click **New Comment**.

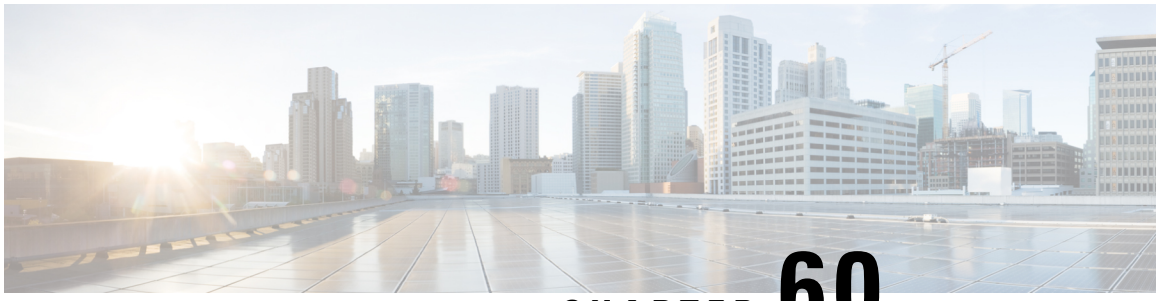
Step 3 Enter your comment and click **OK**. You can edit or delete this comment until you save the rule.

Step 4 Click **Save**.

Step 5 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.



CHAPTER 60

URL Filtering

- URL Filtering Overview, on page 1229
- Best Practices for URL Filtering, on page 1230
- License Requirements for URL Filtering, on page 1233
- Requirements and Prerequisites for URL Filtering, on page 1233
- How to Configure URL Filtering with Category and Reputation, on page 1234
- Manual URL Filtering, on page 1238
- Troubleshoot URL Filtering, on page 1241

URL Filtering Overview

Use the URL filtering feature to control the websites that users on your network can access:

- Category and reputation-based URL filtering—With a URL Filtering license, you can control access to websites based on the URL's general classification (category) and risk level (reputation). This is the recommended option.
- Manual URL filtering—With any license, you can manually specify individual URLs, groups of URLs, and URL lists and feeds to achieve granular, custom control over web traffic. For more information, see Manual URL Filtering, on page 1238.

See also Blocking Traffic with Security Intelligence, on page 1249, a similar but different feature for blocking malicious URLs, domains, and IP addresses.

About URL Filtering with Category and Reputation

With a URL Filtering license, you can control access to websites based on the category and reputation of requested URLs:

- Category—A general classification for the URL. For example, ebay.com belongs to the Auctions category, and monster.com belongs to the Job Search category.

  A URL can belong to more than one category.
- Reputation—How likely the URL is to be used for purposes that might be against your organization's security policy. Reputations range from High Risk (level 1) to Well Known (level 5).

Benefits of Category and Reputation-Based URL Filtering

URL categories and reputations help you quickly configure URL filtering. For example, you can use access control to block high-risk URLs in the Hacking category. Or, you can use QoS to rate limit traffic from sites in the Streaming Media category. There are also categories for types of threats, such as a Spyware and Adware category.

Using category and reputation data simplifies policy creation and administration. It grants you assurance that the system controls web traffic as expected. Because Cisco continually updates its threat intelligence with new URLs, as well as new categories and risks for existing URLs, the system uses up-to-date information to filter requested URLs. Sites that (for example) represent security threats, or that serve undesirable content, may appear and disappear faster than you can update and deploy new policies.

Some examples of how the system can adapt include:

- If an access control rule blocks all gaming sites, as new domains get registered and classified as Games, the system can block those sites automatically. Similarly, if a QoS rule rate limits all streaming media sites, the system can automatically limit traffic to new Streaming Media sites.
- If an access control rule blocks all malware sites and a shopping page gets infected with malware, the system can recategorize the URL from Shopping to Malware Sites and block that site.
- If an access control rule blocks high-risk social networking sites and somebody posts a link on their profile page that contains links to malicious payloads, the system can change the reputation of that page from Benign Sites to High Risk and block it.

Related Topics

Snort® Restart Scenarios, on page 348

Best Practices for URL Filtering

Keep in mind the following guidelines and limitations for URL filtering:

Filter by Category and Reputation

Follow the instructions in How to Configure URL Filtering with Category and Reputation, on page 1234.

Configure Your Policy to Inspect Packets That Must Pass Before a URL Can Be Identified

The system cannot filter URLs before:

- A monitored connection is established between a client and server.
- The system identifies the HTTP or HTTPS application in the session.
- The system identifies the requested URL (for encrypted sessions, from the ClientHello message or the server certificate).

This identification should occur within 3 to 5 packets, or after the server certificate exchange in the TLS/SSL handshake if the traffic is encrypted.

Important! To ensure that your system examines these initial packets that would otherwise pass, see Inspection of Packets That Pass Before Traffic Is Identified, on page 1670 and subtopics.

If early traffic matches all other rule conditions but identification is incomplete, the system allows the packet to pass and the connection to be established (or the TLS/SSL handshake to complete). After the system completes its identification, the system applies the appropriate rule action to the remaining session traffic.

URL Conditions and Rule Order

- Position URL rules after all other rules that *must* be hit.
- URLs can belong to more than one category. It is possible to want to allow one category of websites and block another—whether explicitly or by relying on the default action. In this case, make sure you create and order URL rules so you get the desired effect, depending on whether the allow or the block should take precedence.

For additional guidelines for rules, see the following topics: Best Practices for Access Control Rules, on page 1194 and Rule Condition Mechanics, on page 365.

Uncategorized or Reputationless URLs

When you build a URL rule, you first choose the category you want to match. If you explicitly choose **Uncategorized** URLs, you cannot further constrain by reputation.

You cannot manually assign categories and reputations to URLs, but in access control and QoS policies, you can manually block specific URLs. See Manual URL Filtering, on page 1238.

URL Filtering for Encrypted Web Traffic

When performing URL filtering on encrypted web traffic, the system:

- Disregards the encryption protocol; a rule matches both HTTPS and HTTP traffic if the rule has a URL condition but not an application condition that specifies the protocol.
- Does not use URL lists. You must use URL objects and groups instead.
- Matches HTTPS traffic based on the subject common name in the public key certificate used to encrypt the traffic, and also evaluates the reputation of any other URLs presented at any time during the transaction, including the post-decryption HTTP URL.
- Disregards subdomains within the subject common name.
- Does not display an HTTP response page for encrypted connections blocked by access control rules (or any other configuration); see Limitations to HTTP Response Pages, on page 1243.

HTTP/2

The system can extract HTTP/2 URLs from TLS certificates, but not from a payload.

Manual URL Filtering

- Specify URLs using a custom Security Intelligence list or feed object. Do not use a URL object or directly enter a URL into the rule. For details, see Manual URL Filtering Options, on page 1239.
- If you manually filter specific URLs using URL objects or by entering URLs directly into the rule, carefully consider other traffic that might be affected. To determine whether network traffic matches a URL condition, the system performs a simple substring match. If the requested URL matches any part of the string, the URLs are considered to match.

- If you use manual URL filtering to create exceptions to other rules, position the specific rule with the exceptions above the general rule that would otherwise apply.

Search Query Parameters in URLs

The system does not use search query parameters in the URL to match URL conditions. For example, consider a scenario where you block all shopping traffic. In that case, using a web search to search for amazon.com is not blocked, but browsing to amazon.com is.

URL Filtering in High Availability Deployments

For guidelines for URL filtering with Firepower Management Centers in high availability, see URL Filtering and Security Intelligence, on page 202.

Memory Limitations for Selected Device Models

- If you are using NGIPSv, see the Cisco Firepower NGIPSv Quick Start Guide for VMware for information on allocating the correct amount of memory to perform category and reputation-based URL filtering.
- Device models with less memory store less URL data locally, and the system may therefore check the cloud more frequently to determine category and reputation for sites that are not in the local database.

  Lower-memory devices include:

  - Virtual FTD (FTDv) with 8 GB of RAM
  - ASA 5506-X, ASA 5506H-X, ASA 5506W-X, ASA 5508-X and ASA 5516-X
  - ASA 5512-X, ASA 5515-X and ASA 5525-X
  - 7100 series

Related Topics

Inspection of Packets That Pass Before Traffic Is Identified, on page 1670

Filtering HTTPS Traffic

To filter encrypted traffic, the system determines the requested URL based on information passed during the TLS/SSL handshake: the subject common name in the public key certificate used to encrypt the traffic.

HTTPS filtering, unlike HTTP filtering, disregards subdomains within the subject common name. Do not include subdomain information when manually filtering HTTPS URLs in access control or QoS policies. For example, use example.com rather than www.example.com.

HTTPS filtering also does not support URL lists. You must use URL objects and groups instead.



Tip In an SSL policy, you can handle and decrypt traffic to specific URLs by defining a distinguished name SSL rule condition. The common name attribute in a certificate's subject distinguished name contains the site's URL. Decrypting HTTPS traffic allows access control rules to evaluate the decrypted session, which improves URL filtering.

Controlling Traffic by Encryption Protocol

The system disregards the encryption protocol (HTTP vs HTTPS) when performing URL filtering in access control or QoS policies. This occurs for both manual and reputation-based URL conditions. In other words, URL filtering treats traffic to the following websites identically:

- http://example.com/
- https://example.com/

To configure a rule that matches only HTTP or HTTPS traffic, add an application condition to the rule. For example, you could allow HTTPS access to a site while disallowing HTTP access by constructing two access control rules, each with an application and URL condition.

The first rule allows HTTPS traffic to the website:

Action: Allow
Application: HTTPS
URL: example.com

The second rule blocks HTTP access to the same website:

Action: Block
Application: HTTP
URL: example.com

License Requirements for URL Filtering

FTD License

- Category and reputation filtering—URL Filtering
- Manual filtering—No additional license.

Classic License

- Category and reputation filtering—URL Filtering
- Manual filtering—No additional license.

Requirements and Prerequisites for URL Filtering

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

How to Configure URL Filtering with Category and Reputation

	Do This	More Information
Step	If you will use category and reputation-based URL filtering on an NGIPSv device, allocate the required amount of memory.	Cisco Firepower NGIPSv Quick Start Guide for VMware
Step	Ensure that you have the correct licenses.	Licensing the Firepower System, on page 105, including: . • URL Filtering Licenses for Firepower Threat Defense Devices, on page 112 • URL Filtering Licenses for Classic Devices, on page 128 Assign the URL Filtering license to each managed device that will filter URLs. In order to enable the feature, at least one managed device must have a URL Filtering license assigned to it.
Step	Ensure that your Firepower Management Center can communicate with the cloud to obtain URL filtering data.	Internet Access Requirements, on page 2426 and Communication Port Requirements, on page 2427.
Step	Understand limitations and guidelines and take any necessary actions.	Best Practices for URL Filtering, on page 1230
Step	Enable the URL Filtering feature.	Enable URL Filtering Using Category and Reputation, on page 1235
Step	Configure rules to filter URLs by category and reputation.	Configuring URL Conditions, on page 1236 For the best protection against malicious sites, you must block sites by reputation AND block URLs in all Threat categories. (Optional) Supplement or Selectively Override Category and Reputation-Based URL Filtering, on page 1240
Step	(Optional) Allow users to bypass a website block by clicking through a warning page.	HTTP Response Pages and Interactive Blocking, on page 1243

	Do This	More Information
Step	Order your rules so that traffic hits key rules first.	URL Rule Order, on page 1198
Step	(Optional) Modify advanced options related to URL filtering.	Generally, use the defaults unless you have a specific reason to change them. For information about advanced options, including the following, see Access Control Policy Advanced Settings, on page 1211. • **Maximum URL characters to store in connection events** • **Allow an Interactive Block to bypass blocking for (seconds)** • **Retry URL cache miss lookup**
Step	Deploy your changes.	Deploy Configuration Changes, on page 346
Step	Be sure you have enabled other Firepower features that protect your network from malicious sites	See Blocking Traffic with Security Intelligence, on page 1249.

Enable URL Filtering Using Category and Reputation

You must be an Admin user to perform this task.

Before you begin

Complete prerequisites described in How to Configure URL Filtering with Category and Reputation, on page 1234.

Procedure

Step 1 Choose **System** > **Integration**.

Step 2 Click **Cisco CSI**.

Step 3 Configure URL Filtering Options, on page 1235.

Step 4 Click **Save**.

URL Filtering Options

The following options are on the **System** > **Integration** page:

Enable URL Filtering

Allows traffic filtering based on a website's general classification, or category, and risk level, or reputation. Adding a URL Filtering license automatically enables **Enable URL Filtering**. URL filtering must be enabled before you can choose other URL filtering options.

When you enable URL filtering, depending on how long since URL filtering was last enabled, or if this is the first time you are enabling URL filtering, the Firepower Management Center downloads URL data from Cisco Collective Security Intelligence (Cisco CSI). This process may take some time.

Enable Automatic Updates

Options for updating URL filtering threat data:

- If you enable the **Enable Automatic Updates** option on the **System > Integration** page, the Firepower Management Center checks the cloud every 30 minutes for updates. This option is enabled by default when you add a URL filtering license.
- If you need strict control over when the system contacts external resources, disable automatic updates on this page and instead create a recurring task using the scheduler. See Automating URL Filtering Updates Using a Scheduled Task, on page 193.

Update Now

You can perform a one-time, on-demand update by clicking the **Update Now** button at the top of this dialog box, but you should also either enable automatic updates or create a recurring task using the scheduler. You cannot start an on-demand update if an update is already in progress.

Although daily updates tend to be small, if it has been more than five days since your last update, new URL data may take up to 20 minutes to download, depending on your bandwidth. Then, it may take up to 30 minutes to perform the update itself.

Query Cisco CSI for Unknown URLs

Allows the system to submit URLs to the cloud for threat intelligence evaluation when users browse to a website whose category and reputation are not in the local dataset. Disable this option if you do not want to submit your uncategorized URLs, for example, for privacy reasons.

Connections to uncategorized URLs do **not** match rules with category or reputation-based URL conditions. You cannot assign categories or reputations to URLs manually.

Configuring URL Conditions

Protect your network by controlling access to sites based on URL category and reputation.



Caution Adding the first or removing the last URL or Category category/reputation condition in an access control or SSL (but not a QoS) rule restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Procedure

Step 1 In the rule editor, click the following for URL conditions:

- Access control or QoS—Click **URLs**.
- SSL—Click **Category**.

Step 2 Find and choose the URL categories that you want to control:

In an access control or QoS rule, click **Category**.

Step 3 (Optional) Constrain URL categories by choosing a **Reputation**.

Note that if you explicitly match **Uncategorized** URLs, you cannot further constrain by reputation, because uncategorized URLs do not have reputations. Choosing a reputation level also includes other reputations either more or less severe than the level you choose, depending on the rule action:

- Includes less severe reputations—If the rule allows or trusts web traffic. For example, if you configure an access control rule to allow Benign Sites (level 4), it also automatically allows Well Known (level 5) sites.
- Includes more severe reputations—If the rule rate limits, decrypts, blocks, or monitors web traffic. For example, if you configure an access control rule to block Suspicious Sites (level 2), it also blocks High Risk (level 1) sites.

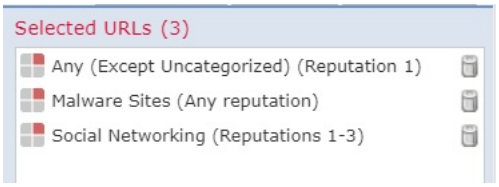
If you change the rule action, the system automatically changes the reputation levels in URL conditions.

Step 4 Click **Add to Rule**, or drag and drop.

Step 5 Save or continue editing the rule.

Example: URL Condition in an Access Control Rule

The following graphic shows the URL condition for an access control rule that blocks all malware sites, all High Risk sites, and all non-benign social networking sites.



The following table summarizes how you build the condition.

Blocked URL	Category	Reputation
Malware sites, regardless of reputation	Malware Sites	Any
Any URL with a high risk (level 1)	Any	1 - High Risk

Blocked URL	Category	Reputation
Social networking sites with a risk greater than benign (levels 1 through 3)	Social Network	3 - Benign sites with security risks

What to do next

- (Optional) Supplement or Selectively Override Category and Reputation-Based URL Filtering, on page 1240
- Return to How to Configure URL Filtering with Category and Reputation, on page 1234.
- If you are done making changes, Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Rules with URL Conditions

The following table lists rules that support URL conditions, and the types of filtering that each rule type supports.

Rule Type	Supports Category and Reputation Filtering?	Supports Manual Filtering?
Access control	Yes	Yes
SSL	Yes	No; use distinguished name conditions instead
QoS	Yes	Yes

URL Rule Order

For the most effective URL matching, place rules that include URL conditions before other rules, particularly if the URL rules are block rules and the other rules meet both of the following criteria:

- They include application conditions.
- The traffic to be inspected is encrypted.

If you configure exceptions to a rule, put the exception above the other rule.

Manual URL Filtering

In access control and QoS rules, you can supplement or selectively override category and reputation-based URL filtering by manually filtering individual URLs, groups of URLs, or URL lists and feeds.

For example, you might use access control to block a category of websites that are not appropriate for your organization. However, if the category contains a website that is appropriate, and to which you want to provide access, you can create a manual Allow rule for that site and place it before the Block rule for the category.

You can perform this type of URL filtering without a special license.

Manual URL filtering is not supported in SSL rules; instead, use distinguished name conditions.

Caution Depending on how you implement manual URL filtering, URL matching may not be what you intend. See Manual URL Filtering Options, on page 1239.

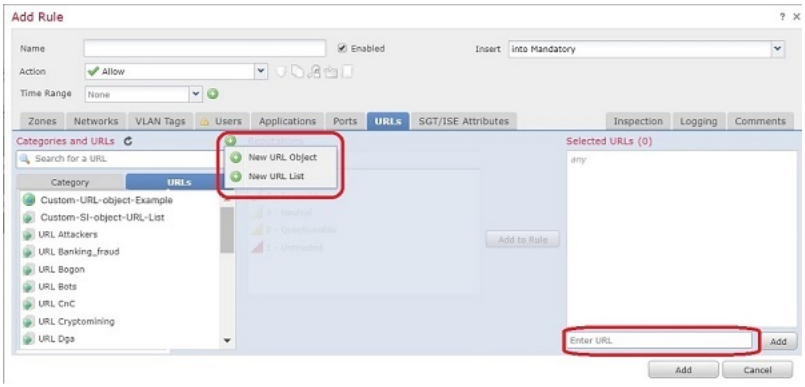
Related Topics

Security Intelligence Lists and Feeds, on page 427

Manual URL Filtering Options

There are several ways to specify URLs for manual URL filtering:

Figure 52: Manual URL Filtering Options in an Access Control Rule



Option	Description
(Best practice) Use custom Security Intelligence URL list or feed objects. This is the **New URL List** option on the rule page in the web interface.	This is the recommended method for manual URL filtering. You can create a new list or feed, or choose an existing one from the URLs sub-tab of the URLs tab in an access control or QoS rule. For more information, see Custom Security Intelligence Lists and Feeds, on page 432 and subtopics.

Option	Description
Use URL objects, individually or as groups. (URL objects are described at URL Objects, on page 408.) Or Enter URLs directly into the access control rule. (The **Enter URL** option on the rule page in the web interface.)	If you do not include a path (that is, there is no / character in the URL), the match is based on the server's hostname only. The hostname is considered a match if it comes after the :// separator, or after any dot in the hostname. For example, ign.com matches ign.com and www.ign.com, but it does not match verisign.com. If you include one or more / character, the entire URL string is used for a substring match, including the server name, path, and any query parameters. However, we recommend that you do not use manual URL filtering to block or allow individual web pages or parts of sites, as servers can be reorganized and pages moved to new paths. Substring matching can also lead to unexpected matches, where the string you include in the URL object also matches paths on unintended servers or strings within query parameters. The **Enter URL** option does not support wildcards.

Supplement or Selectively Override Category and Reputation-Based URL Filtering

In access control or QoS rules, you can use Security Intelligence URL lists and feeds to supplement, or to specify exceptions to, your category and reputation-based URL filtering rules.

(In SSL rules, use distinguished name conditions to serve this purpose.)

Before you begin

- Configure URL filtering using category and reputation. See Configuring URL Conditions, on page 1236.
- Understand important best practices for manual URL filtering. See Best Practices for URL Filtering, on page 1230 and Manual URL Filtering Options, on page 1239.
- Configure one or more Security Intelligence objects (lists or feeds) containing the URLs that you want to use for manual filtering. See Custom Security Intelligence Lists and Feeds, on page 432.

Procedure

Step 1 Navigate to the access control or QoS policy in which you will define your rule.

Step 2 Create or edit the rule in which you will add your new condition:

- If you are supplementing a category- or reputation-based URL filtering rule, edit the existing rule.
- If you are overriding or creating exceptions to a category- or reputation-based URL filtering rule, create a new rule.

Step 3 If you are creating a new rule, configure the rule name, position, action, and other options at the top of the rule.

Important! If the list or feed you are configuring in this procedure contains exceptions to category- or reputation-based rules, put this rule above those rules in the rule order.

Step 4 Click **URLs**.

Step 5 Click **URLs** (beside the **Category** tab.)

Step 6 Select the list or feed you created in the prerequisite to this task.

Step 7 Click **Add to Rule**.

Step 8 Click **Add** or continue editing the rule.

What to do next

(Optional) In SSL rules, use distinguished name conditions to configure parallel behavior.

Troubleshoot URL Filtering

Expected URL Category is Missing from the Categories List

The URL filtering feature uses a different set of categories than the Security Intelligence feature; the category that you expect to see may be a Security Intelligence category. To see those categories, look at the **URLs** tab on the **Security Intelligence** tab in an access control policy.

Initial Packets Are Passing Uninspected

See Inspection of Packets That Pass Before Traffic Is Identified, on page 1670 and subtopics.

How can I find the category and reputation of a particular URL?

Do a manual lookup. See Finding URL Category and Reputation, on page 2136.

Error when attempting a manual lookup: "Cloud Lookup Failure for <URL>"

Make sure the feature is properly enabled. See the prerequisites in Finding URL Category and Reputation, on page 2136.

URL appears to be incorrectly handled based on its URL category and reputation

Problem: The system does not handle the URL correctly based on its URL category and reputation.

Solutions:

- Verify that the URL category and reputation associated with the URL are what you think they are. See Finding URL Category and Reputation, on page 2136.
- The following issues may be addressed by settings described in URL Filtering Options, on page 1235, accessible using Enable URL Filtering Using Category and Reputation, on page 1235.
  - The URL cache may hold stale information. Contact Cisco TAC.
  - The local data set may not be updated with current information from the cloud. See information about the **Enable Automatic Updates** setting in URL Filtering Options, on page 1235.
  - The system may be configured to *not* check the cloud for current data. See information about the **Query Cisco CSI for Unknown URLs** setting in URL Filtering Options, on page 1235.

- Your access control policy may be configured to pass traffic to the URL without checking the cloud. See information about the **Retry URL cache miss lookup** setting in Access Control Policy Advanced Settings, on page 1211.
- See also Best Practices for URL Filtering, on page 1230.
- Verify that the URL is being handled using the access control rule that you think it is being handled by, and that the rule does what you think it does. Consider rule order.
- Verify that the local URL category and reputation database on the Firepower Management Center is successfully being updated from the cloud and that managed devices are successfully being updated from the Firepower Management Center.

  Status of these processes are reported in the Health Monitor, in the **URL Filtering Monitor** module. For details, see Health Monitoring, on page 275.

If you want to immediately update the local URL category and reputation database, go to **System** > **Integration**, click **Cisco CSI**, then click **Update Now**. For more information, see URL Filtering Options, on page 1235.

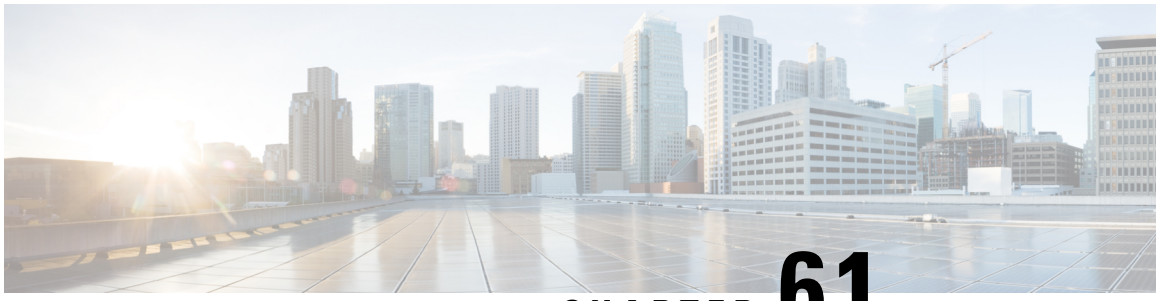
A URL category or reputation is not correct

For access control or QoS rules: Use manual filtering, paying careful attention to rule order. See Manual URL Filtering, on page 1238 and Configuring URL Conditions, on page 1236.

For SSL rules: Manual filtering is not supported. Instead, use distinguished name conditions.

Events Do Not Include URL Category and Reputation

- Make sure you have included applicable URL rules in an access control policy, the rules are active, and the policies have been deployed to the relevant devices.
- URL category and reputation do not appear in an event if the connection is processed before it matches a URL rule.
- The rule that handles the connection must be configured for URL category and reputation.
- Even if you have configured URL categories in the Categories tab in an SSL rule, you must also configure the URLs tab in a rule in your access control policy.



CHAPTER 61

HTTP Response Pages and Interactive Blocking

The following topics describe how to configure custom pages to display when the system blocks web requests:

- About HTTP Response Pages, on page 1243
- Requirements and Prerequisites for HTTP Response Pages, on page 1244
- Choosing HTTP Response Pages, on page 1245
- Interactive Blocking with HTTP Response Pages, on page 1245

About HTTP Response Pages

As part of access control, you can configure an *HTTP response page* to display when the system blocks web requests, using either access control rules or the access control policy default action.

The response page displayed depends on how you block the session:

- **Block Response Page**: Overrides the default browser or server page that explains that the connection was denied.
- **Interactive Block Response Page**: Warns users, but also allows them to click a button (or refresh the page) to load the originally requested site. Users may have to refresh after bypassing the response page to load page elements that did not load.

If you do not choose a response page, the system blocks sessions without interaction or explanation.

Limitations to HTTP Response Pages

Response Pages are for Access Control Rules/Default Action Only

The system displays a response page only for unencrypted or decrypted HTTP/HTTPS connections blocked (or interactively blocked) either by access control rules or by the access control policy default action. The system does not display a response page for connections blocked by any other policy or mechanism.

Displaying the Response Page Disables Connection Reset

The system cannot display a response page if the connection is reset (RST packet sent). If you enable response pages, the system prioritizes that configuration. Even if you choose **Block with reset** or **Interactive Block with reset** as the rule action, the system displays the response page and does not reset matching web connections. To ensure that blocked web connections reset, you must disable response pages.

Note that all non-web traffic that matches the rule *is* blocked with reset.

No Response Page for Encrypted Connections (Must Decrypt)

The system does not display a response page for encrypted connections blocked by access control rules (or any other configuration). Access control rules evaluate encrypted connections if you did not configure an SSL policy, or your SSL policy passes encrypted traffic.

For example, the system cannot decrypt HTTP/2 or SPDY sessions. If web traffic encrypted using one of these protocols reaches access control rule evaluation, the system does not display a response page if the session is blocked.

However, the system does display a response page for connections decrypted by the SSL policy, then blocked (or interactively blocked) either by access control rules or by the access control policy default action. In these cases, the system encrypts the response page and sends it at the end of the reencrypted SSL stream.

No Response Page for "Promoted" Connections

The system does not display a response page when web traffic is blocked as a result of a promoted access control rule (an early-placed blocking rule with only simple network conditions).

No Response Page for Certain Redirected Connections

If a URL is entered without specifying "http" or "https", and the browser initiates the connection on port 80, and the user clicks through a response page, and the connection is subsequently redirected to port 443, the user will not see a second interactive response page because the response to this URL is already cached.

No Response Page Before URL Identification

The system does not display a response page when web traffic is blocked before the system identifies the requested URL; see Best Practices for URL Filtering, on page 1230.

Requirements and Prerequisites for HTTP Response Pages

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Choosing HTTP Response Pages

Reliable display of HTTP response pages depends on your network configuration, traffic loads, and size of the page. Smaller pages are more likely to display successfully.

Procedure

Step 1 In the access control policy editor, click **HTTP Responses**.

If the controls are dimmed, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 Choose the **Block Response Page** and **Interactive Block Response Page**:

- System-provided—Displays a generic response. Click **View** ( ) to view the code for this page.
- Custom—Create a custom response page. A pop-up window appears, prepopulated with system-provided code that you can replace or modify by clicking **Edit** ( ). A counter shows how many characters you have used.
- None—Disables the response page and blocks sessions without interaction or explanation. To quickly disable interactive blocking for the whole access control policy, choose this option.

Step 3 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Interactive Blocking with HTTP Response Pages

When you configure interactive blocking, users can load an originally requested site after reading a warning. Users may have to refresh after bypassing the response page to load page elements that did not load.

Tip To quickly disable interactive blocking for the whole access control policy, display neither the system-provided page nor a custom page. The system then blocks all connections without interaction.

If a user does not bypass an interactive block, matching traffic is denied without further inspection. If a user bypasses an interactive block, the access control rule allows the traffic, although the traffic may still be subject to deep inspection and blocking.

By default, a user bypass is in effect for 10 minutes (600 seconds) without displaying the warning page on subsequent visits. You can set the duration to as long as a year, or you can force the user to bypass the block every time. This limit applies to every Interactive Block rule in the policy. You cannot set the limit per rule.

Logging options for interactively blocked traffic are identical to those in allowed traffic, but if a user does not bypass the interactive block, the system can log only beginning-of-connection events. When the system

initially warns the user, it marks any logged beginning-of-connection event with the `Interactive Block` or `Interactive Block with reset` action. If the user bypasses the block, additional connection events logged for the session have an action of `Allow`.

Configuring Interactive Blocking

Procedure

Step 1 As part of access control, configure an access control rule that matches web traffic; see Create and Edit Access Control Rules, on page 1220:

- Action—Set the rule action to **Interactive Block** or **Interactive Block with reset**; see Access Control Rule Interactive Blocking Actions, on page 1225.
- Conditions—Use URL conditions to specify the web traffic to interactively block; see URL Conditions (URL Filtering), on page 384.
- Logging—Assume users will bypass the block and choose logging options accordingly; see Logging for Allowed Connections, on page 2212.
- Inspection—Assume users will bypass the block and choose deep inspection options accordingly; see Understanding Access Control, on page 1185.

Step 2 (Optional) On access control policy **HTTP Responses**, choose a custom interactive-block HTTP response page; see Choosing HTTP Response Pages, on page 1245.

Step 3 (Optional) On access control policy **Advanced**, change the user bypass timeout; see Setting the User Bypass Timeout for a Blocked Website, on page 1246.

After a user bypasses a block, the system allows the user to browse to that page without warning until the timeout period elapses.

Step 4 Save the access control policy.

Step 5 Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Setting the User Bypass Timeout for a Blocked Website

Procedure

Step 1 Log in to the FMC if you haven't already done so.

Step 2 Click **Policies** > **Access Control** .

Step 3 Click **Edit** ( ).

Step 4 Click **Edit** ( ) next to General Settings.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings.If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

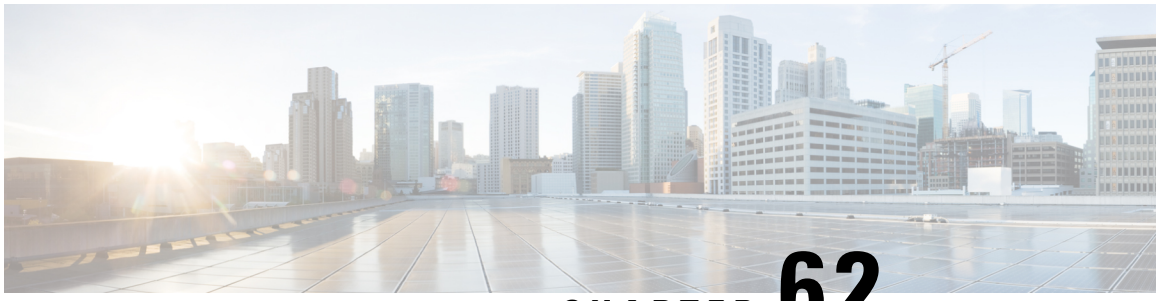
Step 5 In the **Allow an Interactive Block to bypass blocking for (seconds)** field, type the number of seconds that must elapse before the user bypass expires. Setting this value to `0` means the interactive block response is displayed once and the user bypass never expires.

Step 6 Click **OK**.

Step 7 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.



CHAPTER 62

Blocking Traffic with Security Intelligence

The following topics provide an overview of Security Intelligence, including use of lists for blocking and allowing traffic and basic configuration.

- About Security Intelligence, on page 1249
- Best Practices for Security Intelligence, on page 1250
- License Requirements for Security Intelligence, on page 1250
- Requirements and Prerequisites for Security Intelligence, on page 1251
- Security Intelligence Sources, on page 1251
- Configure Security Intelligence, on page 1252
- Security Intelligence Monitoring, on page 1258
- Override Security Intelligence Blocking, on page 1259
- Troubleshooting Security Intelligence, on page 1259
- History for Security Intelligence Block Listing, on page 1260

About Security Intelligence

As an early line of defense against malicious internet content, Security Intelligence uses reputation intelligence to quickly block connections to or from IP addresses, URLs, and domain names. This is called *Security Intelligence block listing*.

Security Intelligence is an early phase of access control, before the system performs more resource-intensive evaluation. Using a Block list improves performance by quickly excluding traffic that does not require inspection.

Note You cannot use a Block list to block fastpathed traffic. 8000 Series fastpathing and prefilter evaluation occur before Security Intelligence filtering. Fastpathed traffic bypasses all further evaluation, including Security Intelligence.

Although you can configure custom Block lists, Cisco provides access to regularly updated intelligence feeds. Sites representing security threats such as malware, spam, botnets, and phishing appear and disappear faster than you can update and deploy custom configurations.

You can refine Security Intelligence Block listing with Do Not Block lists and monitor-only Block lists. These mechanisms exempt traffic from being blocked by a Block list, but do **not** automatically trust or fastpath

matching traffic. Traffic added to a Do Not Block list or monitored at the Security Intelligence stage is intentionally subject to further analysis with the rest of access control.

Related Topics

Security Intelligence Lists and Feeds, on page 427
Other Connections You Can Log, on page 2209
Using Connection and Security Intelligence Event Tables, on page 2245

Best Practices for Security Intelligence

- Configure your access control policies to block threats detected by Cisco-provided Security Intelligence feeds. See Configuration Example: Security Intelligence Blocking, on page 1257.
- If you want to supplement the Cisco-provided Security Intelligence feeds with custom threat data, or manually block emerging threats:
  - For IP addresses, use custom Security Intelligence lists and feeds, or Network objects or groups. To create these, see Security Intelligence Lists and Feeds, on page 427 and Network Objects, on page 402, and their subtopics. To use them for Security Intelligence, see Configure Security Intelligence, on page 1252.
  - For URLs and domains, use custom Security Intelligence lists and feeds, *not* objects or groups. See details at Manual URL Filtering Options, on page 1239.
  - You can also add entries to a Block list from events. See Global and Domain Security Intelligence Lists, on page 428.
- To test new feeds, or for passive deployments, set the action from block to monitor only. See Security Intelligence Monitoring, on page 1258.
- If you need to exclude specific sites or addresses from Security Intelligence blocking, see Override Security Intelligence Blocking, on page 1259.
- System-provided Security Intelligence categories may change over time and without notification; you should plan to check periodically for changes, and modify your policies accordingly.
- You should also configure URL filtering, a separate feature with separate licensing requirements, for further protection against malicious sites. See URL Filtering, on page 1229.

License Requirements for Security Intelligence

FTD License

Threat

Classic License

Protection

Requirements and Prerequisites for Security Intelligence

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Security Intelligence Sources

- System-provided feeds

  Cisco provides access to regularly updated intelligence feeds for domains, URLs and IP addresses. For more information, see Security Intelligence Lists and Feeds, on page 427.

- Third-party feeds

  Optionally, supplement Cisco-provided feeds with third-party reputation feeds, which are dynamic lists that the Firepower Management Center downloads from the internet on a regular basis. See Custom Security Intelligence Feeds, on page 434.

- Custom Block lists or feeds (or objects or groups)

  Block specific IP addresses, URLs, or domain names using a manually-created list or feed (for IP addresses, you can also use network objects or groups.)

  For example, if you become aware of malicious sites or addresses that are not yet blocked by a feed, add these sites to a custom Security Intelligence list and add this custom list to the Block list in the Security Intelligence tab of your access control policy, as described in Custom Security Intelligence Lists, on page 436 and Configure Security Intelligence, on page 1252.

  For IP addresses, you can optionally use network objects rather than lists or feeds for this purpose; for information, see Network Objects, on page 402. (For URLs, using lists and feeds is strongly recommended over other methods.)

- Custom Do Not Block lists or feeds

  Override Security Intelligence blocking for specific sites or addresses. See Override Security Intelligence Blocking, on page 1259.

- Global Block lists (one each for Network, URL and DNS)

While reviewing events, you can immediately add an event's IP address, URL, or domain to the applicable Global Block List so that Security Intelligence will handle future traffic from that source. See Global and Domain Security Intelligence Lists, on page 428.

- Global Do Not Block lists (one each for Network, URL and DNS)

  While reviewing events, you can immediately add an event's IP address, URL, or domain to the applicable Global Do Not Block List if you do not want Security Intelligence to block future traffic from that source. See Global and Domain Security Intelligence Lists, on page 428.

Configure Security Intelligence

Each access control policy has Security Intelligence options. You can add network objects, URL objects and lists, and Security Intelligence feeds and lists to a Block list or Do Not Block list, and constrain any of these by security zone. You can also associate a DNS policy with your access control policy, and add domain names to a Block or Do Not Block list.

The number of objects in the Do Not Block lists plus the number in the Block lists cannot exceed 125 network objects, or 32767 URL objects and lists.



Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.



Caution From Security Intelligence in an access control policy, adding multiple objects to a Block or Do Not Block list, or deleting multiple objects, sometimes restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information. Note that whether the Snort process restarts can vary by device, depending on the memory available for inspection.

Before you begin

- Tip: For guidance on minimum configuration recommendations, see also Configuration Example: Security Intelligence Blocking, on page 1257.
- To ensure that all options are available to select, add at least one managed device to your management center.
- In passive deployments, or if you want to set Security Intelligence filtering to monitor-only, enable logging; see Logging Connections with Security Intelligence, on page 2220.
- Configure a DNS policy to take Security Intelligence action for domains. For more information, see DNS Policies, on page 1261.

Procedure

Step 1 In the access control policy editor, click **Security Intelligence**.

If the controls are dimmed, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 You have the following options:

- Click **Networks** to add network objects (IP addresses).
- Click **URLs** to add URL objects.

Step 3 Find the **Available Objects** you want to add to the Block or Do Not Block list. You have the following options:

- Search the available objects by typing in the **Search by name or value** field. Clear the search string by clicking **Reload** ( ) or **Clear** ( ).
- If no existing list or feed meets your needs, click **Add** ( ), select **New Network List** or **New URL List**, and proceed as described in Creating Security Intelligence Feeds, on page 434 or Uploading New Security Intelligence Lists to the Firepower Management Center, on page 436.
- If no existing object meets your needs, click **Add** ( ), select **New Network Object** or **New URL Object**, and proceed as described in Creating Network Objects, on page 403.

Security Intelligence ignores IP address blocks using a `/0` netmask.

Step 4 Choose one or more **Available Objects** to add.

Step 5 (Optional) Choose an **Available Zone** to constrain the selected objects by zone.

You cannot constrain system-provided Security Intelligence lists by zone.

Step 6 Click **Add to Do Not Block list** or **Add to Block list**, or click and drag the selected objects to either list.

To remove an object from a Block or Do Not Block list, click **Delete** ( ) To remove multiple objects, choose the objects and right-click to **Delete Selected**.

Step 7 (Optional) Set objects on the Block list to monitor-only by right-clicking the object under **Block List**, then choosing **Monitor-only (do not block)**.

You cannot set system-provided global Security Intelligence lists to monitor only.

Step 8 Choose a DNS policy from the **DNS Policy** drop-down list.

Step 9 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Security Intelligence Lists and Feeds, on page 427
Snort® Restart Scenarios, on page 348

Security Intelligence Options

Use the Security Intelligence tab in the access control policy editor to configure network (IP address) and URL Security Intelligence, and to associate the access control policy with a DNS policy in which you have configured Security Intelligence for domains.

Available Objects

Available objects include:

- Security Intelligence categories populated by the system-provided feed.

  For details, see Security Intelligence Categories, on page 1255.

- System-provided Global Block and Do Not Block lists.

  For descriptions, see Security Intelligence Sources, on page 1251.

- Security Intelligence lists and feeds that you create under Object > Object Management > Security Intelligence.

  For descriptions, see Security Intelligence Sources, on page 1251.

- Network and URL objects and groups that are configured on the respective pages under Object > Object Management. These are different from the Security Intelligence objects in the previous bullet.

  For details about network objects, see Network Objects, on page 402. (For URLs, use Security Intelligence lists or feeds rather than objects or groups.)

Available Zones

Except for the system-provided Global lists, you can constrain Security Intelligence filtering by zone.

For example: To improve performance, you may want to target enforcement. As a more specific example, you can block spam only for a security zone that handles email traffic.

To enforce Security Intelligence filtering for an object on multiple zones, you must add the object to the Block or Do Not Block list separately for each zone.

DNS Policy

In order to match DNS traffic using Security Intelligence, you must select a DNS policy for your Security Intelligence configuration.

Using Block or Do Not Block lists, or monitoring traffic based on a DNS list or feed, also requires that you:

- Configure DNS Security Intelligence lists and feeds. See Security Intelligence Lists and Feeds, on page 427.
- Create a DNS policy. See Creating Basic DNS Policies, on page 1264 for more information.
- Configure DNS rules that reference your DNS lists or feeds. See Creating and Editing DNS Rules, on page 1266 for more information.
- Because you deploy the DNS policy as part of your access control policy, you must associate both policies. See DNS Policy Deploy, on page 1272 for more information.

Do Not Block List

See Override Security Intelligence Blocking, on page 1259.

To select all objects in the list, right-click an object.

Block List

See Configuration Example: Security Intelligence Blocking, on page 1257 and other topics in this chapter.

For explanations of the visual indicators in the Block list, see Block List Icons, on page 1256.

To select all objects in the list, right-click an object.

Logging

Security Intelligence logging, enabled by default, logs all blocked and monitored connections handled by an access control policy's target devices. However, the system does not log Do Not Block list matches; logging of connections on the Do Not Block list depends on their eventual disposition. Logging must be enabled for connections on the Block list before you can set objects on that list to monitor-only.

To enable, disable, or view logging settings, right-click an object in the Block list.

Related Topics

Global and Domain Security Intelligence Lists, on page 428
Security Intelligence Lists and Multitenancy, on page 429

Security Intelligence Categories

Security Intelligence categories are determined by the system-provided feeds described in Security Intelligence Lists and Feeds, on page 427.

These categories are used in the following locations:

- The Networks sub-tab on the Security Intelligence tab of an access control policy
- The URLs sub-tab beside the Networks tab on the Security Intelligence tab of an access control policy
- In a DNS policy on the DNS tab in the DNS rule configuration page
- In events generated when traffic matches Block or Monitor configurations in the above locations

Categories are updated by Talos from the cloud, and this list may change independently of Firepower releases.

Table 87: Cisco Talos Intelligence Group (Talos) Feed Categories

Security Intelligence Category	Description
Attackers	Active scanners and hosts known for outbound malicious activity
Banking_fraud	Sites that engage in fraudulent activities that relate to electronic banking
Bogon	Bogon networks and unallocated IP addresses
Bots	Sites that host binary malware droppers

Security Intelligence Category	Description
CnC	Sites that host command-and-control servers for botnets
Cryptomining	Hosts providing remote access to pools and wallets for the purpose of mining cryptocurrency
Dga	Malware algorithms used to generate a large number of domain names acting as rendezvous points with their command-and-control servers
Exploitkit	Software kits designed to identify software vulnerabilities in clients
High_risk	Domains and hostnames that match against the OpenDNS predictive security algorithms from security graph
Ioc	Hosts that have been observed to engage in Indicators of Compromise (IOC)
Link_sharing	Websites that share copyrighted files without permission
Malicious	Sites exhibiting malicious behavior that do not necessarily fit into another, more granular, threat category
Malware	Sites that host malware binaries or exploit kits
Newly_seen	Domains that have recently been registered, or not yet seen via telemetry. **Attention** Currently, this category does not have any active feed and is reserved for future use.
Open_proxy	Open proxies that allow anonymous web browsing
Open_relay	Open mail relays that are known to be used for spam
Phishing	Sites that host phishing pages
Response	IP addresses and URLs that are actively participating in malicious or suspicious activity
Spam	Mail hosts that are known for sending spam
Spyware	Sites that are known to contain, serve, or support spyware and adware activities
Suspicious	Files that appear to be suspicious and have characteristics that resemble known malware
Tor_exit_node	Hosts known to offer exit node services for the Tor Anonymizer network

Block List Icons

The following visual indicators may appear in the Block list on the Security Intelligence tab in an access control policy:

Icon or Visual Indicator	Description
Block ()	The object is set to block.
Monitor ()	The object is set to monitor-only. See Security Intelligence Monitoring, on page 1258
An object is displayed in strikethrough text	The same object is also on the Do Not Block list, which overrides the block.

Configuration Example: Security Intelligence Blocking

Configure your access control policy to block all threats detectable by the system's regularly updated Security Intelligence feeds.

The number of objects in the Block lists plus the number in the Do Not Block lists cannot exceed 125 network objects, or 32767 URL objects and lists.



Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.



Caution From Security Intelligence in an access control policy, adding multiple objects to a Do Not Block list or Block list, or deleting multiple objects, sometimes restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information. Note that whether the Snort process restarts can vary by device, depending on the memory available for inspection.

Before you begin

- To ensure that all options are available to select, add at least one managed device to your management center.
- Configure a DNS policy to block all Security Intelligence threat categories for domains. For more information, see DNS Policies, on page 1261.
- If you have, or will have, custom lists of entities to block, create a Security Intelligence object of each type (URLs, DNS, Networks.) See Security Intelligence Lists and Feeds, on page 427.

Procedure

Step 1 Click **Policies > Access Control**.

Step 2 Create a new access control policy or edit an existing policy.

Step 3 In the access control policy editor, click **Security Intelligence**.

If the controls are dimmed, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 4 Click **Networks** to add blocking criteria for IP addresses.

a) Scroll down in the Networks list and select all of the threat categories listed below the Global lists.
b) If applicable, select the security zones for which you want to block these threats.
c) Click **Add to Block List**.
d) If you have created custom lists or feeds with addresses to block, add those to the Block List using the same steps as above.

Step 5 Click **URLs** to add blocking criteria for URLs, and repeat the steps you followed for Networks.

Step 6 Choose a DNS policy from the **DNS Policy** drop-down list; see DNS Policy Overview, on page 1261.

Step 7 Click **Save**.

What to do next

- Enable logging for these connections; see Logging Connections with Security Intelligence, on page 2220.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.
- For additional protection, configure URL filtering to block malicious URLs. See URL Filtering, on page 1229.

Security Intelligence Monitoring

Monitoring logs connection events for traffic that would have been blocked by Security Intelligence, but does not block the traffic. Monitoring is especially useful for:

- Testing feeds before you implement them.

  Consider a scenario where you want to test a third-party feed before you implement blocking using that feed. When you set the feed to monitor-only, the system allows connections that would have been blocked to be further analyzed by the system, but also logs a record of each of those connections for your evaluation.

- Passive deployments, to optimize performance.

  Managed devices that are deployed passively cannot affect traffic flow; there is no advantage to configuring the system to block traffic. Additionally, because blocked connections are not actually blocked in passive deployments, the system may report multiple beginning-of-connection events for each blocked connection.

To Configure Security Intelligence Monitoring:

After you configure Security Intelligence blocking following the instructions in Configuration Example: Security Intelligence Blocking, on page 1257, right-click each applicable object in the Block list and choose **Monitor-only**. You cannot set system-provided Security Intelligence lists to monitor only.

Override Security Intelligence Blocking

Optionally, you can use Do Not Block lists to exempt specific domains, URLs, or IP addresses from being blocked by Security Intelligence lists or feeds.

For example, you can:

- Override the occasional false-positive block in a reputable Security Intelligence feed
- Inspect specific traffic in depth instead of blocking it early based on reputation
- Exempt otherwise-restricted transactions based on zone from Security Intelligence blocking

  For example, you can add an improperly classified URL to a Do Not Block list, but then restrict the Do Not Block list object using a security zone used by those in your organization who need to access those URLs. That way, only those with a business need can access the URLs on the Do Not Block list.



Note Entries on a Do Not Block list are *not* automatically trusted or fastpathed; this traffic is intentionally subject to further analysis with the rest of access control.

Procedure

Step 1 Option 1: Add an IP address, URL, or domain from an event to the Global Do Not Block List. See Global and Domain Security Intelligence Lists, on page 428.

Step 2 Option 2: Use a custom Security Intelligence list or feed.

a) Create the custom Security Intelligence list or feed. See Custom Security Intelligence Lists, on page 436 or Creating Security Intelligence Feeds, on page 434.

b) For IP addresses (Networks) and URLs: Edit your access control policy, click the Security Intelligence tab, then click the custom list or feed in the Networks or URLs sub-tab, then click **Add to Do Not Block List**.

c) Save your changes.

d) For domains (DNS): See the "DNS Policy" section in the Security Intelligence Options, on page 1254 topic.

e) Deploy your changes.

Troubleshooting Security Intelligence

Security Intelligence Categories Are Missing from the Available Options List

Symptoms: On the Security Intelligence tab of the access control policy, Security Intelligence categories (such as CnC or Exploitkit) are not displayed in the Networks tab under Available Options.

Cause:

- These categories do not appear until you have added at least one managed device to your management center. You must add a device in order to pull all TALOS feeds.
- The URL filtering feature uses a different set of categories than the Security Intelligence feature; the category that you expect to see may be a URL filtering category. To see URL filtering categories, look at the **URLs** tab in an access control rule.

Troubleshooting Memory Use

Symptoms: Connections that should be blocked by a Security Intelligence Block list are instead evaluated by access control rules. The Security Intelligence health module alerts that it is out of memory.

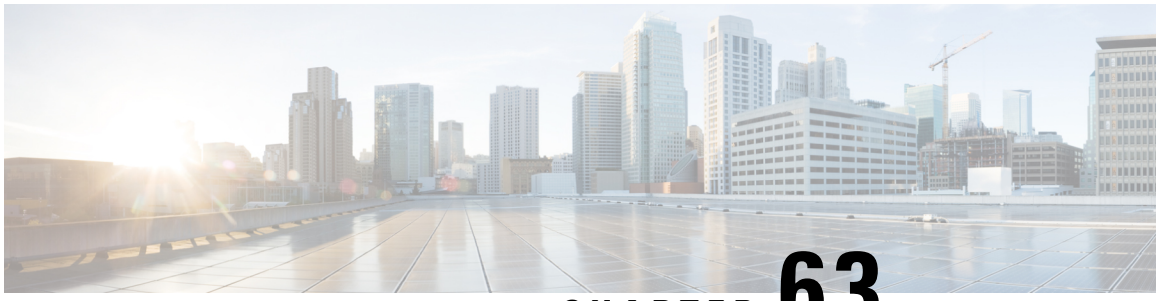
Cause: Memory limitations. Cisco Intelligence Feeds are based on the latest threat intelligence from Cisco Talos Intelligence Group (Talos). These feeds tend to get larger as time passes. When a Firepower device receives a feed update, it loads as many entries as it can into the memory it has allocated for Security Intelligence. When a device cannot load all the entries, it may not block traffic as expected. Some connections that should be blocked by a Block list instead continue to be evaluated by access control rules.

Affected platforms: Lower-memory devices are most likely to have this issue, especially if your Block list includes a lot of Security Intelligence categories or you also filter URLs based on category and reputation. These devices include Firepower 7010, 7020, and 7030; ASA 5506-X, 5508-X, 5516-X, 5512-X , 5515-X, and 5525-X; NGIPSv.

Workaround: If you think this is happening, redeploy configurations to the affected devices. This can allocate more memory to Security Intelligence. If the issue persists, contact Cisco Technical Assistance Center (TAC), who can help you verify the issue and propose a solution appropriate to your deployment.

History for Security Intelligence Block Listing

Feature	Version	Details
New Security Intelligence categories	All	Talos has added the following new Security Intelligence categories: • banking_fraud • ioc • high_risk • link_sharing • malicious • newly_seen • spyware You should update your access control and DNS policies to address the new categories, and check periodically for future changes. New/modified pages: Security Intelligence tab, Networks and URLs sub-tabs; DNS rules in DNS policies Supported platforms: FMC



CHAPTER 63

DNS Policies

The following topics explain DNS policies, DNS rules, and how to deploy DNS policies to managed devices.

- DNS Policy Overview, on page 1261
- DNS Policy Components, on page 1262
- License Requirements for DNS Policies, on page 1263
- Requirements and Prerequisites for DNS Policies, on page 1263
- Managing DNS Policies, on page 1263
- DNS Rules, on page 1265
- DNS Policy Deploy, on page 1272

DNS Policy Overview

DNS-based Security Intelligence allows you to block traffic based on the domain name requested by a client, using a Security Intelligence Block list. Cisco provides domain name intelligence you can use to filter your traffic; you can also configure custom lists and feeds of domain names tailored to your deployment.

Traffic on a DNS policy Block list is immediately blocked and therefore is not subject to any further inspection—not for intrusions, exploits, malware, and so on, but also not for network discovery. You can use a Security Intelligence Do Not Block list to override a Block list and force access control rule evaluation, and, recommended in passive deployments, you can use a "monitor-only" setting for Security Intelligence filtering. This allows the system to analyze connections that would have been blocked by a Block list, but also logs the match to the Block list and generates an end-of-connection Security Intelligence event.

Note DNS-based Security Intelligence may not work as intended for a domain name unless the DNS server deletes a domain cache entry due to expiration, or a client's DNS cache or the local DNS server's cache is cleared or expires.

You configure DNS-based Security Intelligence using a DNS policy and associated DNS rules. To deploy it to your devices, you must associate your DNS policy with an access control policy, then deploy your configuration to managed devices.

DNS Policy Components

A DNS policy allows you to block connections based on domain name, using a Block list, or exempt such connections from this type of blocking using a Do Not Block list. The following list describes the configurations you can change after creating a DNS policy.

Name and Description

Each DNS policy must have a unique name. A description is optional.

In a multidomain deployment, policy names must be unique within the domain hierarchy. The system may identify a conflict with the name of a policy you cannot view in your current domain.

Rules

Rules provide a granular method of handling network traffic based on the domain name. Rules in a DNS policy are numbered, starting at 1. The system matches traffic to DNS rules in top-down order by ascending rule number.

When you create a DNS policy, the system populates it with a default Global Whitelist for DNS rule and a default Global Block List for DNS rule. Both rules are fixed to the first position in their respective categories. You cannot modify these rules, but you can disable them.

In a multidomain deployment, the system also adds Descendant DNS Whitelists and Descendant DNS Block Lists rules to DNS policies in ancestor domains. These rules are fixed to the second position in their respective categories.



Note If multitenancy is enabled for your Firepower Management Center, the system is organized into a hierarchy of domains, including ancestor and descendant domains. These domains are distinct and separate from the domain names used in DNS management.

A descendant list contains the domains on the Block or Do Not Block lists of Firepower System subdomain users. From an ancestor domain, you cannot view the contents of descendant lists. If you do not want subdomain users to add domains to a Block or Do Not Block list:

- disable the descendant list rules, and
- enforce Security Intelligence using the access control policy inheritance settings

The system evaluates rules in the following order:

- Global Whitelist for DNS rule (if enabled)
- Descendant DNS Whitelists rule (if enabled)
- Rules with a Whitelist action
- Global Block List for DNS rule (if enabled)
- Descendant DNS Block Lists rule (if enabled)
- Rules with an action other than Whitelist

Usually, the system handles DN-based network traffic according to the *first* DNS rule where *all* the rule's conditions match the traffic. If no DNS rules match the traffic, the system continues evaluating the traffic based on the associated access control policy's rules. DNS rule conditions can be simple or complex.

License Requirements for DNS Policies

FTD License

Threat

Classic License

Protection

Requirements and Prerequisites for DNS Policies

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Managing DNS Policies

Use the DNS Policy page (**Policies** > **Access Control** > **DNS**) to manage custom DNS policies. In addition to custom policies that you create, the system provides the Default DNS Policy, which uses the default Block list and Do Not Block list. You can edit and use this system-provided custom policy. In a multidomain deployment, this default policy uses the default Global DNS Block List, Global DNS Do Not Block List, Descendant DNS Block lists, and Descendant DNS Do Not Block lists, and can only be edited in the Global domain.

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Policies** > **Access Control** > **DNS**.

Step 2 Manage your DNS policy:

- Compare—To compare DNS policies, click **Compare Policies** and proceed as described in Comparing Policies, on page 355.
- Copy—To copy a DNS policy, click **Copy** () and proceed as described in Editing DNS Policies, on page 1264.
- Create—To create a new DNS policy, click **Add DNS Policy** and proceed as described in Creating Basic DNS Policies, on page 1264.
- Delete—To delete a DNS policy, click **Delete** (), then confirm you want to delete the policy.
- Edit—To modify an existing DNS policy, click **Edit** ()) and proceed as described in Editing DNS Policies, on page 1264.

Creating Basic DNS Policies

Procedure

Step 1 Choose **Policies** > **Access Control** > **DNS**.

Step 2 Click **Add DNS Policy**.

Step 3 Give the policy a unique **Name** and, optionally, a **Description**.

Step 4 Click **Save**.

What to do next

Configure the policy. See Editing DNS Policies, on page 1264.

Editing DNS Policies

Only one person should edit a DNS policy at a time, using a single browser window. If multiple users attempt to save the same policy, only the first set of saved changes are retained.

To protect the privacy of your session, after thirty minutes of inactivity on the policy editor, a warning appears. After sixty minutes, the system discards your changes.

Procedure

Step 1 Choose **Policies** > **Access Control** > **DNS**.

Step 2 Click **Edit** () next to the DNS policy you want to edit.

If **View** (🔍) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Edit your DNS policy:

- Name and Description—To change the name or description, click the field and type the new information.
- Rules—To add, categorize, enable, disable, or otherwise manage DNS rules, click **Rules** and proceed as described in Creating and Editing DNS Rules, on page 1266.

Step 4 Click **Save**.

What to do next

- Optionally, further configure the new policy as described in Logging Connections with Security Intelligence, on page 2220.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

DNS Rules

DNS rules handle traffic based on the domain name requested by a host. As part of Security Intelligence, this evaluation happens after any traffic decryption, and before access control evaluation.

The system matches traffic to DNS rules in the order you specify. In most cases, the system handles network traffic according to the *first* DNS rule where *all* the rule's conditions match the traffic.

In addition to its unique name, each DNS rule has the following basic components:

State

By default, rules are enabled. If you disable a rule, the system does not use it to evaluate network traffic, and stops generating warnings and errors for that rule.

Position

Rules in a DNS policy are numbered, starting at 1. The system matches traffic to rules in top-down order by ascending rule number. With the exception of Monitor rules, the first rule that traffic matches is the rule that handles that traffic.

Conditions

Conditions specify the specific traffic the rule handles. A DNS rule must contain a DNS feed or list condition, and can also match traffic by security zone, network, or VLAN.

Action

A rule's action determines how the system handles matching traffic:

- Traffic with a **Whitelist** action is allowed, subject to further access control inspection.

- Monitored traffic is subject to further evaluation by remaining rules on the DNS Block list. If the traffic does not match a DNS Block list rule, it is inspected with access control rules. The system logs a Security Intelligence event for the traffic.
- Traffic on a Block list is dropped without further inspection. You can also return a Domain Not Found response, or redirect the DNS query to a sinkhole server.

Related Topics

About Security Intelligence, on page 1249

Creating and Editing DNS Rules

In a DNS policy, you can add up to a total of 32767 DNS lists to the Block list and Do Not Block list rules; that is, the number of lists in the DNS policy cannot exceed 32767.

Procedure

Step 1 In the DNS policy editor, you have the following options:

- To add a new rule, click **Add DNS Rule**.
- To edit an existing rule, click **Edit** ( ).

Step 2 Enter a **Name**.

Step 3 Configure the rule components, or accept the defaults:

- Action—Choose a rule **Action**; see DNS Rule Actions, on page 1268.
- Conditions—Configure the rule's conditions; see DNS Rule Conditions, on page 1269.
- Enabled—Specify whether the rule is **Enabled**.

Step 4 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

DNS Rule Management

The **Rules** tab of the DNS policy editor allows you to add, edit, move, enable, disable, delete, and otherwise manage DNS rules within your policy.

For each rule, the policy editor displays its name, a summary of its conditions, and the rule action. Other icons represent **Warning** ( ), **Error** ( ), and other important **Information** ( ). Disabled rules are dimmed and marked `(disabled)` beneath the rule name.

Enabling and Disabling DNS Rules

When you create a DNS rule, it is enabled by default. If you disable a rule, the system does not use it to evaluate network traffic and stops generating warnings and errors for that rule. When viewing the list of rules

in a DNS policy, disabled rules are dimmed, although you can still modify them. Note that you can also enable or disable a DNS rule using the DNS rule editor.

Procedure

Step 1 In the DNS policy editor, right-click the rule and choose a rule state.

Step 2 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

DNS Rule Order Evaluation

Rules in a DNS policy are numbered, starting at 1. The system matches traffic to DNS rules in top-down order by ascending rule number. In most cases, the system handles network traffic according to the *first* DNS rule where *all* the rule's conditions match the traffic:

- For Monitor rules, the system logs the traffic, then continues evaluating traffic against lower-priority DNS Block list rules.
- For non-Monitor rules, the system does **not** continue to evaluate traffic against additional, lower-priority DNS rules after that traffic matches a rule.

Note the following regarding rule order:

- The Global Whitelist for DNS is always first, and takes precedence over all other rules.
- The Descendant DNS Whitelists rule only appears in multidomain deployments, in non-leaf domains. It is always second, and takes precedence over all other rules except the Global Whitelist for DNS.
- The Do-Not-Block List section precedes the Block List section; Do-Not-Block List rules always take precedence over other rules.
- The Global Block List for DNS is always first in the Block List section, and takes precedence over all other Monitor and Block list rules.
- The Descendant DNS Block Lists rule only appears in multidomain deployments, in non-leaf domains. It is always second in the Block List section, and takes precedence over all other Monitor and Block list rules except the Global Block List.
- The Block List section contains Monitor and Block list rules.
- When you first create a DNS rule, the system positions it last in the Do-Not-Block List section if you assign a **Do Not Block** action, or last in the Block List section if you assign any other action.

You can drag and drop rules to reorder them.

DNS Rule Actions

Every DNS rule has an *action* that determines the following for matching traffic:

- handling—foremost, the rule action governs whether the system will block, not block, or monitor traffic that matches the rule's conditions, based on a Block or Do Not Block list
- logging—the rule action determines when and how you can log details about matching traffic

Whitelist Action

The **Whitelist** action allows traffic to pass to the next phase of inspection, which is access control rules.

The system does not log whitelist matches. Logging of these connections depends on their eventual disposition.

Monitor Action

The **Monitor** action is designed to force connection logging; matching traffic is neither immediately allowed nor blocked. Rather, traffic is matched against additional rules to determine whether to permit or deny it. The first non-Monitor DNS rule matched determines whether the system blocks the traffic. If there are no additional matching rules, the traffic is subject to access control evaluation.

For connections monitored by a DNS policy, the system logs end-of-connection Security Intelligence and connection events to the Firepower Management Center database.

Block Actions

These actions block traffic without further inspection of any kind:

- The **Drop** action drops the traffic.
- The **Domain Not Found** action returns a non-existent internet domain response to the DNS query, which prevents the client from resolving the DNS request.
- The **Sinkhole** action returns a sinkhole object's IPv4 or IPv6 address in response to the DNS query (A and AAAA records only). The sinkhole server can log, or log and block, follow-on connections to the IP address. If you configure a **Sinkhole** action, you must also configure a sinkhole object.

For a connection blocked based on the **Drop** or **Domain Not Found** actions, the system logs beginning-of-connection Security Intelligence and connection events. Because blocked traffic is immediately denied without further inspection, there is no unique end of connection to log.

For a connection blocked based on the **Sinkhole** action, logging depends on the sinkhole object configuration. If you configure your sinkhole object to only log sinkhole connections, the system logs end-of-connection connection events for the follow-on connection. If you configure your sinkhole object to log and block sinkhole connections, the system logs beginning-of-connection connection events for the follow-on connection, then blocks that connection.

Note On an ASA FirePOWER device, if you configure a DNS rule with a sinkhole action, and traffic matches the rule, the ASA blocks the follow-on sinkhole connection by default. As a workaround, run the following commands from the ASA command line:

```
asa(config)# policy-map global_policy
asa(config-pmap)# class inspection_default
asa(config-pmap-c)# no inspect dns preset_dns_map
```

If the ASA continues to block the connection, contact Support.

Related Topics

How Rules and Policy Actions Affect Logging, on page 2209

DNS Rule Conditions

A DNS rule's conditions identify the type of traffic that rule handles. Conditions can be simple or complex. You must define a DNS feed or list condition within a DNS rule. You can also optionally control traffic by security zone, network, or VLAN.

When adding conditions to a DNS rule:

- If you do not configure a particular condition for a rule, the system does not match traffic based on that criterion.
- You can configure multiple conditions per rule. Traffic must match **all** the conditions in the rule for the rule to apply to traffic. For example, a rule with a DNS feed or list condition and network condition but no VLAN tag condition evaluates traffic based on the domain name and source or destination, regardless of any VLAN tagging in the session.
- For each condition in a rule, you can add up to 50 criteria. Traffic that matches **any** of a condition's criteria satisfies the condition. For example, you can use a single rule to block traffic based on up to 50 DNS lists and feeds.

Controlling Traffic Based on DNS and Security Zone

Zone conditions in DNS rules allow you to control traffic by its source security zone. A *security zone* is a grouping of one or more interfaces, which may be located across multiple devices.

Procedure

Step 1 In the DNS rule editor, click **Zones**.

Step 2 Find and select the zones you want to add from the **Available Zones**. To search for zones to add, click the **Search by name** prompt above the **Available Zones** list, then type a zone name. The list updates as you type to display matching zones.

Step 3 Click to select a zone, or right-click and then select **Select All**.

Step 4 Click **Add to Source**, or drag and drop.

Step 5 Save or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Controlling Traffic Based on DNS and Network

Network conditions in DNS rules allow you to control traffic by its source IP address. You can explicitly specify the source IP addresses for the traffic you want to control.

Procedure

Step 1 In the DNS rule editor, click **Networks**.

Step 2 Find and select the networks you want to add from the **Available Networks**, as follows:

- To add a network object on the fly, which you can then add to the condition, click **Add** ( ) above the **Available Networks** list and proceed as described in Creating Network Objects, on page 403.
- To search for network objects to add, click the **Search by name or value** prompt above the **Available Networks** list, then type an object name or the value of one of the object's components. The list updates as you type to display matching objects.

Step 3 Click **Add to Source**, or drag and drop.

Step 4 Add any source IP addresses or address blocks that you want to specify manually. Click the **Enter an IP address** prompt below the **Source Networks** list; then type an IP address or address block and click **Add**.

The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Step 5 Save or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Controlling Traffic Based on DNS and VLAN

VLAN conditions in DNS rules allow you to control VLAN-tagged traffic. The system uses the innermost VLAN tag to identify a packet by VLAN.

When you build a VLAN-based DNS rule condition, you can manually specify VLAN tags. Alternately, you can configure VLAN conditions with VLAN tag *objects*, which are reusable and associate a name with one or more VLAN tags.

Procedure

Step 1 In the DNS rule editor, select **VLAN Tags**.

Step 2 Find and select the VLANs you want to add from the **Available VLAN Tags**, as follows:

- To add a VLAN tag object on the fly, which you can then add to the condition, click **Add** ( ) above the Available VLAN Tags list and proceed as described in Creating VLAN Tag Objects, on page 406.
- To search for VLAN tag objects and groups to add, click the **Search by name or value** prompt above the **Available VLAN Tags** list, then type either the name of the object, or the value of a VLAN tag in the object. The list updates as you type to display matching objects.

Step 3 Click **Add to Rule**, or drag and drop.

Step 4 Add any VLAN tags that you want to specify manually. Click the **Enter a VLAN Tag** prompt below the **Selected VLAN Tags** list; then type a VLAN tag or range and click **Add**. You can specify any VLAN tag from 1 to 4094; use a hyphen to specify a range of VLAN tags.

The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal VLAN tags to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Step 5 Save or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Controlling Traffic Based on DNS List, Feed, or Category

DNS conditions in DNS rules allow you to control traffic if a DNS list, feed, or category contains the domain name requested by the client. You must define a DNS condition in a DNS rule.

Regardless of whether you add a global or custom Block or Do Not Block list to a DNS condition, the system applies the configured rule action to the traffic. For example, if you add the Global Do Not Block List to a rule, and configure a **Drop** action, the system blocks all traffic that should have been allowed to pass to the next phase of inspection.

Procedure

Step 1 In the DNS rule editor, click **DNS**.

Step 2 Find and select the DNS lists and feeds you want to add from the **DNS Lists and Feeds**, as follows:

- To add a DNS list or feed on the fly, which you can then add to the condition, click **Add** ( ) above the **DNS Lists and Feeds** list and proceed as described in Creating Security Intelligence Feeds, on page 434 .
- To search for DNS lists, feeds, or categories to add, click the **Search by name or value** prompt above the **DNS Lists and Feeds** list, then type an object name or the value of one of the object's components. The list updates as you type to display matching objects.
- For descriptions of the system-provided threat categories, see Security Intelligence Categories, on page 1255.

Step 3 Click **Add to Rule**, or drag and drop.

Step 4 Save or continue editing the rule.

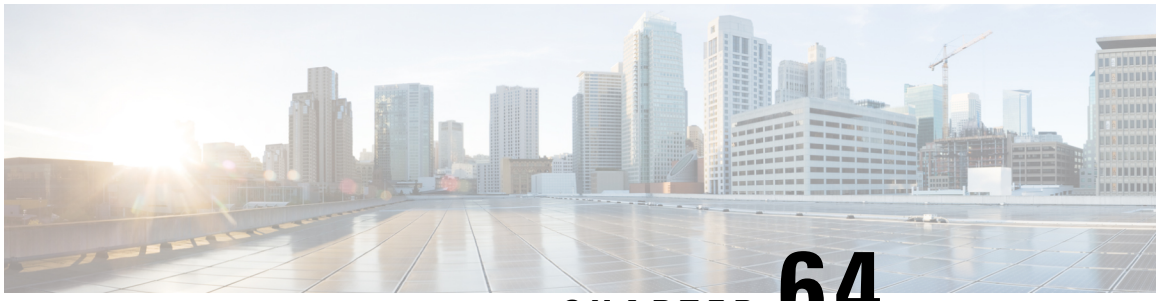
What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

DNS Policy Deploy

After you finish updating your DNS policy configuration, you must deploy it as part of access control configuration.

- Associate your DNS policy with an access control policy, as described in Configure Security Intelligence, on page 1252.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.



CHAPTER 64

Prefiltering and Prefilter Policies

- About Prefiltering, on page 1273
- Best Practices for Prefiltering, on page 1276
- Encapsulated Traffic Handling Best Practices, on page 1277
- Requirements and Prerequisites for Prefilter Policies, on page 1277
- Configure Prefiltering, on page 1278
- Tunnel Zones and Prefiltering, on page 1282
- Large Flow Offloads, on page 1286

About Prefiltering

Prefiltering is the first phase of access control, before the system performs more resource-intensive evaluation. Prefiltering is simple, fast, and early. Prefiltering uses limited outer-header criteria to quickly handle traffic. Compare this to subsequent evaluation, which uses inner headers and has more robust inspection capabilities.

Configure prefiltering to:

- Improve performance— The sooner you exclude traffic that does not require inspection, the better. You can fastpath or block certain types of plaintext, passthrough tunnels based on their outer encapsulation headers, without inspecting their encapsulated connections. You can also fastpath or block any other connections that benefit from early handling.
- Tailor deep inspection to encapsulated traffic—You can rezone certain types of tunnels, so that you can later handle their encapsulated connections using the same inspection criteria. Rezoning is necessary because after prefiltering, access control uses inner headers.

Prefiltering vs Access Control

Prefilter and access control policies both allow you to block and trust traffic, though the prefiltering "trust" functionality is called "fastpathing" because it skips more inspection. The following table explains this and other differences between prefiltering and access control, to help you decide whether to configure custom prefiltering.

If you do not configure custom prefiltering, you can only approximate—not replicate—prefilter functionality with early-placed Block and Trust rules in the access control policy.

Characteristic	Prefiltering	Access Control	For more information, see...
Primary function	Quickly fastpath or block certain types of plaintext, passthrough tunnels (see Encapsulation Conditions, on page 375), or tailor subsequent inspection to their encapsulated traffic. Fastpath or block any other connections that benefit from early handling.	Inspect and control all network traffic, using simple or complex criteria, including contextual information and deep inspection results.	About Prefiltering, on page 1273
Implementation	Prefilter policy. The prefilter policy is invoked by the access control policy.	Access control policy. The access control policy is a main configuration. In addition to invoking subpolicies, access control policies have their own rules.	About Prefilter Policies, on page 1279 Associating Other Policies with Access Control, on page 1213
Sequence within access control	First. The system matches traffic to prefilter criteria before all other access control configurations.	—	—
Rule actions	Fewer. You can stop further inspection (Fastpath and Block) or allow further analysis with the rest of access control (Analyze).	More. Access control rules have a larger variety of actions, including monitoring, deep inspection, block with reset, and interactive blocking.	Tunnel and Prefilter Rule Components, on page 1281 Access Control Rule Actions, on page 1223
Bypass capability	Fastpath rule action. Fastpathing traffic in the prefilter stage bypasses all further inspection and handling, including: • Security Intelligence • authentication requirements imposed by an identity policy • SSL decryption • access control rules • deep inspection of packet payloads • discovery • rate limiting	Trust rule action. Traffic trusted by access control rules is only exempt from deep inspection and discovery.	Introduction to Access Control Rules, on page 1215

Characteristic	Prefiltering	Access Control	For more information, see...
Rule criteria	Limited. Rules in the prefilter policy use simple network criteria: IP address, VLAN tag, port, and protocol. For tunnels, tunnel endpoint conditions specify the IP address of the routed interfaces of the network devices on either side of the tunnel.	Robust. Access control rules use network criteria, but also user, application, requested URL, and other contextual information available in packet payloads. Network conditions specify the IP address of source and destination hosts.	Tunnel vs Prefilter Rules, on page 1280 Rule Condition Types, on page 363
IP headers used (tunnel handling)	Outermost. Using outer headers allows you to handle entire plaintext, passthrough tunnels. For nonencapsulated traffic, prefiltering still uses "outer" headers—which in this case are the only headers.	Innermost possible. For a nonencrypted tunnel, access control acts on its individual encapsulated connections, not the tunnel as a whole.	Passthrough Tunnels and Access Control, on page 1275
Rezone encapsulated connections for further analysis	Rezones tunneled traffic. Tunnel zones allow you to tailor subsequent inspection to prefiltered, encapsulated traffic.	Uses tunnel zones. Access control uses the tunnel zones you assign during prefiltering.	Tunnel Zones and Prefiltering, on page 1282
Connection logging	Fastpathed and blocked traffic only. Allowed connections may still be logged by other configurations.	Any connection.	Other Connections You Can Log, on page 2209
Supported devices	Firepower Threat Defense only.	All.	Best Practices for Prefiltering, on page 1276

Passthrough Tunnels and Access Control

Plaintext (nonencrypted) tunnels can encapsulate multiple connections, often flowing between discontinuous networks. These tunnels are especially useful for routing custom protocols over IP networks, IPv6 traffic over IPv4 networks, and so on.

An outer *encapsulation header* specifies the source and destination IP addresses of the *tunnel endpoints*—the routed interfaces of the network devices on either side of the tunnel. Inner *payload headers* specify the source and destination IP addresses of the encapsulated connections' actual endpoints.

Often, network security devices handle plaintext tunnels as *passthrough* traffic. That is, the device is not one of the tunnel endpoints. Instead, it is deployed between the tunnel endpoints and monitors the traffic flowing between them.

Some network security devices, such as Cisco ASA firewalls running Cisco ASA Software (rather than Firepower Threat Defense), enforce security policies using outer IP headers. Even for plaintext tunnels, these devices have no control over or insight into individual encapsulated connections and their payloads.

By contrast, the Firepower System leverages access control as follows:

- Outer header evaluation—First, prefiltering uses outer headers to handle traffic. You can block or fastpath entire plaintext, passthrough tunnels at this stage.
- Inner header evaluation—Next, the rest of access control (and other features such as QoS) use the innermost detectable level of headers to ensure the most granular level of inspection and handling possible.

  If a passthrough tunnel is not encrypted, the system acts on its individual encapsulated connections at this stage. You must *rezone* a tunnel (see Tunnel Zones and Prefiltering, on page 1282) to act on all its encapsulated connections.

Access control has no insight into encrypted passthrough tunnels. For example, access control rules see a passthrough VPN tunnel as one connection. The system handles the entire tunnel using only the information in its outer, encapsulation header.

Best Practices for Prefiltering

Consider the following guidelines and limitations for prefiltering:

Best Practices for Fastpath Prefiltering

When you use the fastpath action in a prefilter rule, the matching traffic bypasses inspection and is simply transmitted through the device. Use this action for traffic that you can trust and that would not benefit from any of the security features available.

The following types of traffic are ideal for fastpathing. For example, you could configure the rules to fastpath any traffic from or to the IP addresses of the endpoints or servers. You can further limit the rule based on ports used.

- Site-to-site VPN traffic that is going through the device. That is, the device is not an endpoint in the VPN topology.
- Scanner traffic. Scanner probes can create a lot of false-positive responses from intrusion policies.
- Voice/video.
- Backups.
- Management traffic (sftunnel) that traverses Firepower Threat Defense devices. Performing deep inspection on management traffic (using access control policies) can cause issues. You can prefilter based on port TCP/8305 between the management center and managed devices.

Model Requirements

Prefiltering is supported on Firepower Threat Defense devices only. Prefilter configurations have no effect on other devices.

Prefilter-like Capabilities on Non-FTD Devices

For Classic devices (7000/8000 series, ASA FirePOWER, NGIPSv):

- Use early-placed Trust and Block access control rules to approximate prefilter functionality, keeping in mind the differences between the two features. See Prefiltering vs Access Control, on page 1273.
- Match entire GRE-encapsulated tunnels using access control rules, with some limitations. See Port and ICMP Code Conditions, on page 373.
- (8000 series only) Device-specific fastpath rules can bypass access control, but cannot block traffic. See Configure Fastpath Rules (8000 Series), on page 243.

Encapsulated Traffic Handling Best Practices

This topic discusses guidelines for the following types of encapsulated traffic:

- Generic Routing Encapsulation (GRE)
- Point-to-Point Protocol (PPTP)
- IPinIP
- IPv6inIP
- Teredo

Devices cannot set reserved bits in the GRE packet header

Traffic is dropped if any device sets the values of bits 5 through 12 in the packet header to something other than 0 (zero) and if a signature is configured for Snort. For more information about the GRE packet header, see the introduction to RFC 1701.

How tunnels work with interfaces

Prefilter and access control policy rules are applied to all tunnel types on routed, transparent, inline-set, inline-tap, and passive interfaces.

References

For more information about the GRE and PPTP protocols, see the following:

- RFC 1701, RFC 2784, and RFC 2890 (GRE protocol v0)
- RFC 2637 (PPTP and GRE protocol v1)

Requirements and Prerequisites for Prefilter Policies

Model Support

FTD

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Configure Prefiltering

To perform custom prefiltering, configure and deploy prefilter policies to managed devices, as a part of access control.

Only one person should edit a policy at a time, using a single browser window. If multiple users save the same policy, the last saved changes are retained. For your convenience, the system displays information on who (if anyone) is currently editing each policy. To protect the privacy of your session, a warning appears after 30 minutes of inactivity on the policy editor. After 60 minutes, the system discards your changes.

Procedure

Step 1 Choose **Policies** > **Access Control** > **Prefilter**.

Step 2 Click **New Policy** to create a custom prefilter policy.

A new prefilter policy has no rules and a default action of Analyze all tunnel traffic. It performs no logging or tunnel rezoning. You can also **Copy** ( ) or **Edit** ( ) an existing policy.

Step 3 Configure the prefilter policy's default action and its logging options.

- Default action—Choose a default action for supported plaintext, passthrough tunnels: **Analyze all tunnel traffic** (with access control) or **Block all tunnel traffic**.
- Default action logging—Click **Logging** ( ) next to the default action; see Logging Connections with a Policy Default Action, on page 2221. You can configure default action logging for blocked tunnels only.

Step 4 Configure tunnel and prefilter rules.

In a custom prefilter policy, you can use both kinds of rule, in any order. Create rules depending on the specific type of traffic you want to match and the actions or further analysis you want to perform; see Tunnel vs Prefilter Rules, on page 1280.

Caution Exercise caution when using tunnel rules to assign tunnel zones. Connections in rezoned tunnels may not match security zone constraints in later evaluation. For more information, see Tunnel Zones and Prefiltering, on page 1282.

For detailed information on configuring rule components, see Tunnel and Prefilter Rule Components, on page 1281 and Rule Management: Common Characteristics, on page 361.

Step 5 Evaluate rule order. To move a rule, click and drag or use the right-click menu to cut and paste.

Properly creating and ordering rules is a complex task, but one that is essential to building an effective deployment. If you do not plan carefully, rules can preempt other rules or contain invalid configurations. For more information, see Best Practices for Access Control Rules, on page 1194.

Step 6 Save the prefilter policy.

Step 7 For configurations that support tunnel zone constraints, appropriately handle rezoned tunnels.

Match connections in rezoned tunnels by using tunnel zones as source zone constraints; see Configuring Interface Conditions, on page 367.

Step 8 Associate the prefilter policy with the access control policy deployed to your managed devices.

See Associating Other Policies with Access Control, on page 1213.

Step 9 Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Note When you deploy a prefilter policy, its rules are not applied on the existing tunnel sessions. Hence, traffic on an existing connection is not bound by the new policy that is deployed. In addition, the policy hit count is incremented only for the first packet of a connection that matches a policy. Thus, the traffic of an existing connection that could match a policy is omitted from the hit count. To have the policy rules effectively applied, clear the existing tunnel sessions, and then deploy the policy.

About Prefilter Policies

Prefiltering is a policy-based feature. In the Firepower System, an access control policy is a main configuration that invokes subpolicies and other configurations, including a prefilter policy.

Policy Components: Rules and Default Action

In a prefilter policy, *tunnel rules*, *prefilter rules*, and a *default action* handle network traffic:

- Tunnel and prefilter rules—First, rules in a prefilter policy handle traffic in the order you specify. Tunnel rules match specific tunnels only and support rezoning. Prefilter rules have a wider range of constraints and do not support rezoning. For more information, see Tunnel vs Prefilter Rules, on page 1280.
- Default action (tunnels only)—If a tunnel does not match any rules, the default action handles it. The default action can block these tunnels, or continue access control on their individual encapsulated connections. You cannot rezone tunnels with the default action.

  There is no default action for nonencapsulated traffic. If a nonencapsulated connection does not match any prefilter rules, the system continues with access control.

Connection Logging

You can log connections fastpathed and blocked by the prefilter policy; see Other Connections You Can Log, on page 2209.

Connection events contain information on whether and how logged connections—including entire tunnels—were prefiltered. You can view this information in event views (workflows), dashboards, and reports, and use it as correlation criteria. Keep in mind that because fastpathed and blocked connections are not subject to deep inspection, associated connection events contain limited information.

Default Prefilter Policy

Every access control policy has an associated prefilter policy.

The system uses a default policy if you do not configure custom prefiltering. Initially, this system-provided policy passes all traffic to the next phase of access control. You can change the policy's default action and configure its logging options, but you cannot add rules to it or delete it.

Prefilter Policy Inheritance and Multitenancy

Access control uses a hierarchical implementation that complements multitenancy. Along with other advanced settings, you can lock a prefilter policy association, enforcing that association in all descendant access control policies. For more information, see Access Control Policy Inheritance, on page 1191.

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain. The default prefilter policy belongs to the Global domain.

Tunnel vs Prefilter Rules

Whether you configure a tunnel or prefilter rule depends on the specific type of traffic you want to match and the actions or further analysis you want to perform.

Characteristic	Tunnel Rules	Prefilter Rules
Primary function	Quickly fastpath, block, or rezone plaintext, passthough tunnels.	Quickly fastpath or block any other connection that benefits from early handling.
Encapsulation and port/protocol criteria	Encapsulation conditions match only plaintext tunnels over selected protocols, listed in Encapsulation Conditions, on page 375.	Port conditions can use a wider range of port and protocol constraints than tunnel rules; see Port and ICMP Code Conditions, on page 373.
Network criteria	Tunnel endpoint conditions constrain the endpoints of the tunnels you want to handle; see Tunnel Endpoint Conditions, on page 370.	Network conditions constrain the source and destination hosts in each connection; see Network Conditions, on page 368.
Direction	Bidirectional or unidirectional (configurable). Tunnel rules are bidirectional by default, so they can handle all traffic between tunnel endpoints.	Unidirectional only (nonconfigurable). Prefilter rules match source-to-destination traffic only.
Rezone sessions for further analysis	Supported, using tunnel zones; see Tunnel Zones and Prefiltering, on page 1282.	Not supported.

Tunnel and Prefilter Rule Components

State (Enabled/Disabled)

By default, rules are enabled. If you disable a rule, the system does not use it and stops generating warnings and errors for that rule.

Position

Rules are numbered, starting at 1. The system matches traffic to rules in top-down order by ascending rule number. The first rule that traffic matches is the rule that handles that traffic, regardless of rule type (tunnel vs prefilter).

Action

A rule's action determines how the system handles and logs matching traffic:

- Fastpath—Exempts matching traffic from all futher inspection and control, including access control, identity requirements, and rate limiting. Fastpathing a tunnel fastpaths all encapsulated connections.
- Block—Blocks matching traffic without further inspection of any kind. Blocking a tunnel blocks all encapsulated connections.
- Analyze—Allows traffic to continue to be analyzed by the rest of access control, using inner headers. If passed by access control and any related deep inspection, this traffic may also be rate limited. For tunnel rules, enables rezoning with the Assign Tunnel Zone option.

Direction (Tunnel Rules Only)

A tunnel rule's direction determines how the system source and destination criteria:

- Match tunnels only from source (unidirectional)—Match source-to-destination traffic only. Matching traffic must originate from one of the specified source interfaces or tunnel endpoints, and leave through one of the destination interfaces or tunnel endpoints.
- Match tunnels from source and destination (bidirectional)—Match both source-to-destination traffic and destination-to-source traffic. The effect is identical to writing two unidirectional rules, one the mirror of the other.

Prefilter rules are always unidirectional.

Assign Tunnel Zone (Tunnel Rules Only)

In a tunnel rule, assigning a tunnel zone (whether existing or created on the fly), *rezones* matching tunnels. Rezoning requires the Analyze action.

Rezoning a tunnel allows other configurations—such as access control rules—to recognize all the tunnel's encapsulated connections as belonging together. By using a tunnel's assigned tunnel zone as an interface constraint, you can tailor inspection to its encapsulated connections. For more information, see Tunnel Zones and Prefiltering, on page 1282.

Caution Exercise caution when assigning tunnel zones. Connections in rezoned tunnels may not match security zone constraints in later evaluation. See Using Tunnel Zones, on page 1283 for a brief walkthrough of a tunnel zone implementation, and a discussion of the implications of rezoning without explicitly handling rezoned traffic.

Conditions

Conditions specify the specific traffic the rule handles. Traffic must match all a rule's conditions to match the rule. Each condition type has its own tab in the rule editor.

You can prefilter traffic using the following *outer-header* constraints:

- Interface—Interface Conditions, on page 366
- Network—Tunnel Endpoint Conditions, on page 370 or Network Conditions, on page 368
- Port—Encapsulation Conditions, on page 375 or Port and ICMP Code Conditions, on page 373
- VLAN—VLAN Conditions, on page 372

You must constrain tunnel rules by encapsulation protocol.

Logging

A rule's logging settings govern the records the system keeps of the traffic it handles.

In tunnel and prefilter rules, you can log fastpathed and blocked traffic (the Fastpath and Block actions). For traffic subject to further analysis (the Analyze action), logging in the prefilter policy is disabled, although matching connections may still be logged by other configurations. For more information, see Logging Connections with Tunnel and Prefilter Rules, on page 2218.

Comments

Each time you save changes to a rule you can add comments. For example, you might summarize the overall configuration for the benefit of other users, or note when you change a rule and the reason for the change.

You cannot edit or delete these comments after you save the rule.

Related Topics

Best Practices for Access Control Rules, on page 1194

Tunnel Zones and Prefiltering

Tunnel zones allow you to use prefiltering to tailor subsequent traffic handling to encapsulated connections.

A special mechanism is required because usually, the system handles traffic using the innermost detectable level of headers. This ensures the most granular level of inspection possible. But it also means that if a passthrough tunnel is not encrypted, the system acts on its individual encapsulated connections; see Passthrough Tunnels and Access Control, on page 1275.

Tunnel zones solve this problem. During the first phase of access control (prefiltering) you can use outer headers to identify certain types of plaintext, pasthrough tunnels. Then, you can *rezone* those tunnels by assigning a custom *tunnel zone*.

Rezoning a tunnel allows other configurations—such as access control rules—to recognize all the tunnel's encapsulated connections as belonging together. By using a tunnel's assigned tunnel zone as an interface constraint, you can tailor inspection to its encapsulated connections.

Despite its name, a tunnel zone is not a security zone. A tunnel zone does not represent a set of interfaces. It is more accurate to think of a tunnel zone as a tag that, in some cases, replaces the security zone associated with an encapsulated connection.

Caution For configurations that support tunnel zone constraints, connections in rezoned tunnels do **not** match security zone constraints. For example, after you rezone a tunnel, access control rules can match its encapsulated connections with their newly assigned *tunnel* zone, but not with any original *security* zone.

See Using Tunnel Zones, on page 1283 for a brief walkthough of a tunnel zone implementation, and a discussion of the implications of rezoning without explicitly handling rezoned traffic.

Configurations Supporting Tunnel Zone Constraints

Only access control rules support tunnel zone constraints.

No other configurations support tunnel zone constraints. For example, you cannot use QoS to rate limit a plaintext tunnel as a whole; you can only rate limit its individual encapsulated sessions.

Using Tunnel Zones

This example procedure summarizes how you might rezone GRE tunnels for further analysis, using tunnel zones. You can adapt the concepts described in this example to other scenarios where you need to tailor traffic inspection to connections encapsulated in plaintext, passthrough tunnels.

Consider a Firepower System deployment where your organization's internal traffic flows through the Trusted security zone. The Trusted security zone represents a set of sensing interfaces across multiple managed devices deployed in various locations. Your organization's security policy requires that you allow internal traffic after deep inspection for exploits and malware.

Internal traffic sometimes includes plaintext, passthrough, GRE tunnels between particular endpoints. Because the traffic profile of this encapsulated traffic is different from your "normal" interoffice activity—perhaps it is known and benign—you can limit inspection of certain encapsulated connections while still complying with your security policy.

In this example, after you deploy configuration changes:

- Plaintext, passthrough, GRE-encapsulated tunnels detected in the Trusted zone have their individual encapsulated connections evaluated by one set of intrusion and file policies.
- All other traffic in the Trusted zone is evaluated with a different set of intrusion and file policies.

You accomplish this task by *rezoning* GRE tunnels. Rezoning ensures that access control associates GRE-encapsulated connections with a custom *tunnel* zone, rather than their original Trusted *security* zone. Rezoning is required due to the way the Firepower System and access control handle encapsulated traffic; see Passthrough Tunnels and Access Control, on page 1275 and Tunnel Zones and Prefiltering, on page 1282.

Procedure

Step 1 Configure custom intrusion and file policies that tailor deep inspection to encapsulated traffic, and another set of intrusion and file policies tailored to nonencapsulated traffic.

Step 2 Configure custom prefiltering to rezone GRE tunnels flowing through the Trusted security zone.

Create a custom prefilter policy and associate it with access control. In that custom prefilter policy, create a tunnel rule (in this example, **GRE_tunnel_rezone**) and a corresponding tunnel zone (**GRE_tunnel**). For more information, see Configure Prefiltering, on page 1278.

Table 88: GRE_tunnel_rezone Tunnel Rule

Rule Component	Description
Interface object condition	Match internal-only tunnels by using the Trusted security zone as both a Source Interface Object and Destination Interface Object constraint.
Tunnel endpoint condition	Specify the source and destination endpoints for the GRE tunnels used in your organization. Tunnel rules are bidirectional by default. If you do not change the **Match tunnels from...** option, it does not matter which endpoints you specify as source and which as destination.
Encapsulation condition	Match GRE traffic.
Assign Tunnel Zone	Create the **GRE_tunnel** tunnel zone, and assign it to tunnels that match the rule.
Action	Analyze (with the rest of access control).

Step 3 Configure access control to handle connections in rezoned tunnels.

In the access control policy deployed to your managed devices, configure a rule (in this example, **GRE_inspection**) that handles the traffic you rezoned. For more information, see Create and Edit Access Control Rules, on page 1220.

Table 89: GRE_inspection Access Control Rule

Rule Component	Description
Security zone condition	Match rezoned tunnels by using the GRE_tunnel security zone as a Source Zone constraint; see Interface Conditions, on page 366.
Action	Allow, with deep inspection enabled. Choose the file and intrusion policies tailored to inspect encapsulated internal traffic.

Caution If you skip this step, the rezoned connections may match **any** access control rule not constrained by security zone. If the rezoned connections do not match any access control rules, they are handled by the access control policy default action. Make sure this is your intent.

Step 4 Configure access control to handle nonencapsulated connections flowing through the Trusted security zone.

In the same access control policy, configure a rule (in this example, **`internal_default_inspection`**) that handles non-rezoned traffic in the Trusted security zone.

Table 90: internal_default_inspection Access Control Rule

Rule Component	Description
Security zone condition	Match non-rezoned internal-only traffic by using the Trusted security zone as both a Source Zone and Destination Zone constraint.
Action	Allow, with deep inspection enabled. Choose the file and intrusion policies tailored to inspect nonencapsulated internal traffic.

Step 5 Evaluate the position of the new access control rules relative to preexisting rules. Change rule order if necessary.

If you place the two new access control rules next to each other, it does not matter which you place first. Because you rezoned GRE tunnels, the two rules cannot preempt each other.

Step 6 Save all changed configurations.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Creating Tunnel Zones

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Chose **Tunnel Zone** from the list of object types.

Step 3 Click **Add Tunnel Zone**.

Step 4 Enter a **Name** and, optionally, a **Description**.

Step 5 Click **Save**.

What to do next

- Assign tunnel zones to plaintext, passthrough tunnels as part of custom prefiltering; see Configure Prefiltering, on page 1278.

Large Flow Offloads

On devices that run FXOS (such as Firepower 4100/9300 chassis), certain traffic that you configure to be fastpathed by a prefilter policy is handled by the hardware (specifically, in the NIC), not by your Firepower Threat Defense software. Offloading these connection flows results in higher throughput and lower latency, especially for data-intensive applications such as large file transfers. This feature is especially useful for data centers. This is called *static flow offload*.

Offloaded flows continue to receive limited stateful inspection, such as basic TCP flag and option checking. The system can selectively escalate packets to the firewall system for further processing if necessary.

Examples of applications that can benefit from offloading large flows are:

- High Performance Computing (HPC) Research sites, where the Firepower Threat Defense device is deployed between storage and high compute stations. When one research site backs up using FTP file transfer or file sync over NFS, the large amount of data traffic affects all connections. Offloading FTP file transfer and file sync over NFS reduces the impact on other traffic.
- High Frequency Trading (HFT), where the Firepower Threat Defense device is deployed between workstations and the Exchange, mainly for compliance purposes. Security is usually not a concern, but latency is a major concern.

The following flows can be offloaded:

- Connections that are fastpathed by the prefilter policy.
- Standard or 802.1Q tagged Ethernet frames only.



Important For details, exceptions, and limitations to the above, see Flow Offload Limitations, on page 1286.

Use Static Flow Offload

To offload eligible traffic to hardware, create a prefilter policy rule that applies the **Fastpath** action. Use prefilter rules for TCP/UDP, and tunnel rules for GRE.

(Not recommended.) To disable static flow offload and as a by-product, dynamic flow-offload, use FlexConfig to run the **no flow-offload enable** command. After deployment, you will have to reload the device to implement the change. For information about this command, see the *Cisco ASA Series Command Reference*, available from https://www.cisco.com/c/en/us/support/security/asa-5500-series-next-generation-firewalls/products-command-reference-list.html.

Flow Offload Limitations

Not all flows can be offloaded. Even after offload, a flow can be removed from being offloaded under certain conditions. Following are some of the limitations:

Device Limitations

The feature is supported on the following devices:

- Firepower 9300 running FXOS 1.1.3 or higher.

Flows that cannot be offloaded

The following types of flows cannot be offloaded.

- Any flows that do not use IPv4 addressing, such as IPv6 addressing.
- Flows for any protocol other than TCP, UDP, and GRE.



Note PPTP GRE connections cannot be offloaded.

- Flows on interfaces configured in passive, inline, or inline tap mode. Routed and switched interfaces are the only types supported.
- Flows that require inspection by Snort or other inspection engines. In some cases, such as FTP, the secondary data channel can be offloaded although the control channel cannot be offloaded.
- IPsec and TLS/DTLS VPN connections that terminate on the device.
- Multicast flows in routed mode. They are supported in transparent mode if there are only two member interfaces in a bridge group.
- TCP state bypass flows. You cannot configure flow offload and TCP state bypass on the same traffic.
- Flows tagged with security groups.
- Reverse flows that are forwarded from a different cluster node, in case of asymmetric flows in a cluster.
- Centralized flows in a cluster, if the flow owner is not the control unit.

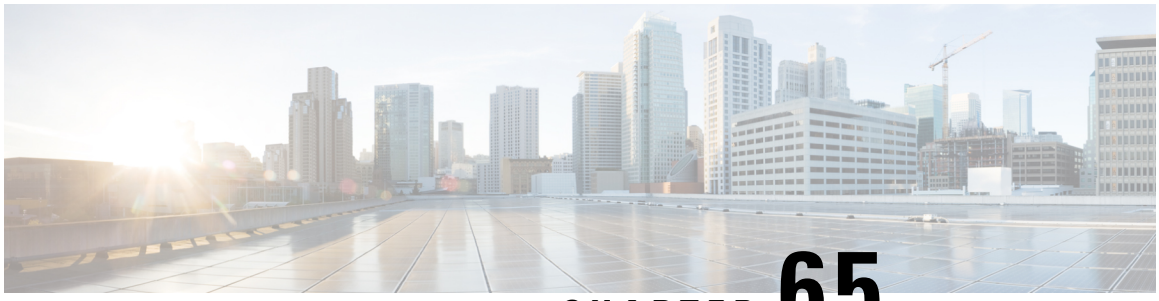
Additional Limitations

- Flow offload and Dead Connection Detection (DCD) are not compatible. Do not configure DCD on connections that can be offloaded.
- If more than one flow that matches flow offload conditions are queued to be offloaded at the same time to the same location on the hardware, only the first flow is offloaded. The other flows are processed normally. This is called a *collision*. Use the **show flow-offload flow** command in the CLI to display statistics for this situation.
- Although offloaded flows pass through FXOS interfaces, statistics for these flows do not appear on the logical device interface. Thus, logical device interface counters and packet rates do not reflect offloaded flows.

Conditions for reversing offload

After a flow is offloaded, packets within the flow are returned to the FTD for further processing if they meet the following conditions:

- They include TCP options other than Timestamp.
- They are fragmented.
- They are subject to Equal-Cost Multi-Path (ECMP) routing, and ingress packets move from one interface to another.



CHAPTER 65

Intelligent Application Bypass

The following topics describe how to configure access control polices to use Intelligent Application Bypass (IAB)

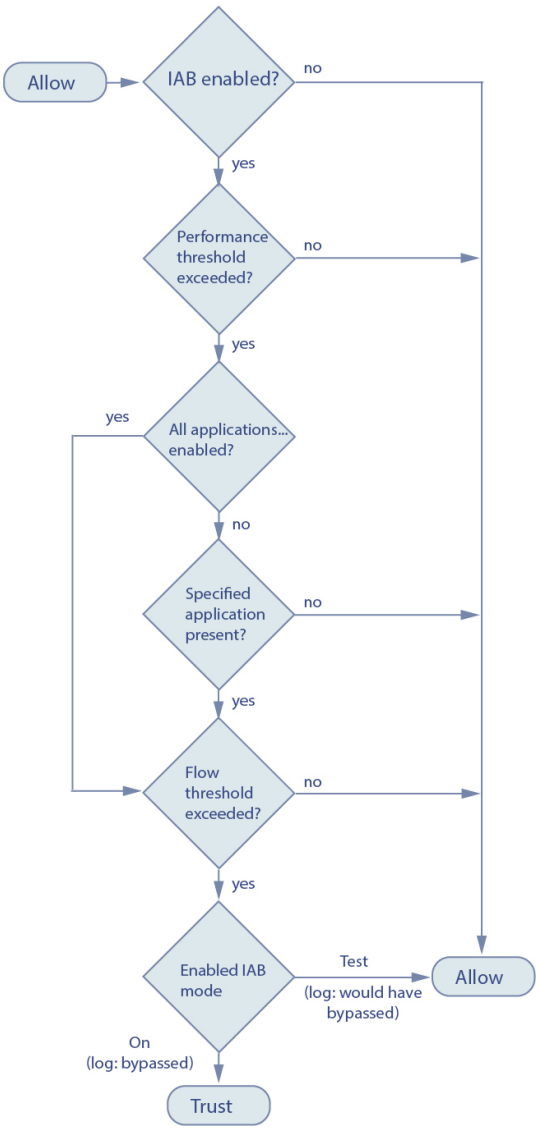
- Introduction to IAB, on page 1289
- IAB Options, on page 1290
- Requirements and Prerequisites for Intelligent Application Bypass, on page 1292
- Configuring Intelligent Application Bypass, on page 1292
- IAB Logging and Analysis, on page 1293

Introduction to IAB

IAB identifies applications that you trust to traverse your network without further inspection if performance and flow thresholds are exceeded. For example, if a nightly backup significantly impacts system performance, you can configure thresholds that, if exceeded, trust traffic generated by your backup application. Optionally, you can configure IAB so that, when an inspection performance threshold is exceeded, IAB trusts all traffic that exceeds any flow bypass threshold, regardless of the application type. This option requires Version 6.2.0.1 or a subsequent 6.2.0.x patch.

The system implements IAB on traffic allowed by access control rules or the access control policy's default action, before the traffic is subject to deep inspection. A test mode allows you to determine whether thresholds are exceeded and, if so, to identify the application flows that would have been bypassed if you had actually enabled IAB (called *bypass mode*).

The following graphic illustrates the IAB decision-making process:



IAB Options

State

Enables or disables IAB.

Performance Sample Interval

Specifies the time in seconds between IAB performance sampling scans, during which the system collects system performance metrics for comparison to IAB performance thresholds. A value of `0` disables IAB.

Bypassable Applications and Filters

This feature provides two mutually exclusive options:

Applications/Filters
Provides an editor where you can specify bypassable applications and sets of applications (filters). See Application Conditions (Application Control), on page 375.

All applications including unidentified applications
When an inspection performance threshold is exceeded, trusts all traffic that exceeds any flow bypass threshold, regardless of the application type. This option requires Version 6.2.0.1 or a subsequent 6.2.0.x patch.

Performance and Flow Thresholds

You must configure at least one inspection performance threshold and one flow bypass threshold. When a performance threshold is exceeded, the system examines flow thresholds and, if one threshold is exceeded, trusts the specified traffic.If you enable more than one of either, only one of each must be exceeded.

Inspection performance thresholds provide intrusion inspection performance limits that, if exceeded, trigger the inspection of flow thresholds. IAB does not use inspection performance thresholds set to `0`. You can configure one or more of the following inspection performance thresholds:

Drop Percentage

Average packets dropped as a percentage of total packets, when packets are dropped because of performance overloads caused by expensive intrusion rules, file policies, decompression, and so on. This does not refer to packets dropped by normal configurations such as intrusion rules. Note that specifying an integer greater than 1 activates IAB when the specified percentage of packets is dropped. When you specify `1`, any percentage from 0 through 1 activates IAB. This allows a small number of packets to activate IAB.

Processor Utilization Percentage

Average percentage of processor resources used.

Package Latency

Average packet latency in microseconds.

Flow Rate

The rate at which the system processes flows, measured as the number of flows per second. Note that this option configures IAB to measure flow *rate*, not flow *count*.

Flow bypass thresholds provide flow limits that, if exceeded, trigger IAB to trust bypassable application traffic in bypass mode or allow application traffic subject to further inspection in test mode. IAB does not use flow bypass thresholds set to `0`. You can configure one or more of the following flow bypass thresholds:

Bytes per Flow

The maximum number of kilobytes a flow can include.

Packets per Flow

The maximum number of packets a flow can include.

Flow Duration

The maximum number of seconds a flow can remain open.

Flow Velocity

The maximum transfer rate in kilobytes per second.

Requirements and Prerequisites for Intelligent Application Bypass

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Configuring Intelligent Application Bypass

Caution Not all deployments require IAB, and those that do might use it in a limited fashion. Do not enable IAB unless you have expert knowledge of your network traffic, especially application traffic, and system performance, including the causes of predictable performance issues. Before you run IAB in bypass mode, make sure that trusting the specified traffic does not expose you to risk.

Before you begin

For Classic devices, you must have the Control license.

Procedure

Step 1 In the access control policy editor, click **Advanced**, then click **Edit** () next to **Intelligent Application Bypass Settings**.

If **View** () appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 Configure IAB options:

- State—Turn IAB **Off** or **On**, or enable IAB in **Test** mode.
- Performance Sample Interval—Enter the time in seconds between IAB performance-sampling scans. If you enable IAB, even in test mode, enter a non-zero value. Entering `0` disables IAB.
- Bypassable Applications and Filters—Choose from:

- Click the number of bypassed applications and filters and specify the applications whose traffic you want to bypass; see Configuring Application Conditions and Filters, on page 376.
- Click **All applications including unidentified applications** so that, when an inspection performance threshold is exceeded, IAB trusts all traffic that exceeds any flow bypass threshold, regardless of the application type. This option requires Version 6.2.0.1 or a subsequent 6.2.0.x patch.

- Inspection Performance Thresholds—Click **Configure** and enter at least one threshold value.
- Flow Bypass Thresholds—Click **Configure** and enter at least one threshold value.

You must specify at least one inspection performance threshold and one flow bypass threshold; both must be exceeded for IAB to trust traffic. If you enter more than one threshold of each type, only one of each type must be exceeded. For detailed information, see IAB Options, on page 1290.

Step 3 Click **OK** to save IAB settings.

Step 4 Click **Save** to save the policy.

What to do next

- Because some packets must be allowed to pass before an application can be detected, you must configure your system to examine those packets.

  See Best Practices for Handling Packets That Pass Before Traffic Identification, on page 1670 and Specify a Policy to Handle Packets That Pass Before Traffic Identification, on page 1670.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

IAB Logging and Analysis

IAB forces an end-of-connection event that logs bypassed flows and flows that would have been bypassed, regardless of whether you have enabled connection logging. Connection events indicate flows that are bypassed in bypass mode or that would have been bypassed in test mode. Custom dashboard widgets and reports based on connection events can display long-term statistics for bypassed and would-have-bypassed flows.

IAB Connection Events

Action

When **Reason** includes `Intelligent App Bypass`:

`Allow` -

indicates that the applied IAB configuration was in test mode and traffic for the application specified by **Application Protocol** remains available for inspection.

`Trust` -

indicates that the applied IAB configuration was in bypass mode and traffic for the application specified by **Application Protocol** has been trusted to traverse the network without further inspection.

Reason

`Intelligent App Bypass` indicates that IAB triggered the event in bypass or test mode.

Application Protocol

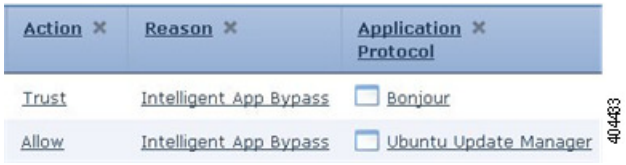
This field displays the application protocol that triggered the event.

Example

In the following truncated graphic, some fields are omitted. The graphic shows the **Action**, **Reason**, and **Application Protocol** fields for two connection events resulting from different IAB settings in two separate access control policies.

For the first event, the `Trust` action indicates that IAB was enabled in bypass mode and Bonjour protocol traffic was trusted to pass without further inspection.

For the second event, the `Allow` action indicates that IAB was enabled in test mode, so Ubuntu Update Manager traffic was subject to further inspection but would have been bypassed if IAB had been in bypass mode.



Example

In the following truncated graphic, some fields are omitted. The flow in the second event was both bypassed (**Action**: `Trust`; **Reason**: `Intelligent App Bypass`) and inspected by an intrusion rule (**Reason**: `Intrusion Monitor`). The `Intrusion Monitor` reason indicates that an intrusion rule set to **Generate Events** detected but did not block an exploit during the connection. In the example, this happened before the application was detected. After the application was detected, IAB recognized the application as bypassable and trusted the flow.



IAB Custom Dashboard Widgets

You can create a Custom Analysis dashboard widget to display long-term IAB statistics based on connection events. Specify the following when creating the widget:

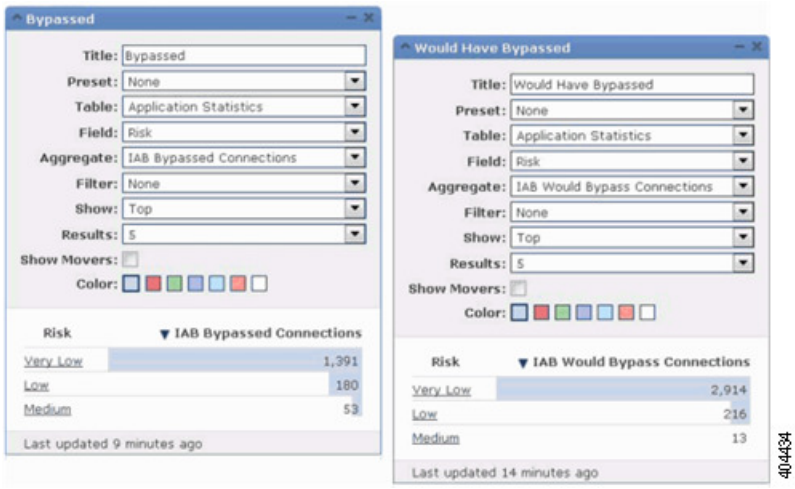
- **Preset**: `None`
- **Table**: `Application Statistics`
- **Field**: any
- **Aggregate**: either of:
  - `IAB Bypassed Connections`
  - `IAB Would Bypass Connections`

- **Filter**: any

Examples

In the following Custom Analysis dashboard widget examples:

- The *Bypassed* example shows statistics for application traffic bypassed because the applications were specified as bypassable and IAB was enabled in bypass mode in the deployed access control policy.
- The *Would Have Bypassed* example shows statistics for application traffic that would have been bypassed because the applications were specified as bypassable and IAB was enabled in test mode in the deployed access control policy. .



IAB Custom Reports

You can create a custom report to display long-term IAB statistics based on connection events. Specify the following when creating the report:

- **Table**: `Application Statistics`
- **Preset**: `None`
- **Filter**: any
- **X-Axis**: any
- **Y-Axis**: either of:
  - `IAB Bypassed Connections`
  - `IAB Would Bypass Connections`

Examples

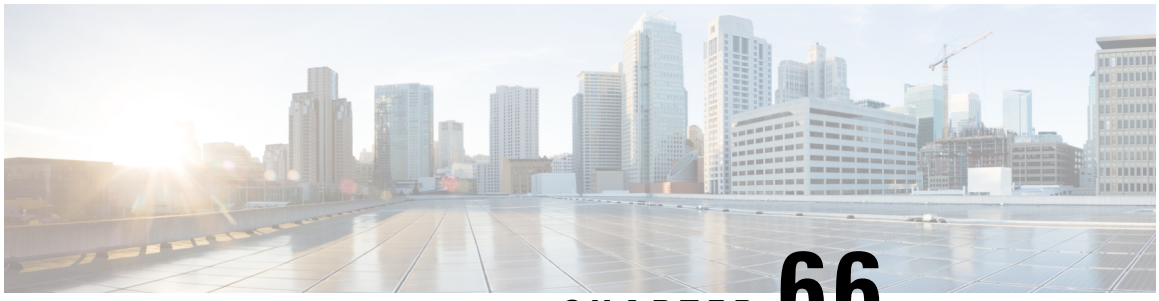
The following graphic shows two abbreviated report examples:

- The *Bypassed* example shows statistics for application traffic bypassed because the applications were specified as bypassable and IAB was enabled in bypass mode in the deployed access control policy.
- The *Would Have Bypassed* example shows statistics for application traffic that would have been bypassed because the applications were specified as bypassable and IAB was enabled in test mode in the deployed access control policy.



Related Topics

Connection and Security Intelligence Event Fields, on page 2225
The Custom Analysis Widget, on page 260
Adding Widgets to a Dashboard, on page 269
Report Templates, on page 2051



CHAPTER 66

Access Control Using Content Restriction

The following topics describe how to configure access control policies to use content restriction features:

- About Content Restriction, on page 1297
- Requirements and Prerequisites for Content Restriction, on page 1298
- Using Access Control Rules to Enforce Content Restriction, on page 1299
- Using a DNS Sinkhole to Enforce Content Restriction, on page 1301

About Content Restriction

Major search engines and content delivery services provide features that allow you to restrict search results and website content. For example, schools use content restriction features to comply with the Children's Internet Protection Act (CIPA).

When implemented by search engines and content delivery services, you can enforce content restriction features only for individual browsers or users. The Firepower System allows you to extend these features to your entire network.

The system allows you to enforce:

- *Safe Search*—Supported in many major search engines, this service filters out explicit and adult-oriented content that business, government, and education environments classify as objectionable. The system does not restrict a user's ability to access the home pages for supported search engines.
- *YouTube EDU*—This service filters YouTube content for an educational environment. It allows schools to set access for educational content while limiting access to noneducational content. YouTube EDU is a different feature than YouTube Restricted Mode, which enforces restrictions on YouTube searches as part of Google's Safe Search feature. YouTube Restricted Mode is a subfeature of Safe Search. With YouTube EDU, users access the YouTube EDU home page, rather than the standard YouTube home page.

You can use two methods to configure the system to enforce these features:

Method: Access Control Rules
: Content restriction features communicate the restricted status of a search or content query via an element in the request URI, an associated cookie, or a custom HTTP header element. You can configure access control rules to modify these elements as the system processes traffic.

Method: DNS Sinkhole
: For Google searches, you can configure the system to redirect traffic to the Google SafeSearch Virtual IP Address (VIP), which imposes filters for Safe Search (including YouTube Restricted Mode).

The table below describes the differences between these enforcement methods.

Table 91: Comparison of Content Restriction Methods

Attribute	Method: Access Control Rules	Method: DNS Sinkhole
Supported devices	Any	Firepower Threat Defense only
Search engines supported	Any tagged `safesearch supported` in the **Applications** tab of the rule editor	Google only
YouTube Restricted Mode supported	Yes	Yes
YouTube EDU supported	Yes	No
SSL policy required	Yes	No
Hosts must be using IPv4	No	Yes
Connection event logging	Yes	Yes

When determining which method to use, consider the following limitations:

- The access control rules method requires an SSL policy, which impacts performance.
- The Google SafeSearch VIP supports IPv4 traffic only. If you configure a DNS sinkhole to manage Google searches, any hosts on the affected network must be using IPv4.

The system logs different values for the **Reason** field in connection events, depending on the method:

- Access Control Rules—`Content Restriction`
- DNS Sinkhole—`DNS Block`

Requirements and Prerequisites for Content Restriction

Model Support

Any, or as indicated in the procedure.

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Using Access Control Rules to Enforce Content Restriction

Caution To avoid rule preemption, position rules governing YouTube EDU above rules governing Safe Search in both SSL and access control policies; see Content Restriction Rule Order, on page 1196.

Note When safe search or YouTube EDU is enabled in an access control rule, inline normalization is enabled automatically. For more information, see The Inline Normalization Preprocessor, on page 1762.

Before you begin

For Classic devices, you must have the Control license.

Procedure

Step 1 Create an SSL policy; see Create Basic SSL Policies, on page 1341.

Step 2 Add SSL rules for handling Safe Search and YouTube EDU traffic:

- Choose **Decrypt - Resign** as the **Action** for the rules. The system does not support any other action for content restriction handling.
- In **Applications**, add selections to the **Selected Applications and Filters** list:
  - YouTube EDU—Add the `YouTube` and `YouTube Upload` applications.
  - Safe Search—Add the `Category: search engine` filter.

For more information, see Adding an Application Condition to a TLS/SSL Rule, on page 1377.

Step 3 Set rule positions for the SSL rules you added. Click and drag, or use the right-click menu to cut and paste.

To avoid preemption, position the Safe Search rule after the YouTube EDU rule.

Step 4 Create or edit an access control policy, and associate the SSL policy with the access control policy.

For more information, see Associating Other Policies with Access Control, on page 1213.

Step 5 In the access control policy, add rules for handling Safe Search and YouTube EDU traffic:

- Choose **Allow** as the **Action** for the rules. The system does not allow any other action for content restriction handling.
- In **Applications**, click dimmed for either **Safe search** ( ) or **YouTube EDU** ( ), and set related options; see Safe Search Options for Access Control Rules, on page 1300 and YouTube EDU Options for Access Control Rules, on page 1300.

  These options are disabled, rather than dimmed, if you choose any **Action** other than **Allow** for the rule.

  You cannot enable Safe Search and YouTube EDU restrictions for the same access control rule.

- In **Applications**, refine application selections in the **Selected Applications and Filters** list.

  In most cases, enabling Safe Search or YouTube EDU populates the **Selected Applications and Filters** list with the appropriate values. The system does not automatically populate the list if a Safe Search or YouTube application is already present in the list when you enable the feature. If applications do not populate as expected, manually add them as follows:

  - YouTube EDU—Add the `YouTube` and `YouTube Upload` applications.
  - Safe Search—Add the `Category: search engine` filter.

For more information, see Configuring Application Conditions and Filters, on page 376.

Step 6 Set rule positions for the access control rules you added. Click and drag, or use the right-click menu to cut and paste.

To avoid preemption, position the Safe Search rule after the YouTube EDU rule.

Step 7 Configure the HTTP response page that the system displays when it blocks restricted content; see Choosing HTTP Response Pages, on page 1245.

Step 8 Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Safe Search Options for Access Control Rules

The Firepower System supports Safe Search filtering for specific search engines only. For a list of supported search engines, see applications tagged `safesearch supported` in the **Applications** tab of the access control rule editor. For a list of unsupported search engines, see applications tagged `safesearch unsupported`.

When enabling Safe Search for an access control rule, set the following parameters:

Enable Safe Search
Enables Safe Search filtering for traffic that matches this rule.

Unsupported Search Traffic
Specifies the action you want the system to take when it processes traffic from unsupported search engines. If you choose **Block** or **Block with Reset**, you must also configure the HTTP response page that the system displays when it blocks restricted content; see Choosing HTTP Response Pages, on page 1245.

YouTube EDU Options for Access Control Rules

When enabling YouTube EDU for an access control rule, set the following parameters:

Enable YouTube EDU
Enables YouTube EDU filtering for traffic that matches this rule.

Custom ID
Specifies the value that uniquely identifies a school or district network in the YouTube EDU initiative. YouTube provides this ID when a school or district registers for a YouTube EDU account.

Note If you check **Enable YouTube EDU**, you must enter a **Custom ID**. This ID is defined externally by YouTube. The system does not validate what you enter against the YouTube system. If you enter an invalid ID, YouTube EDU restrictions may not perform as expected.

Using a DNS Sinkhole to Enforce Content Restriction

Typically, a DNS sinkhole directs traffic away from a particular target. This procedure describes how to configure a DNS sinkhole to redirect traffic to the Google SafeSearch Virtual IP Address (VIP), which imposes content filters on Google and YouTube search results.

Because Google SafeSearch uses a single IPv4 address for the VIP, hosts must use IPv4 addressing.

Caution If your network includes proxy servers, this content restriction method is not effective unless you position your Firepower Threat Defense devices between the proxy servers and the Internet.

This procedure describes enforcing content restriction for Google searches only. To enforce content restriction for other search engines, see Using Access Control Rules to Enforce Content Restriction, on page 1299.

Before you begin

This procedure applies to Firepower Threat Defense only, and requires the Threat license.

Procedure

Step 1 Obtain a list of supported Google domains via the following URL: https://www.google.com/supported_domains.

Step 2 Create a custom DNS list on your local computer, and add the following entries:

- To enforce Google SafeSearch, add an entry for each supported Google domain.
- To enforce YouTube Restricted Mode, add a "youtube.com" entry.

The custom DNS list must be in text file (.txt) format. Each line of the text file must specify an individual domain name, stripped of any leading periods. For example, the supported domain ".google.com" must appear as "google.com".

Step 3 Upload the custom DNS list to the Firepower Management Center; see Uploading New Security Intelligence Lists to the Firepower Management Center, on page 436.

Step 4 Determine the IPv4 address for the Google SafeSearch VIP. For example, run `nslookup` on forcesafesearch.google.com.

Step 5 Create a sinkhole object for the SafeSearch VIP; see Creating Sinkhole Objects, on page 438.

Use the following values for this object:

- IPv4 Address—Enter the SafeSearch VIP address.
- IPv6 Address—Enter the IPv6 loopback address (`::1`).

- Log Connections to Sinkhole—Click Log Connections.
- Type—Choose **None**.

Step 6 Create a basic DNS policy; see Creating Basic DNS Policies, on page 1264.

Step 7 Add a DNS rule for the sinkhole; see Creating and Editing DNS Rules, on page 1266.

For this rule:

- Check the **Enabled** check box.
- Choose `Sinkhole` from the **Action** drop-down list.
- Choose the sinkhole object you created from the **Sinkhole** drop-down list.
- Add the custom DNS list you created to the **Selected Items** list on **DNS**.
- (Optional) Choose a network in **Networks** to limit content restriction to specific users. For example, if you want to limit content restriction to student users, assign students to a different subnet than faculty, and specify that subnet in this rule.

Step 8 Associate the DNS policy with an access control policy; see Associating Other Policies with Access Control, on page 1213.

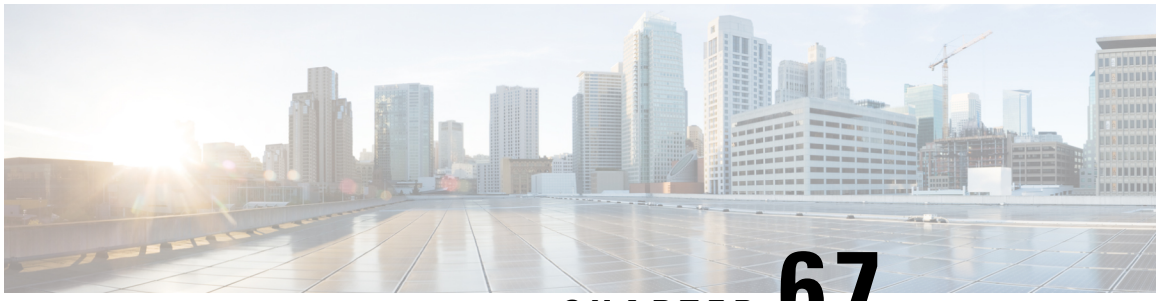
Step 9 Deploy configuration changes; see Deploy Configuration Changes, on page 346.



PART XVI

Encrypted Traffic Handling

- Understanding Traffic Decryption, on page 1305
- Start Creating SSL Policies, on page 1337
- Get Started with TLS/SSL Rules, on page 1345
- Decryption Tuning Using TLS/SSL Rules, on page 1365
- Troubleshoot TLS/SSL Rules, on page 1393



CHAPTER 67

Understanding Traffic Decryption

The following topics provide an overview of Transport Layer Security/Secure Sockets Layer (TLS/SSL) inspection, discuss the prerequisites for TLS/SSL inspection configuration, and detail deployment scenarios.

Note Because TLS and SSL are often used interchangeably, we use the expression *TLS/SSL* to indicate that either protocol is being discussed. The SSL protocol has been deprecated by the IETF in favor of the more secure TLS protocol, so you can usually interpret *TLS/SSL* as referring to TLS only.

The exception is SSL policies. Because the FMC configuration option is **Policies** > **Access Control** > **SSL**, we use the term *SSL policies* although these policies are used to define rules for TLS and SSL traffic.

For more information about SSL and TLS protocols, see a resource such as SSL vs. TLS - What's the Difference?.

- Traffic Decryption Explained, on page 1305
- TLS/SSL Handshake Processing, on page 1307
- TLS/SSL Best Practices, on page 1310
- How to Configure TLS/SSL Policies and Rules, on page 1319
- TLS/SSL Inspection Appliance Deployment Scenarios, on page 1321
- History for TLS/SSL, on page 1334

Traffic Decryption Explained

By default, the Firepower System cannot inspect traffic encrypted with the Secure Socket Layer (SSL) protocol or its successor, the Transport Layer Security (TLS) protocol. *TLS/SSL inspection* enables you to either block encrypted traffic without inspecting it, or inspect encrypted or decrypted traffic with access control. As the system handles encrypted sessions, it logs details about the traffic. The combination of inspecting encrypted traffic and analyzing encrypted session data allows greater awareness and control of the encrypted applications and traffic in your network.

TLS/SSL inspection is a policy-based feature. In the Firepower System, an access control policy is a main configuration that invokes subpolicies and other configurations, including an SSL policy. If you associate an SSL policy with access control, the system uses that SSL policy to handle encrypted sessions before it evaluates them with access control rules. If you do not configure TLS/SSL inspection, or your devices do not support it, access control rules handle all encrypted traffic.

Access control rules also handle encrypted traffic when your TLS/SSL inspection configuration allows it to pass. However, some access control rule conditions require unencrypted traffic, so encrypted traffic might match fewer rules. Also, by default, the system disables intrusion and file inspection of encrypted payloads. This helps reduce false positives and improves performance when an encrypted connection matches an access control rule that has intrusion and file inspection configured.

If the system detects a TLS/SSL handshake over a TCP connection, it determines whether it can decrypt the detected traffic. If it cannot, it applies a configured action:

- Block the encrypted traffic
- Block the encrypted traffic and reset the TCP connection
- Not decrypt the encrypted traffic

If the system cannot decrypt the traffic, it blocks the traffic without further inspection, evaluates undecrypted traffic with access control; otherwise, the system decrypts it using one of the following methods:

- Decrypt with a known private key. When an external host initiates a TLS/SSL handshake with a server on your network, the system matches the exchanged server certificate with a server certificate previously uploaded to the system. It then uses the uploaded private key to decrypt the traffic.
- Decrypt by resigning the server certificate. When a host on your network initiates a TLS/SSL handshake with an external server, the system resigns the exchanged server certificate with a previously uploaded certificate authority (CA) certificate. It then uses the uploaded private key to decrypt the traffic.

Note The Firepower System does not support mutual authentication; that is, you cannot upload a client certificate to the FMC and use it for either **Decrypt - Resign** or **Decrypt - Known Key** TLS/SSL rule actions. For more information, see Decrypt and Resign (Outgoing Traffic), on page 1313. and Known Key Decryption (Incoming Traffic), on page 1313.

Decrypted traffic is subject to the same traffic handling and analysis as originally unencrypted traffic: network, reputation, and user-based access control; intrusion detection and prevention; Cisco Advanced Malware Protection (Cisco AMP); and discovery. If the system does not block the decrypted traffic post-analysis, it re-encrypts the traffic before passing it to the destination host.

Note Set up decrypt rules *only* if your managed device handles encrypted traffic. Decryption rules require processing overhead that can impact performance.

The Firepower System does not currently support TLS version 1.3 encryption or decryption. When users visit a web site that negotiates TLS 1.3 encryption, users might see errors similar to the following in their web browser:

- **ERR_SSL_PROTOCOL_ERROR**
- **SEC_ERROR_BAD_SIGNATURE**
- **ERR_SSL_VERSION_INTERFERENCE**

For more information about how to control this behavior, contact Cisco TAC.

If you set the value of TCP maximum segment size (MSS) using FlexConfig, the observed MSS could be less than your setting. For more information, see About the TCP MSS, on page 643.

TLS/SSL Handshake Processing

In this documentation, the term *TLS/SSL handshake* represents the two-way handshake that initiates encrypted sessions in both the SSL protocol and its successor protocol, TLS.

In a passive deployment, the Firepower System observes a copy of the handshake, but does not process the actual handshake.

In an inline deployment, the Firepower System processes the TLS/SSL handshake, potentially modifying the ClientHello message and acting as a TCP proxy server for the session.

After the client establishes a TCP connection with the server (after it successfully completes the TCP three-way handshake), the managed device monitors the TCP session for any attempt to initiate an encrypted session. The TLS/SSL handshake establishes an encrypted session via the exchange of specialized packets between client and server. In the SSL and TLS protocols, these specialized packets are called *handshake messages*. The handshake messages communicate which encryption attributes both the client and server support:

- ClientHello—The client specifies multiple supported values for each encryption attribute.
- ServerHello—The server specifies a single supported value for each encryption attribute, which determines which encryption method the system uses during the secure session.

Although the data transmitted in the session is encrypted, the handshake messages are not.

After a TLS/SSL handshake completes, the managed device caches encrypted session data, which allows session resumption without requiring the full handshake. The managed device also caches server certificate data, which allows faster handshake processing in subsequent sessions.

ClientHello Message Handling

The client sends the ClientHello message to the server that acts as the packet destination if a secure connection can be established. The client sends the message to initiate the TLS/SSL handshake or in response to a Hello Request message from the destination server.

If you configure TLS/SSL decryption, when a managed device receives a ClientHello message, the system attempts to match the message to TLS/SSL rules that have the **Decrypt - Resign** action. The match relies on data from the ClientHello message and from cached server certificate data. Possible data includes:

Table 92: Data Availability for TLS/SSL Rule Conditions

TLS/SSL Rule Condition	Data Present In
Zones	ClientHello
Networks	ClientHello
VLAN Tags	ClientHello
Ports	ClientHello
Users	ClientHello

TLS/SSL Rule Condition	Data Present In
Applications	ClientHello (Server Name Indicator extension)
Categories	ClientHello (Server Name Indicator extension)
Certificate	server Certificate (potentially cached)
Distinguished Names	server Certificate (potentially cached)
Certificate Status	server Certificate (potentially cached)
Cipher Suites	ServerHello
Versions	ServerHello

If the ClientHello message does not match a **Decrypt - Resign** rule, the system does not modify the message. It then determines whether the message passes access control evaluation (which can include deep inspection). If the message passes, the system transmits it to the destination server.

If the message matches a **Decrypt - Resign** rule, the system modifies the ClientHello message as follows:

- Compression methods—Strips the `compression_methods` element, which specifies the compression methods the client supports. The Firepower System cannot decrypt compressed sessions. This modification reduces the Compressed Session type of undecryptable traffic.
- Cipher suites—Strips cipher suites from the `cipher_suites` element if the Firepower System does not support them. If the Firepower System does not support any of the specified cipher suites, the system transmits the original, unmodified element. This modification reduces the Unknown Cipher Suite and Unsupported Cipher Suite types of undecryptable traffic.
- Session identifiers—Strips any value from the `Session Identifier` element and the SessionTicket extension that does not match cached session data. If a ClientHello value matches cached data, an interrupted session can resume without the client and server performing the full TLS/SSL handshake. This modification increases the chances of session resumption and reduces the Session Not Cached type of undecryptable traffic.
- Elliptic curves—Strips elliptic curves from the Supported Elliptic Curves extension if the Firepower System does not support them. If the Firepower System does not support any of the specified elliptic curves, the managed device removes the extension and strips any related cipher suites from the `cipher_suites` element.
- ALPN extensions—The Firepower System strips the Application-Layer Protocol Negotiation (ALPN) extension, which is used in protocols the system does not support (for example, the SPDY and HTTP/2 protocols) *only* if the message matches an SSL rule associated with content restriction features. For more information about content restriction, see About Content Restriction, on page 1297.

  If content restriction is not being enforced, you can use the **system support ssl-client-hello-enabled ALPN true** command to strip the ALPN extension. Stripping the ALPN extension increases the likelihood of traffic decryption and inspection.
- Other Extensions—Strips the Extended Master Secret, Next Protocol Negotiation (NPN), and TLS Channel IDs extensions.

SSL rules with a **Decrypt - Resign** or **Decrypt - Known Key** action now natively support the Extended Master Secret (EMS) extension during ClientHello negotiation, enabling more secure communications. The EMS extension is defined by RFC 7627.



Note The system performs these ClientHello modifications by default. If your SSL policy is configured correctly, this default behavior results in more frequent decryption of traffic. To tune the default behavior for your individual network, contact Cisco TAC.



Note Use the **Cipher Suite** and **Version** rule conditions *only* in rules with either the **Block** or **Block with reset** rule actions. The use of these conditions in rules with other rule actions can interfere with the system's ClientHello processing, resulting in unpredictable performance.

After the system modifies the ClientHello message, it determines whether the message passes access control evaluation (which can include deep inspection). If the message passes, the system transmits it to the destination server.

Direct communication between the client and server is no longer possible during the TLS/SSL handshake, because after message modification the Message Authentication Codes (MACs) computed by the client and server no longer match. For all subsequent handshake messages (and for the encrypted session once established), the managed device acts as a man-in-the-middle (MITM). It creates two TLS/SSL sessions, one between client and managed device, one between managed device and server. As a result, each session contains different cryptographic session details.



Note The cipher suites that the Firepower System can decrypt are frequently updated and do not correspond directly to the cipher suites you can use in TLS/SSL rule conditions. For the current list of decryptable cipher suites, contact Cisco TAC.

Related Topics

Default Handling Options for Undecryptable Traffic, on page 1339
Encrypted Traffic Inspection with a Re-signed Certificate in an Inline Deployment, on page 1332

ServerHello and Server Certificate Message Handling

The ServerHello message is the response to a ClientHello message in a successful TLS/SSL handshake.

After a managed device processes a ClientHello message and transmits it to the destination server, the server determines whether it supports the decryption attributes the client specified in the message. If it does not support those attributes, the server sends a handshake failure alert to the client. If it supports those attributes, the server sends the ServerHello message. If the agreed-upon key exchange method uses certificates for authentication, the server Certificate message immediately follows the ServerHello message.

When the managed device receives these messsages, it attempts to match them with TLS/SSL rules. These messages contain information that was absent from either the ClientHello message or the session data cache. Specifically, the system can potentially match these messages on Distinguished Names, Certificate Status, Cipher Suites, and Versions conditions.

If the messages do not match any TLS/SSL rules, the managed device performs the default action for the SSL policy. For more information, see SSL Policy Default Actions, on page 1338.

If the messages match an SSL rule, the managed device continues as appropriate:

Action: Monitor

The TLS/SSL handshake continues to completion. The managed device tracks and logs but does not decrypt encrypted traffic.

Action: Block or Block with Reset

The managed device blocks the TLS/SSL session. If appropriate, it also resets the TCP connection.

Action: Do Not Decrypt

The TLS/SSL handshake continues to completion. The managed device does not decrypt the application data exchanged during the TLS/SSL session.

Action: Decrypt - Known Key

The managed device attempts to match the server certificate data to an Internal Certificate object previously imported into the Firepower Management Center. Because you cannot generate an Internal Certificate object, and you must possess its private key, we assume you own the server on which you're using known key decryption.

If the certificate matches a known certificate, the TLS/SSL handshake continues to completion. The managed device uses the uploaded private key to decrypt and reencrypt the application data exchanged during the TLS/SSL session.

If the server changes its certificate between the initial connection with the client and subsequent connections, you must import the new server certificate in the Firepower Management Center for future connections to be decrypted.



Note Because the Firepower System does not support the Extended Master Secret extension defined by RFC 7627, you must disable this extension on the server for which you're decrypting with a known key; otherwise, the connection is reset during decryption, resulting in connection errors.

Action: Decrypt - Resign

The managed device processes the server certificate message and re-signs the server certificate with the previously imported or generated certificate authority (CA). The TLS/SSL handshake continues to completion. The managed device then uses the uploaded private key to decrypt and reencrypt the application data exchanged during the TLS/SSL session.



Note The Firepower System does not support mutual authentication; that is, you cannot upload a client certificate to the FMC and use it for either **Decrypt - Resign** or **Decrypt - Known Key** TLS/SSL rule actions. For more information, see Decrypt and Resign (Outgoing Traffic), on page 1313. and Known Key Decryption (Incoming Traffic), on page 1313.

TLS/SSL Best Practices

This section discusses information you should keep in mind when creating your decryption policies and rules.

Note Because TLS and SSL are often used interchangeably, we use the expression *TLS/SSL* to indicate that either protocol is being discussed. The SSL protocol has been deprecated by the IETF in favor of the more secure TLS protocol, so you can usually interpret *TLS/SSL* as referring to TLS only.

The exception is SSL policies. Because the FMC configuration option is **Policies** > **Access Control** > **SSL**, we use the term *SSL policies* although these policies are used to define rules for TLS and SSL traffic.

For more information about SSL and TLS protocols, see a resource such as SSL vs. TLS - What's the Difference?.

Related Topics

The Case for Decryption, on page 1311
When to Decrypt Traffic, When Not to Decrypt, on page 1312
Other TLS/SSL Rule Actions, on page 1314
TLS/SSL Rule Components, on page 1316
TLS/SSL Rule Order Evaluation, on page 1317

The Case for Decryption

Only decrypted traffic takes advantage of the Firepower System's threat defense and policy enforcement features. Traffic that is encrypted when it passes through the Firepower System can be allowed or blocked only but it *cannot* be subjected to deep inspection or the full range of policy enforcement (such as intrusion prevention).

All encrypted connections are:

- Sent through the TLS/SSL decryption policy to determine if they should be decrypted or blocked.

  You can also configure TLS/SSL decryption rules to block encrypted traffic of types you know you do not want on your network, such as traffic that uses the nonsecure SSL protocol or traffic with an expired or invalid certificate.

- Any unblocked connections, whether or not decrypted, then go through the access control policy for a final allow or block decision.

Keep in mind that decrypting and then re-encrypting traffic adds a processing load on the device, which can reduce overall system performance.

In summary:

- Encrypted traffic can be allowed or blocked by policy; encrypted traffic *cannot* be inspected
- Decrypted traffic is subject to threat defense and policy enforcement; decrypted traffic can be allowed or blocked by policy

Related Topics

Deep Inspection Using File and Intrusion Policies, on page 1187

When to Decrypt Traffic, When Not to Decrypt

This section provides guidelines on when you should decrypt traffic and when you should allow it to pass through the firewall encrypted.

When not to decrypt traffic

You should not decrypt traffic if doing so is forbidden by:

- Law; for example, some jurisdictions forbid decrypting financial information
- Company policy; for example, your company might forbid decrypting privileged communications
- Privacy regulations
- Traffic that uses certificate pinning (also referred to as *TLS/SSL pinning*) must remain encrypted to prevent breaking the connection

If you elect to bypass decryption for certain types of traffic, no processing is done on the traffic. The encrypted traffic is first evaluated by SSL policy and then proceeds to the access control policy, where a final allow or block decision is made. Encrypted traffic can be allowed or blocked on any TLS/SSL rule condition, including, but not limited to:

- Certificate status (for example, expired or invalid certificate)
- Protocol (for example, the nonsecure SSL protocol)
- Network (security zone, IP address, VLAN tag, and so on)
- Exact URL or URL category
- Port
- User group

SSL policies provide a **Do Not Decrypt** action for this traffic; for more information, see TLS/SSL Rule Do Not Decrypt Action, on page 1360.

Note Neither TLS server identity discovery nor the decryption of TLS 1.3 traffic is supported on 8000 Series devices.

Note The related information links at the end of this topic explain how some aspects of rule evaluation work. Conditions such as URL and application filtering have limitations with respect to encrypted traffic. Make sure you understand those limitations.

When to decrypt traffic

All encrypted traffic must be decrypted to take advantage of the Firepower System's threat protection and policy enforcement features. To the extent your managed device allows traffic to be decrypted (subject to its memory and processing power), you should decrypt traffic that is not protected by law or regulation. If you must decide what traffic to decrypt, base your decision on the risk of allowing the traffic on your network.

The Firepower System provides a flexible framework for classifying traffic using rule conditions, which include URL reputation, cipher suite, protocol, and many other factors.

The Firepower System provides two methods of decryption, which are discussed in the following sections.

Related Topics

Decrypt and Resign (Outgoing Traffic), on page 1313
Known Key Decryption (Incoming Traffic), on page 1313
TLS/SSL Rule Guidelines and Limitations, on page 1345
SSL Rule Order, on page 1197
URL Conditions (URL Filtering), on page 384
Application Rule Order, on page 1197

Decrypt and Resign (Outgoing Traffic)

The **Decrypt - Resign** TLS/SSL rule action enables the Firepower System to act as a man in the middle, intercepting, decrypting, and (if the traffic is allowed) inspecting, and re-encrypting it. The **Decrypt - Resign** rule action is used with outgoing traffic; that is, the destination server is outside your protected network.

The FTD device negotiates with the client using an internal Certificate Authority (CA) object specified in the rule and builds an SSL tunnel between the client and the FTD device. At the same time, the device connects to the destination web site and creates an SSL tunnel between the server and the FTD device.

Thus, the client sees the CA certificate configured for the SSL decryption rule instead of the certificate from the destination server. The client must trust the certificate to complete the connection. The FTD device then performs decryption/re-encryption in both directions for traffic between the client and the destination server.

Prerequisite

To use the **Decrypt - Resign** rule action, you must create an internal CA object using a CA file and paired private key file. You can generate a CA and private key in the Firepower System if you don't already have them.

Note The Firepower System does not support mutual authentication; that is, you cannot upload a client certificate to the FMC and use it for either **Decrypt - Resign** or **Decrypt - Known Key** TLS/SSL rule actions. For more information, see Decrypt and Resign (Outgoing Traffic), on page 1313. and Known Key Decryption (Incoming Traffic), on page 1313.

Related Topics

TLS/SSL Rule Decrypt Actions, on page 1361
External Certificate Objects, on page 454

Known Key Decryption (Incoming Traffic)

The **Decrypt - Known Key** TLS/SSL rule action uses a server's private key to decrypt traffic. The **Decrypt - Known Key** rule action is used with incoming traffic; that is, the destination server is inside your protected network.

The main purpose of decrypting with a known key is to protect your servers from external attacks.

Prerequisite

To use the **Decrypt - Known Key** rule action, you must create an internal certificate object using the server's certificate file and paired private key file.

Note The Firepower System does not support mutual authentication; that is, you cannot upload a client certificate to the FMC and use it for either **Decrypt - Resign** or **Decrypt - Known Key** TLS/SSL rule actions. For more information, see Decrypt and Resign (Outgoing Traffic), on page 1313. and Known Key Decryption (Incoming Traffic), on page 1313.

Related Topics

TLS/SSL Rule Decrypt Actions, on page 1361
Internal Certificate Objects, on page 455

Other TLS/SSL Rule Actions

The following sections discuss other TLS/SSL rule actions.

Related Topics

TLS/SSL Rule Blocking Actions, on page 1360
TLS/SSL Rule Monitor Action, on page 1359

TLS/SSL Rule Examples

The following sections provide examples of setting up recommended TLS/SSL rules.

Related Topics

Block Nonsecure Protocols, on page 1314

Block Nonsecure Protocols

This example shows how to block TLS and SSL protocols on your network that are no longer considered secure, such as TLS 1.0, TLS 1.1, and SSLv3.

You should exclude nonsecure protocols from your network because they are all exploitable. In this example:

- You can block some protocols using **Version** page on the SSL rule.
- Because the Firepower System considers SSLv2 as undecryptable, you can block it using the **Undecryptable Actions** on the SSL policy.

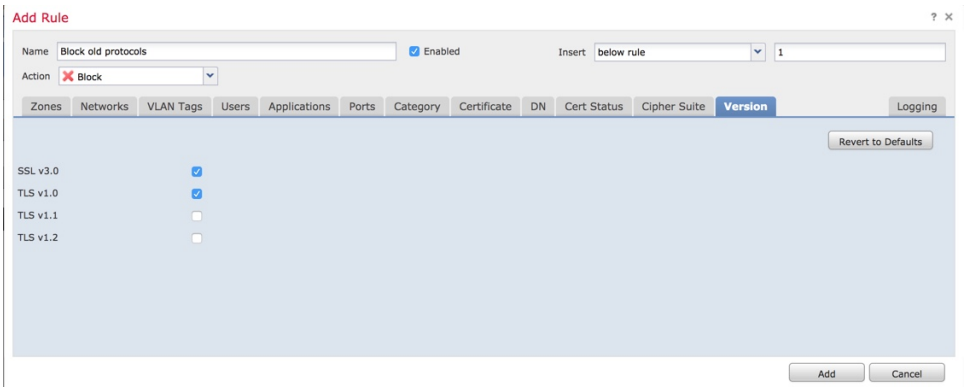
Procedure

Step 1 Log in to the Firepower Management System if you have not already done so.
Step 2 Click **Policies** > **Access Control** > **SSL**.
Step 3 Add or edit an SSL policy.
Step 4 Click **Add Rule**.
Step 5 In the **Name** field, enter a name for the rule.
Step 6 From the **Action** list, click **Block** or **Block with reset**.

Step 7 Click **Version** page.

Step 8 Check the check boxes for protocols that are no longer secure, such as **SSL v3.0**, **TLS 1.0**, and **TLS 1.1**. Clear the check boxes for any protocols that are still considered secure.

The following figure shows an example.

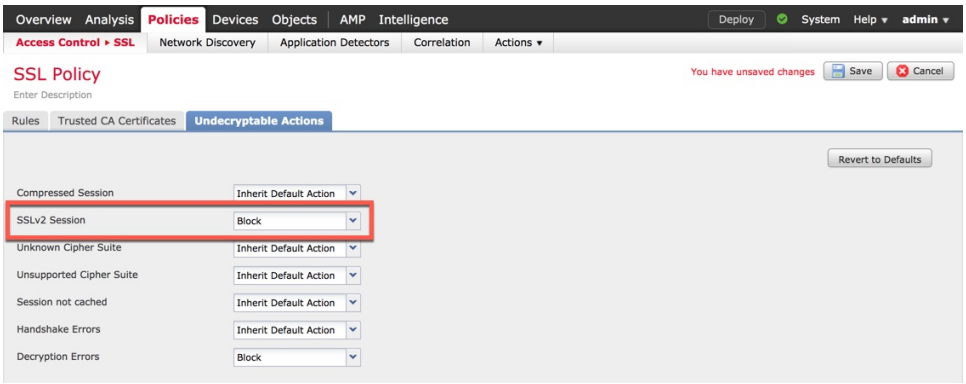


Step 9 Choose other rule conditions as needed.

Step 10 Save the rule.

Step 11 On the SSL policy page, click **Undecryptable Actions**.

Step 12 From the **SSLv2 Session** list, click **Block** or **Block with reset**. The following figure shows an example.



Step 13 Click **Save**.

Step 14 Because this is a specific rule, order it earlier in your policy than more general rules such as application-matching rules.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

TLS/SSL Rule Conditions, on page 1357

TLS/SSL Rule Components

Each TLS/SSL rule has the following components.

State

By default, rules are enabled. If you disable a rule, the system does not use it to evaluate network traffic, and stops generating warnings and errors for that rule.

Position

Rules in an SSL policy are numbered, starting at 1. The system matches traffic to rules in top-down order by ascending rule number. With the exception of Monitor rules, the first rule that traffic matches is the rule that handles that traffic.

Conditions

Conditions specify the specific traffic the rule handles. Conditions can match traffic by security zone, network or geographical location, VLAN, port, application, requested URL, user, certificate, certificate subject or issuer, certificate status, cipher suite, or encryption protocol version. The use of conditions can depend on target device licenses.

Action

A rule's action determines how the system handles matching traffic. You can monitor, allow, block, or decrypt encrypted matching traffic. Decrypted and allowed encrypted traffic is subject to further inspection. Note that the system does **not** perform inspection on blocked encrypted traffic.

Logging

A rule's logging settings govern the records the system keeps of the traffic it handles. You can keep a record of traffic that matches a rule. You can log a connection when the system blocks an encrypted session or allows it to pass without decryption, according to the settings in an SSL policy. You can also force the system to log connections that it decrypts for further evaluation by access control rules, regardless of how the system later handles or inspects the traffic. You can log connections to the Firepower Management Center database, as well as to the system log (syslog) or to an SNMP trap server.

For more information about logging, see Best Practices for Connection Logging, on page 2216.



Tip Properly creating and ordering TLS/SSL rules is a complex task. If you do not plan your policy carefully, rules can preempt other rules, require additional licenses, or contain invalid configurations. To help ensure that the system handles traffic as you expect, the SSL policy interface has a robust warning and error feedback system for rules.

Related Topics

Interface Conditions, on page 366
Network Conditions, on page 368
VLAN Conditions, on page 372
Port and ICMP Code Conditions, on page 373
Application Conditions (Application Control), on page 375

URL Conditions (URL Filtering), on page 384
User, Realm, and ISE Attribute Conditions (User Control), on page 384
Best Practices for Access Control Rules, on page 1194

TLS/SSL Rule Order Evaluation

When you create the TLS/SSL rule in an SSL policy, you specify its position using the **Insert** list in the rule editor. TLS/SSL rules in an SSL policy are numbered, starting at 1. The system matches traffic to TLS/SSL rules in top-down order by ascending rule number.

In most cases, the system handles network traffic according to the *first* TLS/SSL rule where *all* the rule's conditions match the traffic. Except in the case of Monitor rules (which log traffic but do not affect traffic flow), the system does *not* continue to evaluate traffic against additional, lower-priority rules after that traffic matches a rule. Conditions can be simple or complex; you can control traffic by security zone, network or geographical location, VLAN, port, application, requested URL, user, certificate, certificate distinguished name, certificate status, cipher suite, or encryption protocol version.

Each rule also has an *action*, which determines whether you monitor, block, or inspect matching encrypted or decrypted traffic with access control. Note that the system does *not* further inspect encrypted traffic it blocks. It does subject encrypted and undecryptable traffic to access control. However, access control rule conditions require unencrypted traffic, so encrypted traffic matches fewer rules.

Rules that use *specific* conditions (such as network and IP addresses) should be ordered *before* rules that use general conditions (such as applications). If you're familiar with the Open Systems Interconnect (OSI) model, use similar numbering in concept. Rules with conditions for layers 1, 2, and 3 (physical, data link, and network) should be ordered first in your rules. Conditions for layers 5, 6, and 7 (session, presentation, and application) should be ordered later in your rules. For more information about the OSI model, see this Wikipedia article.



Tip Proper TLS/SSL rule order reduces the resources required to process network traffic, and prevents rule preemption. Although the rules you create are unique to every organization and deployment, there are a few general guidelines to follow when ordering rules that can optimize performance while still addressing your needs.

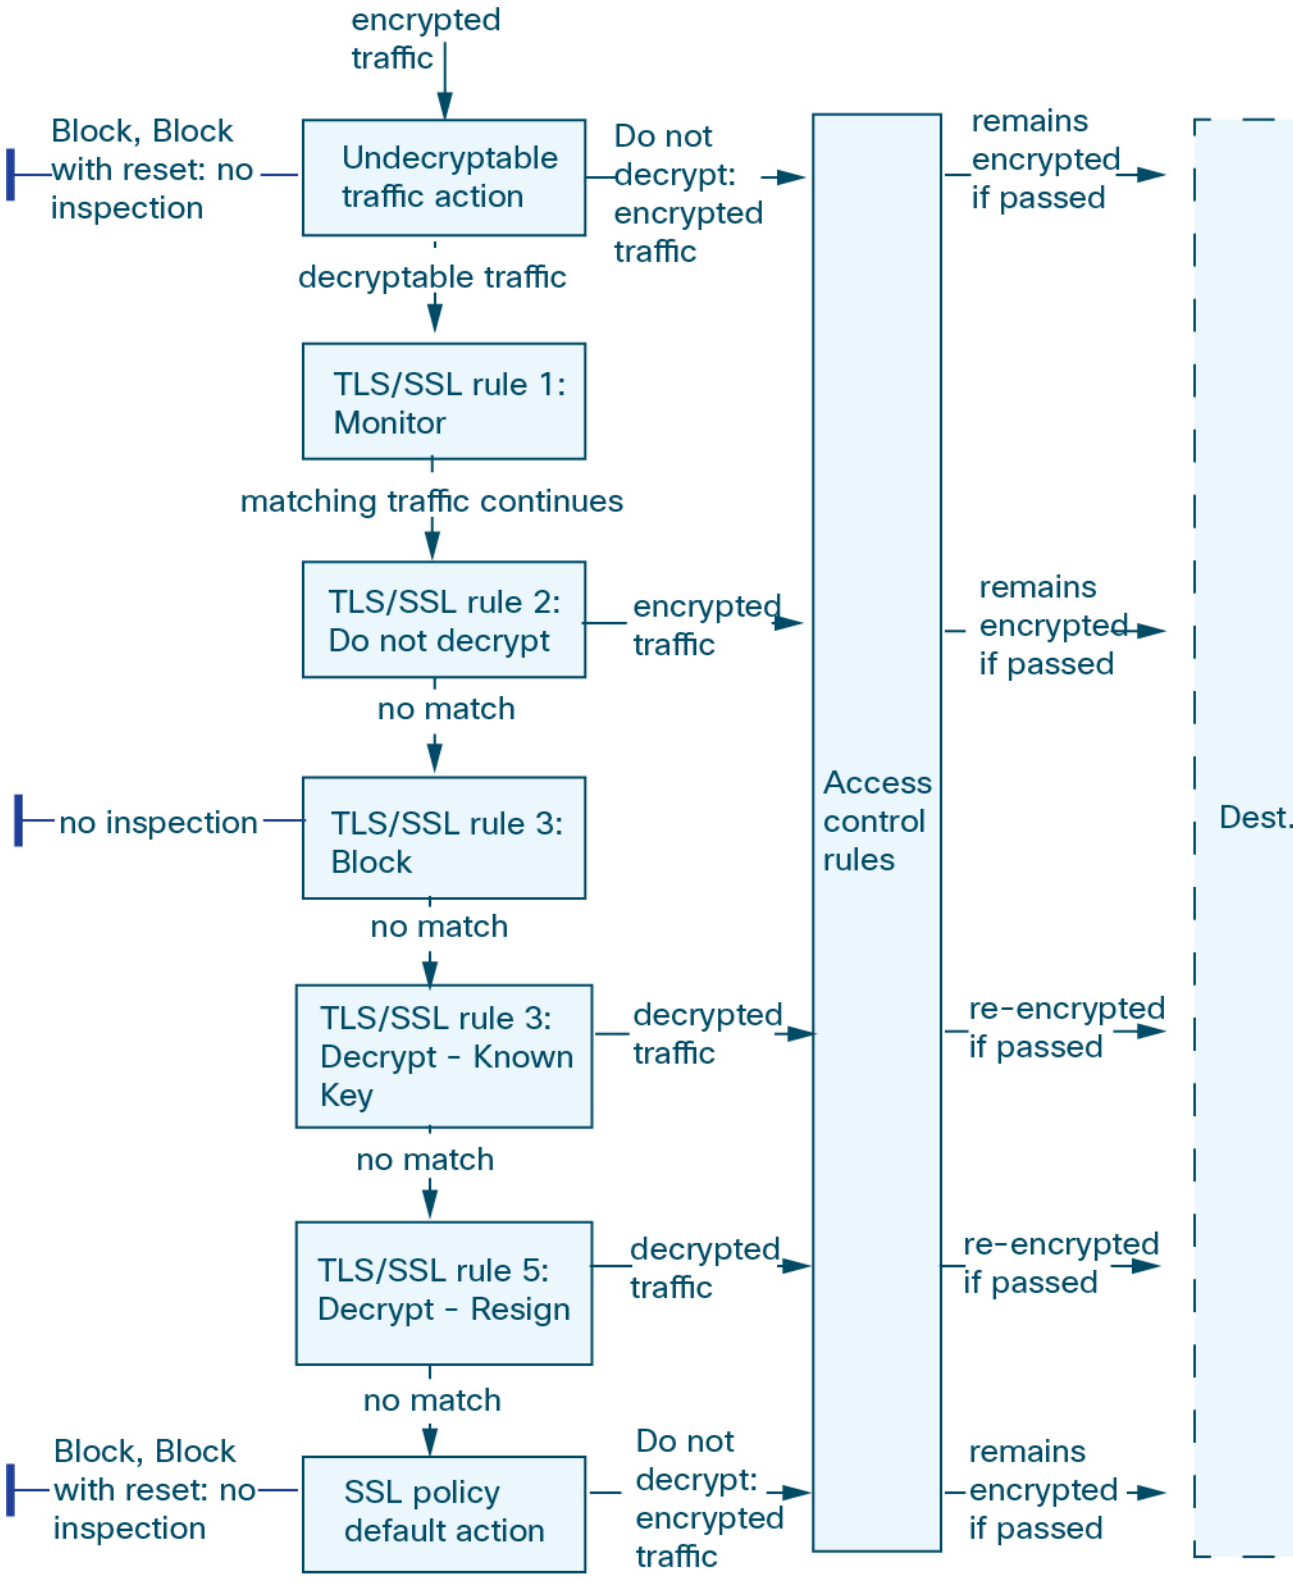
In addition to ordering rules by number, you can group rules by category. By default the system provides three categories: Administrator, Standard, and Root. You can add custom categories, but you cannot delete the system-provided categories or change their order.

Related Topics

Best Practices for Access Control Rules, on page 1194
Default Handling Options for Undecryptable Traffic, on page 1339
SSL Rule Order, on page 1197

Multi-Rule Example

The following scenario summarizes the ways that SSL rules handle traffic in an inline deployment.


encrypted traffic
Block, Block with reset: no inspection
Undecryptable traffic action
Do not decrypt: encrypted traffic
remains encrypted if passed
decryptable traffic
TLS/SSL rule 1: Monitor
matching traffic continues
TLS/SSL rule 2: Do not decrypt
encrypted traffic
remains encrypted if passed
no match
no inspection
TLS/SSL rule 3: Block
Access control rules
Dest.
no match
TLS/SSL rule 3: Decrypt - Known Key
decrypted traffic
re-encrypted if passed
no match
TLS/SSL rule 5: Decrypt - Resign
decrypted traffic
re-encrypted if passed
no match
Block, Block with reset: no inspection
SSL policy default action
Do not decrypt: encrypted traffic
remains encrypted if passed

In this scenario, traffic is evaluated as follows:

- **Undecryptable Traffic Action** evaluates encrypted traffic first. For traffic the system cannot decrypt, the system either blocks it without further inspection or passes it for access control inspection. Encrypted traffic that does not match continues to the next rule.
- **TLS/SSL Rule 1: Monitor** evaluates encrypted traffic next. Monitor rules track and log encrypted traffic but do not affect traffic flow. The system continues to match traffic against additional rules to determine whether to permit or deny it.
- **TLS/SSL Rule 2: Do Not Decrypt** evaluates encrypted traffic third. Matching traffic is not decrypted; the system inspects this traffic with access control, but not file or intrusion inspection. Traffic that does not match continues to the next rule.
- **TLS/SSL Rule 3: Block** evaluates encrypted traffic fourth. Matching traffic is blocked without further inspection. Traffic that does not match continues to the next rule.
- **TLS/SSL Rule 4: Decrypt - Known Key** evaluates encrypted traffic fifth. Matching traffic incoming to your network is decrypted using a private key you upload. The decrypted traffic is then evaluated against access control rules. Access control rules handle decrypted and unencrypted traffic identically. The system can block traffic as a result of this additional inspection. All remaining traffic is reencrypted before being allowed to the destination. Traffic that does not match the SSL rule continues to the next rule.
- **TLS/SSL Rule 5: Decrypt - Resign** is the final rule. If traffic matches this rule, the system re-signs the server certificate with an uploaded CA certificate, then acts as a man-in-the-middle to decrypt traffic. The decrypted traffic is then evaluated against access control rules. Access control rules treat decrypted and unencrypted traffic identically. The system can block traffic as a result of this additional inspection. All remaining traffic is reencrypted before being allowed to the destination. Traffic that does not match the SSL rule continues to the next rule.
- **SSL Policy Default Action** handles all traffic that does not match any of the TLS/SSL rules. The default action either blocks encrypted traffic without further inspection or does not decrypt it, passing it for access control inspection.

How to Configure TLS/SSL Policies and Rules

This topic provides a high-level overview of tasks you must complete to configure SSL policies and TLS/SSL rules in those policies to block, monitor, or allow TLS/SSL traffic on your network.

You must be an Admin, Access Admin, or Network Admin to perform this task. You can configure SSL policies on any device type except NGIPSv.

Procedure

	Command or Action	Purpose
Step 1	Create an SSL policy.	An SSL policy is a container for one or more rules. To use an SSL policy and its rules for access control, you must later associate the SSL policy with an access control policy. For more information, see See Create Basic SSL Policies, on page 1341.

	Command or Action	Purpose
Step 2	Set a default action for your SSL policy.	The default action is taken when traffic matches no rules defined by the SSL policy. See SSL Policy Default Actions, on page 1338.
Step 3	Specify how undecryptable traffic should be handled.	Traffic can be undecryptable for a number of reasons, including unsecure protocols, uses and unknown cipher suite, or in the event of errors with the handshake or decryption. See Default Handling Options for Undecryptable Traffic, on page 1339.
Step 4	For **Decrypt - Known Key** (to decrypt inbound traffic to a server in your network) TLS/SSL rules, create an internal certificate object.	The internal certificate object uses your server's certificate and private key. See Internal Certificate Objects, on page 455.
Step 5	For **Decrypt - Resign** (to decrypt outbound traffic to a server outside of your network) TLS/SSL rules, create an internal certificate authority (CA) object.	The internal CA object uses a CA and private key. See Internal Certificate Authority Objects, on page 448.
Step 6	Create your TLS/SSL rules:	• **Block, Block with reset, Interactive block**: Configuring TLS/SSL Rule Actions, on page 1361. • **Do Not Decrypt**, see Configuring TLS/SSL Rule Actions, on page 1361. • **Decrypt - Resign**, see Configuring a Decrypt - Resign Action, on page 1362. • **Decrypt - Known Key**, see Configuring a Decrypt - Known Key Action, on page 1362. • **Monitor**, see Configuring TLS/SSL Rule Actions, on page 1361.
Step 7	Associate the SSL policy with an access control policy.	Unless you associate your SSL policy with an access control policy, it has no effect. After you do this, you can choose to allow or block traffic that matches the access control rule and take other actions. See Associating Other Policies with Access Control, on page 1213.
Step 8	Configure your access control rules to allow or block decrypted traffic.	See Access Control Policy Components, on page 1201.
Step 9	Deploy the access control policy to managed devices.	Before your policy can take effect, it must be deployed to managed devices. See Deploy Configuration Changes, on page 346.

TLS/SSL Inspection Appliance Deployment Scenarios

This section presents several scenarios in which the Life Insurance Example, Inc. life insurance company (LifeIns) uses SSL inspection on encrypted traffic to help audit their processes. Based on their business processes, LifeIns plans to deploy:

- one FTD device in a passive deployment for the Customer Service department
- one FTD device in an inline deployment for the Underwriting Department
- one Firepower Management Center to manage both devices

Customer Service Business Processes

LifeIns created a customer-facing website for their customers. LifeIns receives encrypted questions and requests regarding policies from prospective customers through their website and through e-mail. LifeIns's Customer Service department processes them and returns the requested information within 24 hours. Customer Service wants to expand its incoming contact metrics collection. LifeIns has an established internal audit review for Customer Service.

LifeIns also receives encrypted applications online. The Customer Service department processes the applications within 24 hours before sending the case file to the Underwriting department. Customer Service filters out any obvious false applications sent through the online form, which consumes a fair portion of their time.

Underwriting Business Processes

LifeIns's underwriters submit encrypted medical information requests online to the Medical Repository Example, LLC medical data repository (MedRepo). MedRepo reviews the requests and transmits the encrypted records to LifeIns within 72 hours. The underwriters subsequently underwrite an application and submit policy and rate decisions. Underwriting wants to expand its metrics collection.

Lately, an unknown source has been sending spoofed responses to LifeIns. Though LifeIns's underwriters receive training on proper Internet use, LifeIns's IT department first wants to analyze all encrypted traffic that takes the form of medical responses, then wants to block all spoof attempts.

LifeIns places junior underwriters on six-month training periods. Lately, these underwriters have been incorrectly submitting encrypted medical regulation requests to MedRepo's customer service department. MedRepo has submitted multiple complaints to LifeIns in response. LifeIns plans on extending their new underwriter training period to also audit underwriter requests to MedRepo.

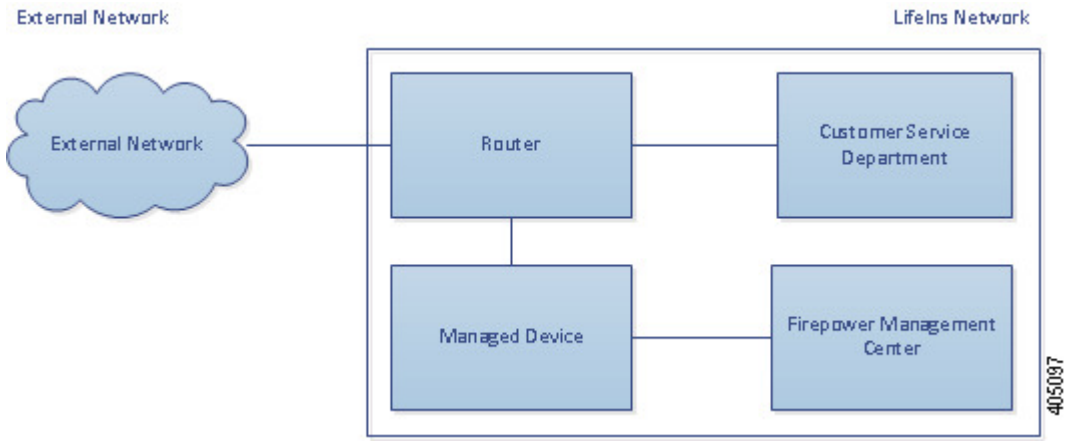
Traffic Decryption in a Passive Deployment

LifeIns's business requirements state that Customer Service must:

- process all requests and applications within 24 hours
- improve its incoming contact metrics collection process
- identify and discard incoming false applications

Customer Service does not require additional audit review.

LifeIns plans to passively deploy a Customer Service managed device.



Traffic from an external network goes to LifeIns's router. The router routes traffic to the Customer Service department, and mirrors a copy of the traffic to the managed device for inspection.

On the managing Firepower Management Center, a user in the access control configures TLS/SSL inspection to:

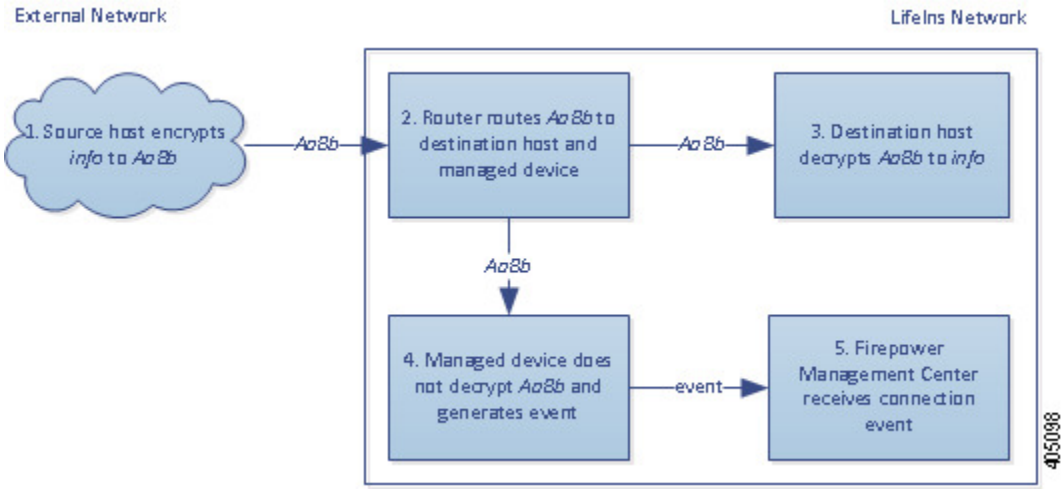
- log all encrypted traffic sent to the Customer Service department
- decrypt encrypted traffic sent using the online application form to Customer Service
- not decrypt all other encrypted traffic sent to Customer service, including traffic sent using the online request form

The user also configures access control to inspect the decrypted application form traffic for fake application data and log when fake data is detected.

In the following scenarios, the user submits an online form to Customer Service. The user's browser establishes a TCP connection with the server, then initiates a TLS/SSL handshake. The managed device receives a copy of this traffic. The client and server complete the TLS/SSL handshake, establishing the encrypted session. Based on handshake and connection details, the system logs the connection and acts upon the copy of the encrypted traffic.

Encrypted Traffic Monitoring in a Passive Deployment

For all TLS/SSL-encrypted traffic sent to Customer Service, the managed device logs the connection.



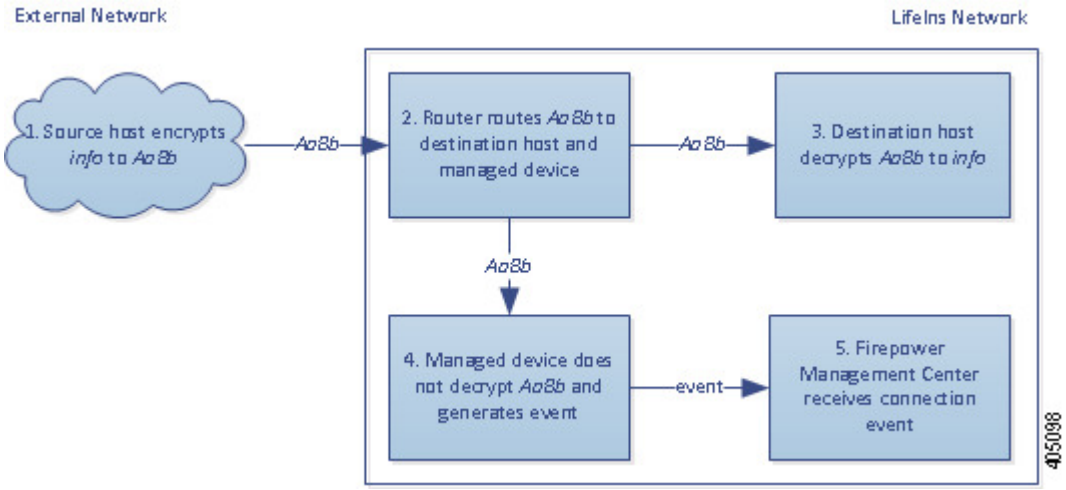
The following steps occur:

1. The user submits the plain text request (`info`). The client encrypts this (`AaBb`) and sends the encrypted traffic to Customer Service.
2. LifeIns's router receives the encrypted traffic and routes it to the Customer Service department server. It also mirrors a copy to the managed device.
3. The Customer Service department server receives the encrypted information request (`AaBb`) and decrypts it to plain text (`info`).
4. The managed device does not decrypt the traffic.

   The access control policy continues to process the encrypted traffic and allows it. The device generates a connection event after the session ends.
5. The Firepower Management Center receives the connection event.

Undecrypted Encrypted Traffic in a Passive Deployment

For all TLS/SSL-encrypted traffic that contains requests about policies, the managed device allows the traffic without decrypting it and logs the connection.



The following steps occur:

1. The user submits the plain text request (`info`). The client encrypts this (`AaBb`) and sends the encrypted traffic to Customer Service.
2. LifeIns's router receives the encrypted traffic and routes it to the Customer Service department server. It also mirrors a copy to the managed device.
3. The Customer Service department server receives the encrypted information request (`AaBb`) and decrypts it to plain text (`info`).
4. The managed device does not decrypt the traffic.

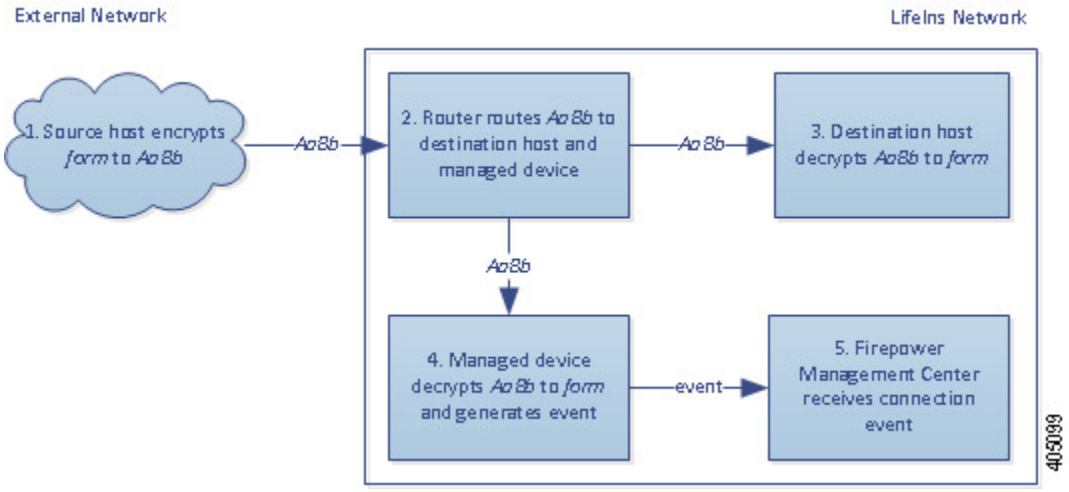
   The access control policy continues to process the encrypted traffic and allows it. The device generates a connection event after the session ends.
5. The Firepower Management Center receives the connection event.

Encrypted Traffic Inspection with a Private Key in a Passive Deployment

For all TLS/SSL-encrypted traffic that contains application form data, the system decrypts the traffic and logs the connection.

Note In a passive deployment, if traffic is encrypted with either the DHE or ECDHE cipher suite, you cannot decrypt it with a known private key.

For traffic with legitimate application form information, the system logs the connection.

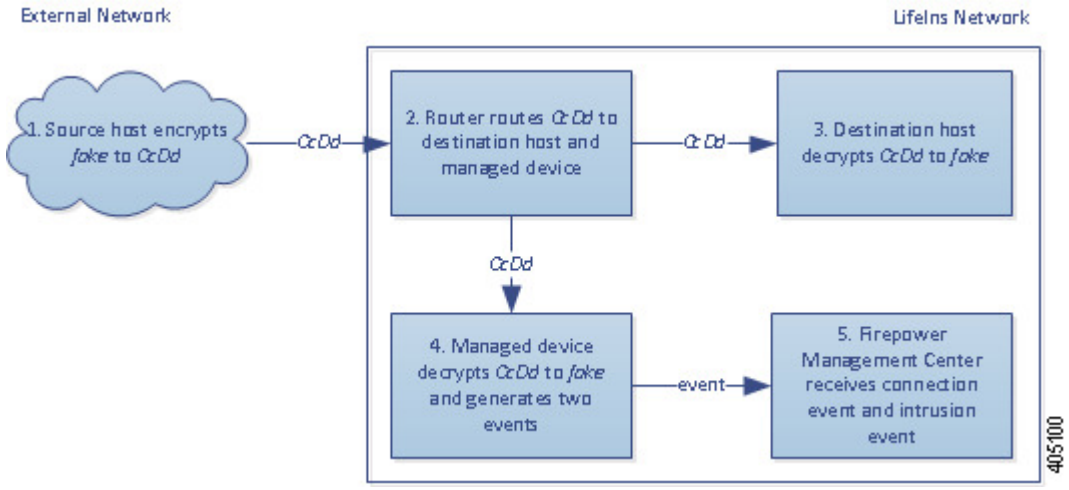


The following steps occur:

1. The user submits the plain text request (`form`). The client encrypts this (`AaBb`) and sends the encrypted traffic to Customer Service.
2. LifeIns's router receives the encrypted traffic and routes it to the Customer Service department server. It also mirrors a copy to the managed device.
3. The Customer Service department server receives the encrypted information request (`AaBb`) and decrypts it to plain text (`form`).
4. The managed device uses the session key obtained with the uploaded known private key to decrypt the encrypted traffic to plain text (`form`).

   The access control policy continues to process the decrypted traffic and does not find fake application information. The device generates a connection event after the session ends.
5. The Firepower Management Center receives a connection event with information about the encrypted and decrypted traffic.

In contrast, if the decrypted traffic contains fake application data, the system logs the connection and the fake data.



The following steps occur:

1. The user submits the plain text request (`fake`). The client encrypts this (`CcDd`) and sends the encrypted traffic to Customer Service.
2. LifeIns's router receives the encrypted traffic and routes it to the Customer Service department server. It also mirrors a copy to the managed device.
3. The Customer Service department server receives the encrypted information request (`CcDd`) and decrypts it to plain text (`fake`).
4. The managed device uses the session key obtained with the uploaded known private key to decrypt the encrypted traffic to plain text (`fake`).

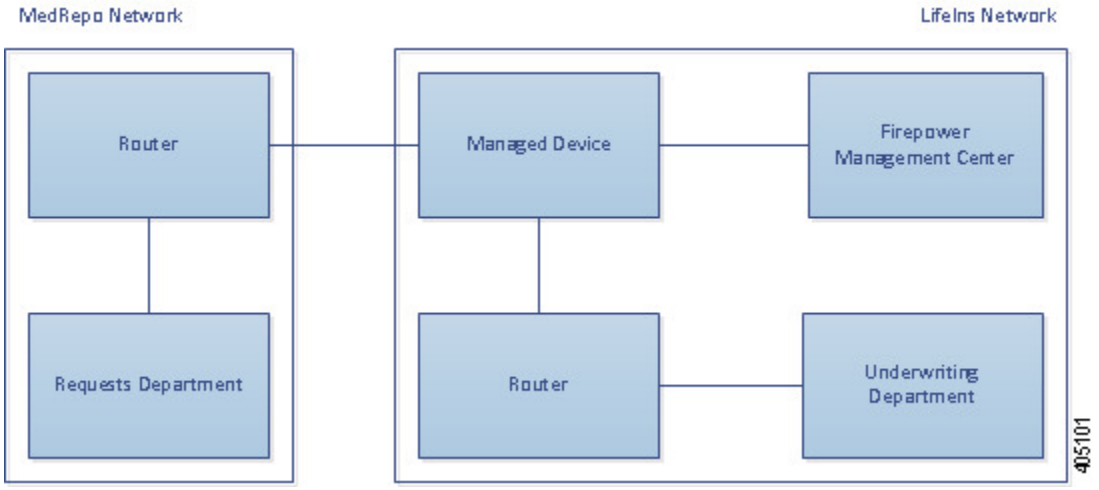
   The access control policy continues to process the decrypted traffic and finds fake application information. The device generates an intrusion event. After the session ends, it generates a connection event.
5. The Firepower Management Center receives a connection event with information about the encrypted and decrypted traffic, and an intrusion event for the fake application data.

Traffic Decryption in an Inline Deployment

LifeIns's business requirements state that Underwriting must:

- audit new and junior underwriters, verifying that their information requests to MedRepo comply with all applicable regulations
- improve its underwriting metrics collection process
- examine all requests that appear to come from MedRepo, then drop any spoofing attempts
- drop all improper regulatory requests to MedRepo's Customer Service department from the Underwriting department
- not audit senior underwriters

LifeIns plans to deploy a device in an inline deployment for the Underwriting department.



Traffic from MedRepo's network goes to MedRepo's router. It routes traffic to LifeIns's network. The managed device receives the traffic, passes allowed traffic to LifeIns's router, and sends events to the managing Firepower Management Center. LifeIns's router routes traffic to the destination host.

On the managing Firepower Management Center, a user in the Access Control and SSL Editor custom role configures an SSL access control rule to:

- log all encrypted traffic sent to the Underwriting department
- block all encrypted traffic incorrectly sent from LifeIns's underwriting department to MedRepo's customer service department
- decrypt all encrypted traffic sent from MedRepo to LifeIns's underwriting department, and from LifeIns's junior underwriters to MedRepo's requests department
- not decrypt encrypted traffic sent from the senior underwriters

The user also configures access control to inspect decrypted traffic with a custom intrusion policy and:

- block decrypted traffic if it contains a spoof attempt, and log the spoof attempt
- block decrypted traffic that contains information not compliant with regulations, and log the improper information
- allow all other encrypted and decrypted traffic

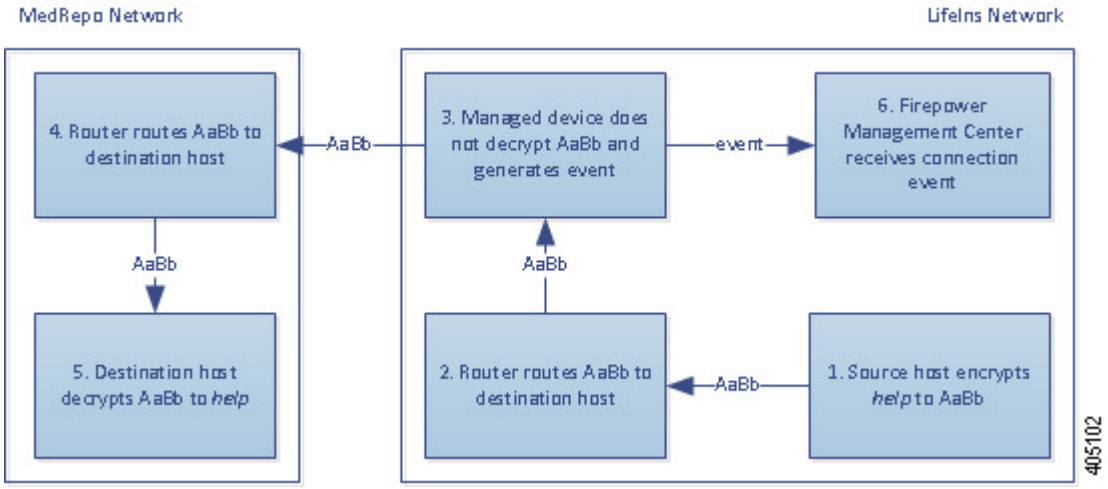
The system reencrypts allowed decrypted traffic before sending it to the destination host.

You can also cause the system to decrypt and resign the traffic using a TLS/SSL control rule with the action **Decrypt - Resign**. If traffic matches the TLS/SSL rule, after the system modifies the ClientHello message, it determines whether the message passes access control evaluation (which can include deep inspection). If the message passes, the system transmits it to the destination server. For more information, see ClientHello Message Handling, on page 1307

In the following scenarios, the user submits information online to a remote server. The user's browser establishes a TCP connection with the server, then initiates an SSL handshake. The managed device receives this traffic; based on handshake and connection details, the system logs the connection and acts on the traffic. If the system blocks the traffic, it also closes the TCP connection. Otherwise, the client and server complete the SSL handshake, establishing the encrypted session.

Encrypted Traffic Monitoring in an Inline Deployment

For all SSL-encrypted traffic sent to and from the Underwriting department, the system logs the connection.



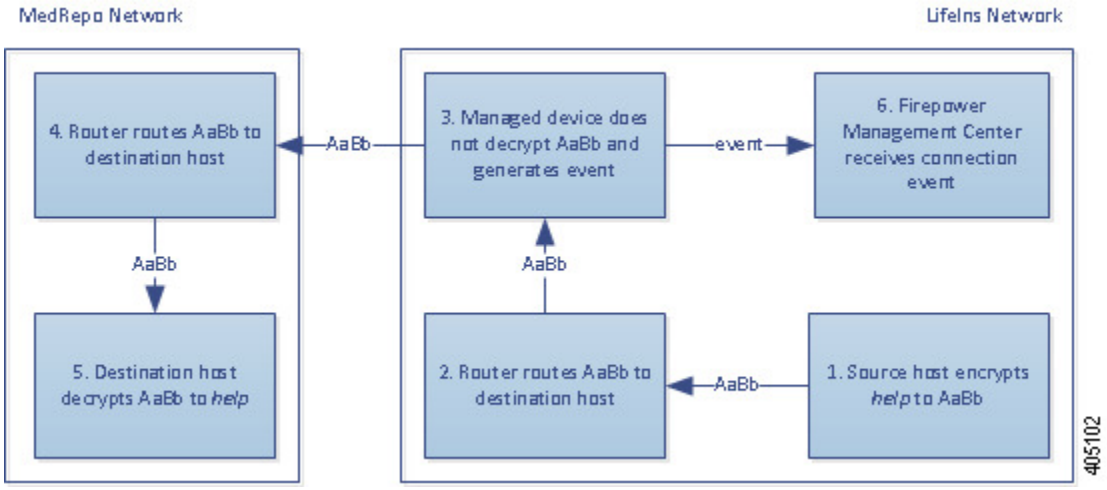
The following steps occur:

1. The user submits the plain text request (`help`). The client encrypts this (`AaBb`) and sends the encrypted traffic to MedRepo's Requests department server.
2. LifeIns's router receives the encrypted traffic and routes it to the Requests department server.
3. The managed device does not decrypt the traffic.

   The access control policy continues to process the encrypted traffic and allows it, then generates a connection event after the session ends.
4. The external router receives the traffic and routes it to the Requests department server.
5. The Underwriting department server receives the encrypted information request (`AaBb`) and decrypts it to plain text (`help`).
6. The Firepower Management Center receives the connection event.

Undecrypted Encrypted Traffic in an Inline Deployment

For all TLS/SSL-encrypted traffic originating from the senior underwriters, the managed device allows the traffic without decrypting it and logs the connection.



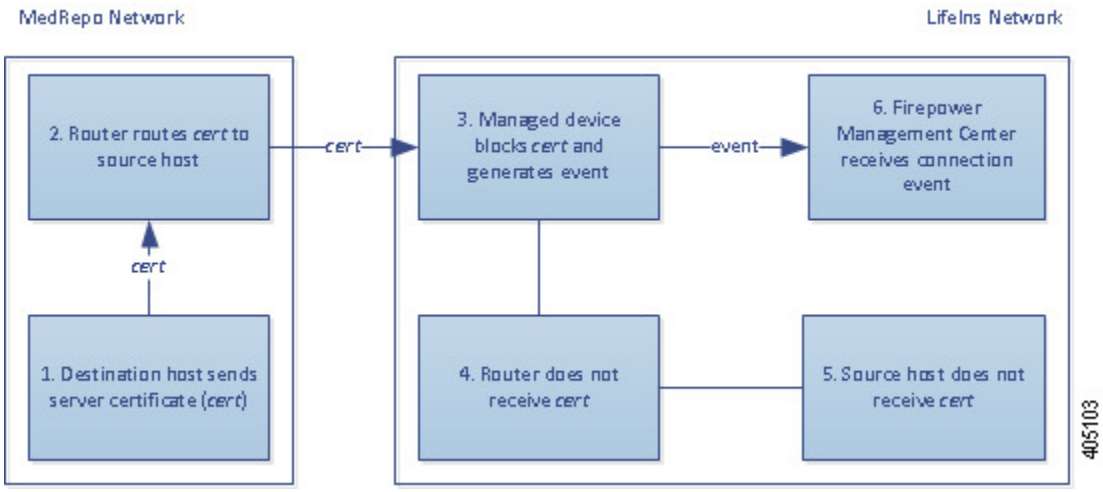
The following steps occur:

1. The user submits the plain text request (`help`). The client encrypts this (`AaBb`) and sends the encrypted traffic to MedRepo's Requests department server.
2. LifeIns's router receives the encrypted traffic and routes it to the Requests department server.
3. The managed device does not decrypt this traffic.

   The access control policy continues to process the encrypted traffic and allows it, then generates a connection event after the session ends.
4. The external router receives the traffic and routes it to the Requests department server.
5. The Requests department server receives the encrypted information request (`AaBb`) and decrypts it to plain text (`help`).
6. The Firepower Management Center receives the connection event.

Encrypted Traffic Blocking in an Inline Deployment

For all SMTPS email traffic improperly sent from LifeIns's underwriting department to MedRepo's Customer Service department, the system blocks the traffic during the SSL handshake without further inspection and logs the connection.

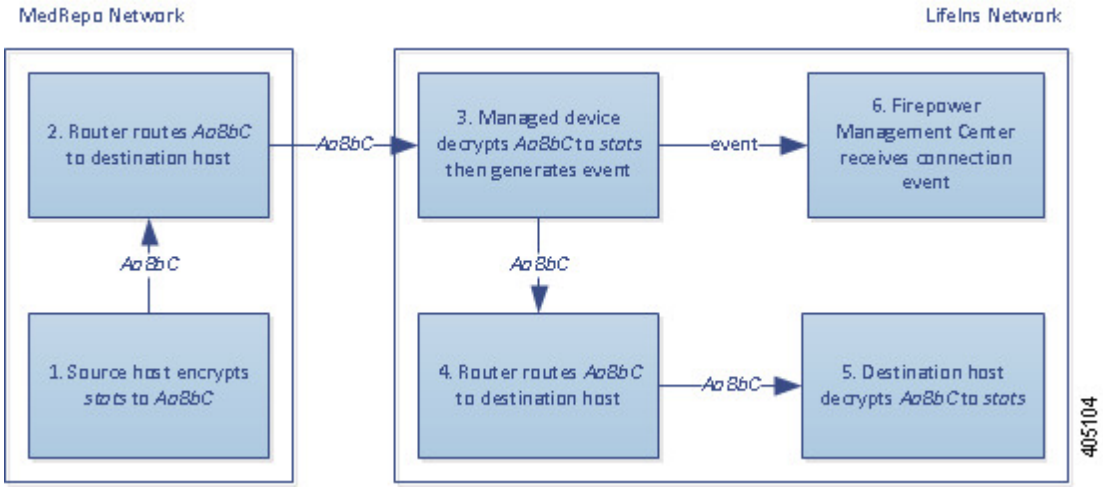


The following steps occur:

1. Having received the request to establish a TLS/SSL handshake from a client's browser, the Customer Service department server sends the server certificate (`cert`) as the next step in the TLS/SSL handshake to the LifeIns underwriter.
2. MedRepo's router receives the certificate and routes it to the LifeIns underwriter.
3. The managed device blocks the traffic without further inspection and ends the TCP connection. It generates a connection event.
4. The internal router does not receive the blocked traffic.
5. The underwriter does not receive the blocked traffic.
6. The Firepower Management Center receives the connection event.

Encrypted Traffic Inspection with a Private Key in an Inline Deployment

For all TLS/SSL-encrypted traffic sent from MedRepo to LifeIns's underwriting department, the system uses an uploaded server private key to obtain session keys, then decrypts the traffic and logs the connection. Legitimate traffic is allowed and reencrypted before being sent to the Underwriting department.

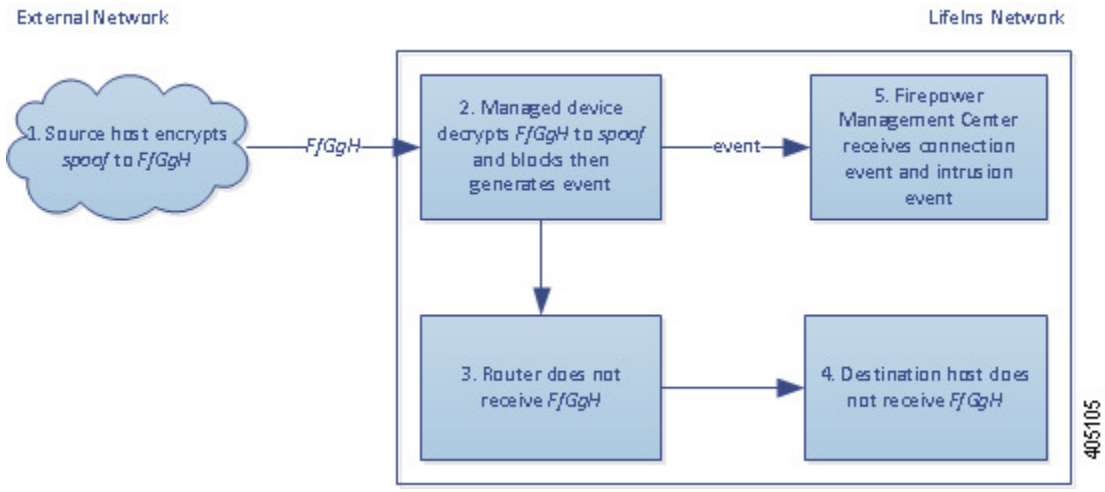


The following steps occur:

1. The user submits the plain text request (`stats`). The client encrypts this (`AaBbC`) and sends the encrypted traffic to the Underwriting department server.
2. The external router receives the traffic and routes it to the Underwriting department server.
3. The managed device uses the session key obtained with the uploaded known private key to decrypt this traffic to plain text (`stats`).

   The access control policy continues to process the decrypted traffic with the custom intrusion policy and does not find a spoof attempt. The device passes the encrypted traffic (`AaBbC`), then generates a connection event after the session ends.
4. The internal router receives the traffic and routes it to the Underwriting department server.
5. The Underwriting department server receives the encrypted information (`AaBbC`) and decrypts it to plain text (`stats`).
6. The Firepower Management Center receives the connection event with information about the encrypted and decrypted traffic.

In contrast, any decrypted traffic that is a spoof attempt is dropped. The system logs the connection and the spoof attempt.



The following steps occur:

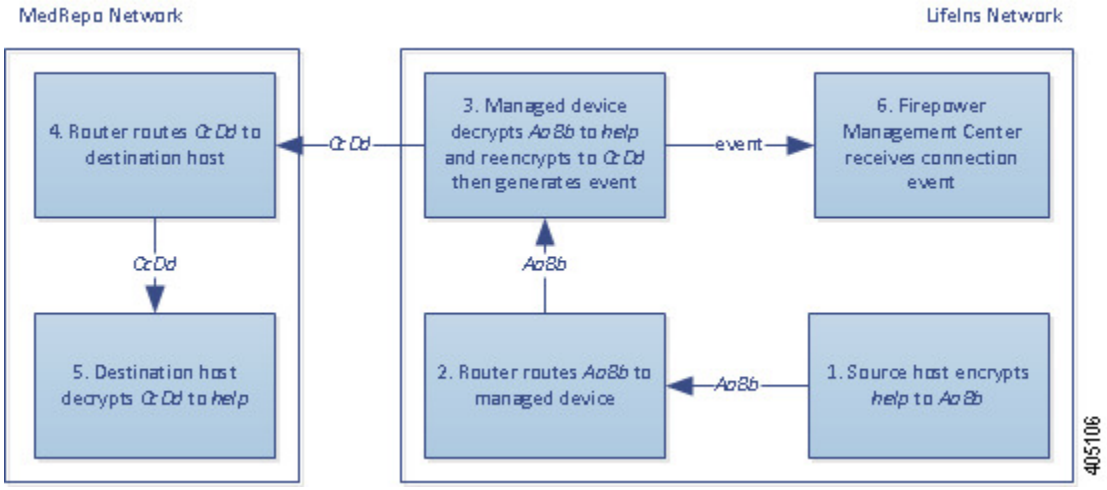
1. The user submits the plain text request (`spoof`), altering the traffic to appear to originate from MedRepo, LLC. The client encrypts this (`FfGgH`) and sends the encrypted traffic to the Underwriting department server.
2. The managed device uses the session key obtained with the uploaded known private key to decrypt this traffic to plain text (`spoof`).

   The access control policy continues to process the decrypted traffic with the custom intrusion policy and finds a spoof attempt. The device blocks the traffic, then generates an intrusion event. It generates a connection event after the session ends.
3. The internal router does not receive the blocked traffic.
4. The Underwriting department server does not receive the blocked traffic.
5. The Firepower Management Center receives a connection event with information about the encrypted and decrypted traffic, and an intrusion event for the spoofing attempt.

Encrypted Traffic Inspection with a Re-signed Certificate in an Inline Deployment

For all TLS/SSL-encrypted traffic sent from the new and junior underwriters to MedRepo's requests department, the system uses a re-signed server certificate to obtain session keys, then decrypts the traffic and logs the connection. Legitimate traffic is allowed and reencrypted before being sent to MedRepo.

Note When decrypting traffic in an inline deployment by re-signing the server certificate, the device acts as a man-in-the-middle. It creates two TLS/SSL sessions, one between client and managed device, one between managed device and server. As a result, each session contains different cryptographic session details.



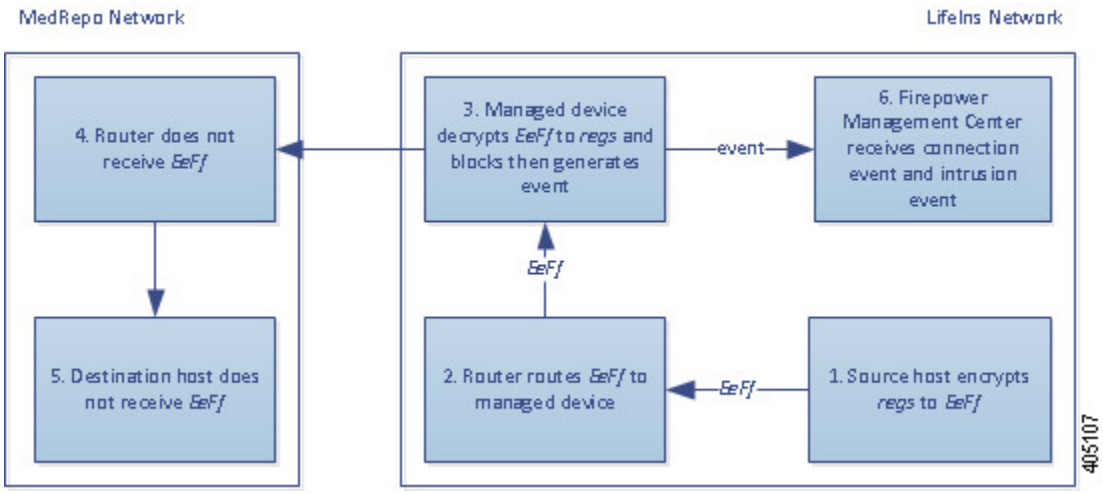
The following steps occur:

1. The user submits the plain text request (`help`). The client encrypts this (`AaBb`) and sends the encrypted traffic to the Requests department server.
2. The internal router receives the traffic and routes it to the Requests department server.
3. The managed device uses the session key obtained with a re-signed server certificate and private key to decrypt this traffic to plain text (`help`).

   The access control policy continues to process the decrypted traffic with the custom intrusion policy and does not find an improper request. The device reencrypts the traffic (`CcDd`), allowing it to pass. It generates a connection event after the session ends.
4. The external router receives the traffic and routes it to the Requests department server.
5. The Requests department server receives the encrypted information (`CcDd`) and decrypts it to plain text (`help`).
6. The Firepower Management Center receives the connection event with information about the encrypted and decrypted traffic.

Note Traffic encrypted with a re-signed server certificate causes client browsers to warn that the certificate is not trusted. To avoid this, add the CA certificate to the organization's domain root trusted certificates store or the client trusted certificates store.

In contrast, any decrypted traffic that contains information that does not meet regulatory requirements is dropped. The system logs the connection and the non-conforming information.



The following steps occur:

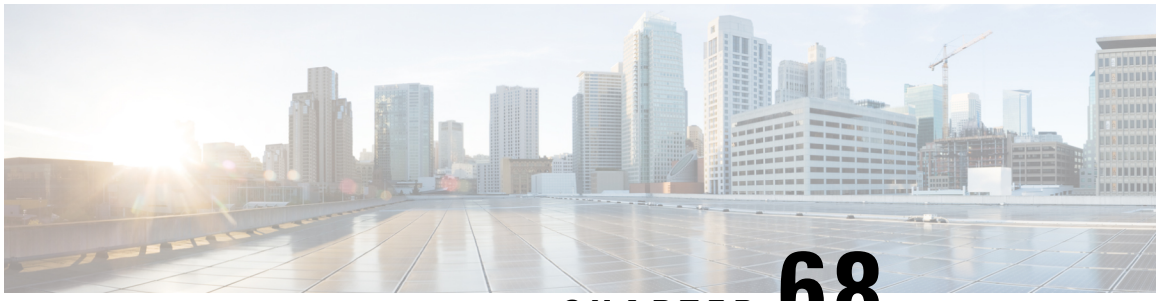
1. The user submits the plain text request (`regs`), which does not comply with regulatory requirements. The client encrypts this (`EeFf`) and sends the encrypted traffic to the Requests department server.
2. The internal router receives the traffic and routes it to the Requests department server.
3. The managed device uses the session key obtained with a re-signed server certificate and private key to decrypt this traffic to plain text (`regs`).

   The access control policy continues to process the decrypted traffic with the custom intrusion policy and finds an improper request. The device blocks the traffic, then generates an intrusion event. It generates a connection event after the session ends.
4. The external router does not receive the blocked traffic.
5. The Requests department server does not receive the blocked traffic.
6. The Firepower Management Center receives a connection event with information about the encrypted and decrypted traffic, and an intrusion event for the improper request.

History for TLS/SSL

Feature	Version	Details
Extended Master Secret extension supported (see RFC 7627)	6.3.0.1	The TLS Extended Master Secret extension is supported for SSL policies; specifically, policies with a rule action of **Decrypt - Resign** or **Decrypt - Known Key**.
Extended Master Secret extension supported (see RFC 7627)	6.2.3.9	The TLS Extended Master Secret extension is supported for SSL policies; specifically, policies with a rule action of **Decrypt - Resign** or **Decrypt - Known Key**.

Feature	Version	Details
SafeSearch supported	6.1.0	• The system displays an HTTP response page for connections decrypted by the SSL policy, then blocked (or interactively blocked) either by access control rules or by the access control policy default action. In these cases, the system encrypts the response page and sends it at the end of the reencrypted SSL stream. • SafeSearch filters objectionable content and stops people from searching adult sites.



CHAPTER 68

Start Creating SSL Policies

The following topics provide an overview of SSL policy creation, configuration, management, and logging.

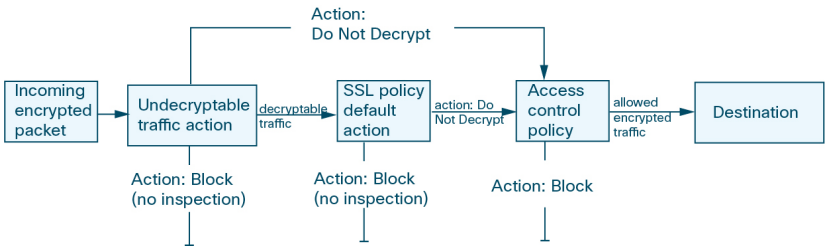
- SSL Policies Overview, on page 1337
- SSL Policy Default Actions, on page 1338
- Default Handling Options for Undecryptable Traffic, on page 1339
- Requirements and Prerequisites for SSL Policies, on page 1340
- Manage SSL Policies, on page 1340
- Create Basic SSL Policies, on page 1341
- Set Default Handling for Undecryptable Traffic, on page 1342
- Editing an SSL Policy, on page 1343

SSL Policies Overview

An *SSL policy* determines how the system handles encrypted traffic on your network. You can configure one or more SSL policies, associate an SSL policy with an access control policy, then deploy the access control policy to a managed device. When the device detects a TCP handshake, the access control policy first handles and inspects the traffic. If it subsequently identifies a TLS/SSL-encrypted session over the TCP connection, the SSL policy takes over, handling and decrypting the encrypted traffic.

Caution Adding or removing an SSL policy restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

The simplest SSL policy, as shown in the following diagram, directs the device where it is deployed to handle encrypted traffic with a single default action. You can set the default action to block decryptable traffic without further inspection, or to inspect undecrypted decryptable traffic with access control. The system can then either allow or block the encrypted traffic. If the device detects undecryptable traffic, it either blocks the traffic without further inspection or does not decrypt it, inspecting it with access control.



A more complex SSL policy can handle different types of undecryptable traffic with different actions, control traffic based on whether a certificate authority (CA) issued or trusts the encryption certificate, and use SSL rules to exert granular control over encrypted traffic logging and handling. These rules can be simple or complex, matching and inspecting encrypted traffic using multiple criteria.

Note Because TLS and SSL are often used interchangeably, we use the expression *TLS/SSL* to indicate that either protocol is being discussed. The SSL protocol has been deprecated by the IETF in favor of the more secure TLS protocol, so you can usually interpret *TLS/SSL* as referring to TLS only.

The exception is SSL policies. Because the FMC configuration option is **Policies** > **Access Control** > **SSL**, we use the term *SSL policies* although these policies are used to define rules for TLS and SSL traffic.

For more information about SSL and TLS protocols, see a resource such as SSL vs. TLS - What's the Difference?.

Related Topics

TLS/SSL Rule Conditions, on page 1357

SSL Policy Default Actions

The default action for an SSL policy determines how the system handles decryptable encrypted traffic that does not match any non-monitor rule in the policy. When you deploy an SSL policy that does not contain any TLS/SSL rules, the default action determines how all decryptable traffic on your network is handled. Note that the system does not perform any kind of inspection on encrypted traffic blocked by the default action.

Table 93: SSL Policy Default Actions

Default Action	Effect on Encrypted Traffic
Block	Block the TLS/SSL session without further inspection.
Block with reset	Block the TLS/SSL session without further inspection and reset the TCP connection. Choose this option if traffic uses a connectionless protocol like UDP. In that case, the connectionless protocol tries to reestablish the connection until it is reset. This action also displays a connection reset error in the browser so the user is informed that the connection is blocked.
Do not decrypt	Inspect the encrypted traffic with access control.

Related Topics

Create Basic SSL Policies, on page 1341

Default Handling Options for Undecryptable Traffic

Table 94: Undecryptable Traffic Types

Type	Description	Default Action	Available Action
Compressed Session	The TLS/SSL session applies a data compression method.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
SSLv2 Session	The session is encrypted with SSL version 2. Note that traffic is decryptable if the ClientHello message is SSL 2.0, and the remainder of the transmitted traffic is SSL 3.0.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
Unknown Cipher Suite	The system does not recognize the cipher suite.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
Unsupported Cipher Suite	The system does not support decryption based on the detected cipher suite.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
Session not cached	The TLS/SSL session has session reuse enabled, the client and server reestablished the session with the session identifier, and the system did not cache that session identifier.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
Handshake Errors	An error occurred during TLS/SSL handshake negotiation.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
Decryption Errors	An error occurred during traffic decryption.	Block	Block Block with Reset

When you first create an SSL policy, logging connections that are handled by the default action is disabled by default. Because the logging settings for the default action also apply to undecryptable traffic handling, logging connections handled by the undecryptable traffic actions is disabled by default.

Note that if your browser uses certificate pinning to verify a server certificate, you cannot decrypt this traffic by re-signing the server certificate.

Related Topics

Set Default Handling for Undecryptable Traffic, on page 1342

Requirements and Prerequisites for SSL Policies

Model Support

Any except NGIPSv.

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Manage SSL Policies

In the SSL policy editor, you can:

- Configure your policy.
- Add, edit, delete, enable, disable, and organize TLS/SSL rules.
- Add trusted CA certificates.
- Determine the handling for encrypted traffic the system cannot decrypt.
- Log traffic that is handled by the default action and undecryptable traffic actions.

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Policies** > **Access Control** > **SSL**.

Step 2 Manage SSL policies:

- Associate—To associate an SSL policy with an access control policy, see Associating Other Policies with Access Control, on page 1213.
- Compare—Click **Compare Policies**; see Comparing Policies, on page 355.
- Copy—Click **Copy** ( ).
- Create—Click **New Policy**; see Create Basic SSL Policies, on page 1341.
- Delete—Click **Delete** ( ). If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Deploy—Click **Deploy**; see Deploy Configuration Changes, on page 346.
- Edit—Click **Edit** ( ); see Editing an SSL Policy, on page 1343. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Import/Export—See About Configuration Import/Export, on page 173.
- Report—Click **Report** ( ); see Generating Current Policy Reports, on page 356.

Create Basic SSL Policies

To configure an SSL policy, you must give the policy a unique name and specify a default action.

Procedure

Step 1 Choose **Policies** > **Access Control** > **SSL**.

Step 2 Click **New Policy**.

Step 3 Give the policy a unique **Name** and, optionally, a **Description**.

Step 4 Specify the **Default Action**; see SSL Policy Default Actions, on page 1338.

Step 5 Configure logging options for the default action as described in Logging Connections with a Policy Default Action, on page 2221.

Step 6 Click **Save**.

What To Do Next

- Configure rules to add to your SSL policy; see Creating and Modifying TLS/SSL Rules, on page 1354.
- Set the default handling for undecryptable traffic; see Set Default Handling for Undecryptable Traffic, on page 1342.
- Configure logging options for default handling of undecryptable traffic; see Logging Connections with a Policy Default Action, on page 2221.

- Associate the SSL policy with an access control policy as described in Associating Other Policies with Access Control, on page 1213.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Set Default Handling for Undecryptable Traffic

You can set undecryptable traffic actions at the SSL policy level to handle certain types of encrypted traffic the system cannot decrypt or inspect. When you deploy an SSL policy that contains no TLS/SSL rules, the undecryptable traffic actions determine how all undecryptable encrypted traffic on your network is handled.

Depending on the type of undecryptable traffic, you can choose to:

- Block the connection.
- Block the connection, then reset it. This option is preferrable for connectionless protocols like UDP, which keep trying to connect until the connection is blocked.
- Inspect the encrypted traffic with access control.
- Inherit the default action from the SSL policy.

Procedure

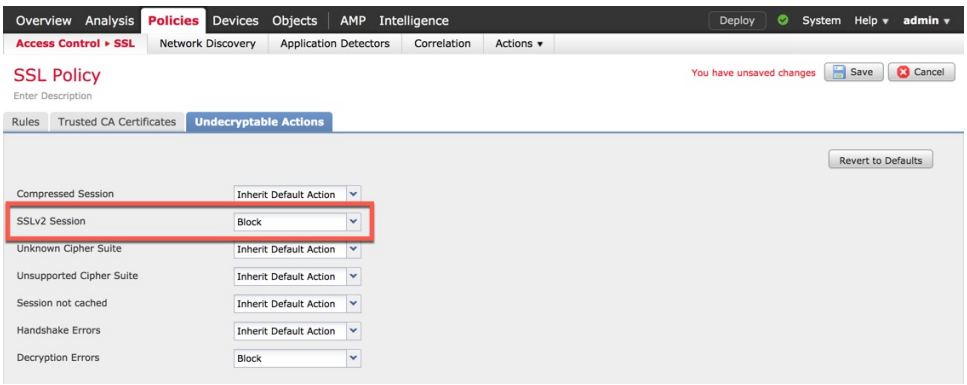
Step 1 In the SSL policy editor, click **Undecryptable Actions**.

Step 2 For each field, choose either the SSL policy's default action or another action you want to take on the type of undecryptable traffic. See Default Handling Options for Undecryptable Traffic, on page 1339 and SSL Policy Default Actions, on page 1338 for more information.

Step 3 Click **Save** to save the policy.

Example

For example, to block all SSLv2 traffic, set the options as follows:



What to do next

- Configure default logging for connections handled by the undecryptable traffic actions; see Logging Connections with a Policy Default Action, on page 2221.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Editing an SSL Policy

Only one person should edit a policy at a time, using a single browser window. If multiple users save the same policy, the last saved changes are retained. For your convenience, the system displays information on who (if anyone) is currently editing each policy. To protect the privacy of your session, a warning appears after 30 minutes of inactivity on the policy editor. After 60 minutes, the system discards your changes.

Procedure

Step 1 Choose **Policies** > **Access Control** > **SSL**.

Step 2 Click **Edit** ( ) next to the SSL policy you want to configure.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Configure the SSL policy:

- Describe—If you want to update your SSL policy description, click the **Description** field and enter the new description.
- Log—If you want to log connections for undecryptable traffic handling and traffic that does not match SSL rules, see Logging Connections with a Policy Default Action, on page 2221.
- Rename—If you want to rename your SSL policy, click the **Name** field and enter the new name.
- Set the default action—If you want to configure how your SSL policy handles traffic that does not match SSL rules, see SSL Policy Default Actions, on page 1338.
- Set the default action for undecryptable traffic—If you want to configure how your SSL policy handles undecryptable traffic, see Set Default Handling for Undecryptable Traffic, on page 1342.
- Trust—If you want to add trusted CA certificates to your SSL policy, see Trusting External Certificate Authorities, on page 1386.

Step 4 Edit the rules in your SSL policy:

- Add—If you want to add a rule, click **Add Rule**.
- Copy—If you want to copy a rule, right-click a selected rule and choose **Copy**.
- Cut—If you want to cut a rule, right-click a selected rule and choose **Cut**.
- Delete—To delete a rule, click **Delete** ( ) next to the rule, then click **OK**.
- Disable—To disable an enabled rule, right-click a selected rule, choose **State**, then choose **Disable**.

- Display—To display the configuration page for a specific rule attribute, click the name or value in the column for the condition on the row for the rule. For example, click the name or value in the **Source Networks** column to display the Networks page for the selected rule. See Network-Based TLS/SSL Rule Conditions, on page 1366 for more information.
- Edit—To edit a rule, click **Edit** () next to the rule.
- Enable—To enable a disabled rule, right-click a selected rule, choose **State**, then choose **Enable**. Disabled rules are dimmed and marked `(disabled)` beneath the rule name.
- Paste—To paste a cut or copied rule, right-click a selected rule and choose **Paste Above** or **Paste Below**.

Step 5 Save or discard your configuration:

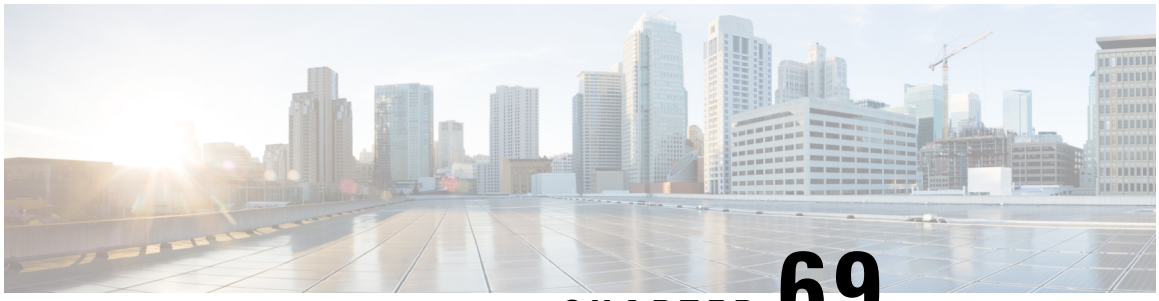
- To save your changes and continue editing, click **Save**.
- To discard your changes, click **Cancel** and, if prompted, click **OK**.

What to do next

- If the SSL policy is not already associated with an access control policy, associate it as described in Associating Other Policies with Access Control, on page 1213.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Creating and Modifying TLS/SSL Rules, on page 1354



CHAPTER 69

Get Started with TLS/SSL Rules

The following topics provide an overview of creating, configuring, managing, and troubleshooting TLS/SSL rules:

Note Because TLS and SSL are often used interchangeably, we use the expression *TLS/SSL* to indicate that either protocol is being discussed. The SSL protocol has been deprecated by the IETF in favor of the more secure TLS protocol, so you can usually interpret *TLS/SSL* as referring to TLS only.

The exception is SSL policies. Because the FMC configuration option is **Policies** > **Access Control** > **SSL**, we use the term *SSL policies* although these policies are used to define rules for TLS and SSL traffic.

For more information about SSL and TLS protocols, see a resource such as SSL vs. TLS - What's the Difference?.

- TLS/SSL Rules Overview, on page 1345
- TLS/SSL Rule Guidelines and Limitations, on page 1345
- Requirements and Prerequisites for TLS/SSL Rules, on page 1352
- Encrypted Traffic Inspection Configuration, on page 1353
- Creating and Modifying TLS/SSL Rules, on page 1354
- TLS/SSL Rule Conditions, on page 1357
- TLS/SSL Rule Actions, on page 1359
- TLS/SSL Rules Management, on page 1363

TLS/SSL Rules Overview

TLS/SSL rules provide a granular method of handling encrypted traffic across multiple managed devices, whether blocking the traffic without further inspection, not decrypting the traffic and inspecting it with access control, or decrypting the traffic for access control analysis.

TLS/SSL Rule Guidelines and Limitations

Keep the following points in mind when setting up your TLS/SSL rules. Properly configuring TLS/SSL rules is a complex task, but one that is essential to building an effective deployment that handles encrypted traffic. Many factors influence how you configure rules, including certain application behavior that you cannot control.

In addition, rules can preempt each other, require additional licenses, or contain invalid configurations. Thoughtfully configured rules can also reduce the resources required to process network traffic. Creating overly complex rules and ordering rules the wrong way can adversely affect performance.

For detailed information, see Best Practices for Access Control Rules, on page 1194.

Related Topics

Rule and Other Policy Warnings, on page 392
Best Practices for Access Control Rules, on page 1194
Guideline for Using TLS/SSL Decryption, on page 1346
TLS/SSL Rule Unsupported Features, on page 1346
TLS/SSL Do Not Decrypt Guidelines, on page 1347
TLS/SSL Decrypt - Resign Guidelines, on page 1347
TLS/SSL Decrypt - Known Key Guidelines, on page 1349
TLS/SSL Block Guidelines, on page 1350
TLS/SSL Certificate Pinning Guidelines, on page 1351
TLS/SSL Heartbeat Guidelines, on page 1351
TLS/SSL Anonymous Cipher Suite Limitation, on page 1351
TLS/SSL Normalizer Guidelines, on page 1351
Other TLS/SSL Rule Guidelines, on page 1352
SSL Rule Order, on page 1197

Guideline for Using TLS/SSL Decryption

Set up **Decrypt - Resign** or **Decrypt - Known Key** rules *only* if your managed device handles encrypted traffic. Decryption rules require processing overhead that can impact performance.

TLS/SSL Rule Unsupported Features

Zone conditions and passive interfaces

If you create a **Decrypt - Resign** rule, and later add a security zone with passive interfaces to a zone condition, the system displays a warning icon next to the rule. Because you cannot decrypt traffic by re-signing a certificate in a passive deployment, the rule has no effect until you remove the passive interfaces from the rule or change the rule action.

Unsupported characters in rule names

Do not use accented characters (for example, Comunicación) in TLS/SSL rule rule names; doing so prevents the policy from being deployed to managed devices.

TLS 1.3 not supported

The Firepower System does not currently support TLS version 1.3 encryption or decryption. When users visit a web site that negotiates TLS 1.3 encryption, users might see errors similar to the following in their web browser:

- **ERR_SSL_PROTOCOL_ERROR**
- **SEC_ERROR_BAD_SIGNATURE**
- **ERR_SSL_VERSION_INTERFERENCE**

For more information about how to control this behavior, contact Cisco TAC.

TLS/SSL Do Not Decrypt Guidelines

You should not decrypt traffic if doing so is forbidden by:

- Law; for example, some jurisdictions forbid decrypting financial information
- Company policy; for example, your company might forbid decrypting privileged communications
- Privacy regulations
- Traffic that uses certificate pinning (also referred to as *TLS/SSL pinning*) must remain encrypted to prevent breaking the connection

If you elect to bypass decryption for certain types of traffic, no processing is done on the traffic. The encrypted traffic is first evaluated by SSL policy and then proceeds to the access control policy, where a final allow or block decision is made. Encrypted traffic can be allowed or blocked on any TLS/SSL rule condition, including, but not limited to:

- Certificate status (for example, expired or invalid certificate)
- Protocol (for example, the nonsecure SSL protocol)
- Network (security zone, IP address, VLAN tag, and so on)
- Exact URL or URL category
- Port
- User group

TLS/SSL Decrypt - Resign Guidelines

You can associate one internal Certificate Authority (CA) certificate and private key with the **Decrypt - Resign** action. If traffic matches this rule, the system re-signs the server certificate with the CA certificate, then acts as a man-in-the-middle. It creates two TLS/SSL sessions, one between client and managed device, one between managed device and server. Each session contains different cryptographic session details, and allows the system to decrypt and reencrypt traffic.

Best practices

We recommend the following:

- Use the **Decrypt - Resign** rule action for decrypting *outgoing* traffic, as opposed to incoming traffic for which we recommend the **Decrypt - Known Key** rule action.

  For more information about **Decrypt - Known Key**, see TLS/SSL Decrypt - Known Key Guidelines, on page 1349.

- Always check the **Replace Key Only** check box when you set up a **Decrypt - Resign** rule action.

  When a user browses to a web site that uses a *self-signed* certificate, the user sees a security warning in the web browser and is aware that they are communicating with an unsecure site.

  When a user browses to a web site that uses a trusted certificate, the user does not see a security warning.

Details

If you configure a rule with the **Decrypt - Resign** action, the rule matches traffic based on the referenced internal CA certificate's signature algorithm type, in addition to any configured rule conditions. Because you associate one CA certificate with a **Decrypt - Resign** action, you cannot create a TLS/SSL rule that decrypts multiple types of outgoing traffic encrypted with different signature algorithms. In addition, any external certificate objects and cipher suites you add to the rule must match the associated CA certificate encryption algorithm type.

For example, outgoing traffic encrypted with an elliptic curve (EC) algorithm matches a **Decrypt - Resign** rule only if the action references an EC-based CA certificate; you must add EC-based external certificates and cipher suites to the rule to create certificate and cipher suite rule conditions.

Similarly, a **Decrypt - Resign** rule that references an RSA-based CA certificate matches only outgoing traffic encrypted with an RSA algorithm; outgoing traffic encrypted with an EC algorithm does not match the rule, even if all other configured rule conditions match.

Guidelines and limitations

Also note the following:

Anonymous cipher suite unsupported

By nature, anonymous cipher suites are not used for authentication and do not use key exchanges. There are limited uses for anonymous cipher suites; for more information, see RFC 5246, appendix F.1.1.1. (Replaced for TLS 1.3 by RFC 8446 appendix C.5.)

You cannot use the **Decrypt - Resign** or **Decrypt - Known Key** action in the rule because anonymous cipher suites are not used for authentication.

Inline or tap mode limitation

You cannot use the **Decrypt - Resign** action in a passive or inline (tap mode) deployment because the device does not directly inspect traffic. If you create a rule with the **Decrypt - Resign** action that contains passive or inline (tap mode) interfaces within a security zone, the policy editor displays a **Warning** ( ) next to the rule.

If your SSL policy targets a device with passive or inline (tap mode) interfaces, and contains a **Decrypt - Resign** rule, the system displays an **Information** ( ) next to the rule. If you later add a zone condition to the TLS/SSL rule that contains passive or inline (tap mode) interfaces, the system displays a **Warning** ( ). If you deploy an SSL policy that contains a **Decrypt - Resign** rule to a device with passive or inline (tap mode) interfaces, any TLS/SSL sessions that match the rule fail.

Decrypt - Resign rule action and a Certificate Signing Request

To use a **Decrypt - Resign** rule action, you should create a Certificate Signing Request (CSR) and have it signed by a trusted certificate authority. (You can use the FMC to create a CSR: **Objects** > **Object Management** > **PKI** > **Internal CAs**.)

To be used in a **Decrypt - Resign** rule, your certificate authority (CA) must have at least one of the following extensions:

- **CA: TRUE**

  For more information, see the discussion of Basic Constraints in RFC3280, section 4.2.1.10.

- **KeyUsage=CertSign**

For more information see RFC 5280, section 4.2.1.3.

To verify your CSR or CA has at least one of the preceding extensions, you can use the **openssl** command as discussed in a reference such as the openssl documentation.

This is necessary because for **Decrypt - Resign** inspection to work, the certificate that used in the TLS/SSL policy generates certificates on-the-fly and signs them so as to act as man-in-the middle and proxy all TLS/SSL connections.

Non-matching cipher suite

The following error is displayed if you attempt to save a TLS/SSL rule with a cipher suite that does not match the certificate. To resolve the issue, see Verify TLS/SSL Cipher Suites, on page 1397.

```
Traffic cannot match this rule; none of your selected cipher suites contain a
signature algorithm that the resigning CA's signature algorithm
```

Untrusted Certificate Authority

If the client does not trust the Certificate Authority (CA) used to re-sign the server certificate, it warns the user that the certificate should not be trusted. To prevent this, import the CA certificate into the client trusted CA store. Alternatively, if your organization has a private PKI, you can issue an intermediate CA certificate signed by the root CA which is automatically trusted by all clients in the organization, then upload that CA certificate to the device.

HTTP proxy limitation

The system cannot decrypt traffic if an HTTP proxy is positioned between a client and your managed device, and the client and server establish a tunneled TLS/SSL connection using the CONNECT HTTP method. The **Handshake Errors** undecryptable action determines how the system handles this traffic.

Upload signed CA

If you create an internal CA object and choose to generate a certificate signing request (CSR), you cannot use this CA for a **Decrypt - Resign** action until you upload the signed certificate to the object.

Use of certain rule conditions



Note Use the **Cipher Suite** and **Version** rule conditions *only* in rules with either the **Block** or **Block with reset** rule actions. The use of these conditions in rules with other rule actions can interfere with the system's ClientHello processing, resulting in unpredictable performance.

TLS/SSL Decrypt - Known Key Guidelines

When you configure the **Decrypt - Known Key** action, you can associate one or more server certificates and paired private keys with the action. If traffic matches the rule, and the certificate used to encrypt the traffic matches the certificate associated with the action, the system uses the appropriate private key to obtain the session encryption and decryption keys. Because you must have access to the private key, this action is best suited to decrypt traffic incoming to servers your organization controls.

Also note the following:

Limitation for DHE and ECDHE key exchanges

You cannot use the **Decrypt - Known Key** action in a passive deployment if the cipher suite used to establish the TLS/SSL connection applies either the Diffie-Hellman ephemeral (DHE) or the elliptic

curve Diffie-Hellman ephemeral (ECDHE) key exchange algorithm. If your SSL policy targets a device with passive or inline (tap mode) interfaces, and contains a **Decrypt - Known Key** rule with a cipher suite condition containing either a DHE or an ECDHE cipher suite, the system displays an **Information** () next to the rule. If you later add a zone condition to the TLS/SSL rule that contains passive or inline (tap mode) interfaces, the system displays a **Warning** ().

Anonymous cipher suite unsupported

By nature, anonymous cipher suites are not used for authentication and do not use key exchanges. There are limited uses for anonymous cipher suites; for more information, see RFC 5246, appendix F.1.1.1. (Replaced for TLS 1.3 by RFC 8446 appendix C.5.)

You cannot use the **Decrypt - Resign** or **Decrypt - Known Key** action in the rule because anonymous cipher suites are not used for authentication.

Cannot match on Distinguished Name or Certificate

You cannot match on **Distinguished Name** or **Certificate** conditions when creating a TLS/SSL rule with a **Decrypt - Known Key** action. The assumption is that if this rule matches traffic, the certificate, subject DN, and issuer DN already match the certificate associated with the rule.

Mismatched signature algorithm

If you configure a rule with the **Decrypt - Resign** action, and mismatch signature algorithm type for one or more external certificate objects or cipher suites, the policy editor displays an **Information** () next to the rule. If you mismatch signature algorithm type for all external certificate objects, or all cipher suites, the policy displays a warning icon **Warning** () next to the rule, and you cannot deploy the access control policy associated with the SSL policy.

Certificate pinning

If the customer's browser uses certificate pinning to verify a server certificate, you cannot decrypt this traffic by re-signing the server certificate. To allow this traffic, configure a TLS/SSL rule with the **Do not decrypt** action to match the server certificate common name or distinguished name.

Use of certain rule conditions



Note Use the **Cipher Suite** and **Version** rule conditions *only* in rules with either the **Block** or **Block with reset** rule actions. The use of these conditions in rules with other rule actions can interfere with the system's ClientHello processing, resulting in unpredictable performance.

TLS/SSL Block Guidelines

If decrypted traffic matches an access control rule with an action of **Interactive Block** or **Interactive Block with reset**, the system displays a customizable response page.

Provided you enabled logging in your rule, two connection events are displayed (in **Analysis** > **Events** > **Connections**): One event for the interactive block and another event to indicate whether or not the user chose to continue to the site or not.

TLS/SSL Certificate Pinning Guidelines

Some applications use a technique referred to as *TLS/SSL pinning* or *certificate pinning*, which embeds the fingerprint of the original server certificate in the application itself. As a result, if you configured a TLS/SSL rule with a **Decrypt - Resign** action, when the application receives a resigned certificate from a managed device, validation fails and the connection is aborted.

Because TLS/SSL pinning is used to avoid man-in-the-middle attacks, there is no way to prevent or work around it. You have the following options:

- Create a **Do Not Decrypt** for those applications rule ordered before **Decrypt - Resign** rules.
- Instruct users to access the applications using a web browser.

For more information about rule ordering, see SSL Rule Order, on page 1197.

TLS/SSL Heartbeat Guidelines

Some applications use the *TLS heartbeat* extension to the Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS) protocols defined by RFC6520. TLS heartbeat provides a way to confirm the connection is still alive—either the client or server sends a specified number of bytes of data and requests the other party echo the response. If this is successful, encrypted data is sent.

You can configure a **Max Heartbeat Length** in a Network Analysis Policy (NAP) to determine how to handle TLS heartbeats. For more information, see The SSL Preprocessor, on page 1746.

TLS/SSL Anonymous Cipher Suite Limitation

By nature, anonymous cipher suites are not used for authentication and do not use key exchanges. There are limited uses for anonymous cipher suites; for more information, see RFC 5246, appendix F.1.1.1. (Replaced for TLS 1.3 by RFC 8446 appendix C.5.)

You cannot use the **Decrypt - Resign** or **Decrypt - Known Key** action in the rule because anonymous cipher suites are not used for authentication.

You can add an anonymous cipher suite to the **Cipher Suite** condition in a TLS/SSL rule, but the system automatically strips anonymous cipher suites during ClientHello processing. For the system to use the rule, you must also configure your TLS/SSL rules in an order that prevents ClientHello processing. For more information, see SSL Rule Order, on page 1197.

TLS/SSL Normalizer Guidelines

If you enable the **Normalize Excess Payload** option in the inline normalization preprocessor, when the preprocessor normalizes decrypted traffic, it might drop a packet and replace it with a trimmed packet. This does not end the TLS/SSL session. If the traffic is allowed, the trimmed packet is encrypted as part of the TLS/SSL session.

Other TLS/SSL Rule Guidelines

Users and groups

If you add a group or user to a rule, then change your realm settings to exclude that group or user, the rule has no effect. (The same applies to disabling the realm.) For more information about realms, see Create a Realm, on page 1949.

Categories in TLS/SSL rules

If your SSL policy has a **Decrypt - Resign** action but web sites are not being decrypted, check **Category** page on rules associated with that policy.

In some cases, a web site redirects to another site for authentication or other purposes and the redirected site might have a different URL categorization than the site you're trying to decrypt. For example, `gmail.com` (**Web based email** category) redirects to `accounts.gmail.com` (**Internet Portals** category) for authentication. Be sure to include all relevant categories in the SSL rule.



Note In order to fully process traffic based on URL category, you must also configure URL filtering. See the URL Filtering, on page 1229 chapter.

Query for URLs not in the local database

If you create a **Decrypt - Resign** rule and users browse to a web site whose category and reputation are not in the local database, data might not be decrypted. Some web sites are not categorized in the local database and, if not, data from those web sites is not decrypted by default.

You can control this behavior with the setting **System > Integration > Cisco CSI** , and check **Query Cisco CSI for Unknown URLs**.

For more information about this option, see Cisco Clouds, on page 2425.

Requirements and Prerequisites for TLS/SSL Rules

Model Support

Any except NGIPSv.

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Encrypted Traffic Inspection Configuration

You must create reusable public key infrastructure (PKI) objects to control encrypted traffic based on encrypted session characteristics and decrypt encrypted traffic. You can add this information on the fly when uploading trusted certificate authority (CA) certificates to the SSL policy and creating SSL rule conditions, creating the associated object in the process. However, configuring these objects ahead of time reduces the chance of improper object creation.

Decrypting Encrypted Traffic with Certificates and Paired Keys

The system can decrypt incoming encrypted traffic if you configure an internal certificate object by uploading the server certificate and private key used to encrypt the session. If you reference that object in an SSL rule with an action of **Decrypt - Known Key** and traffic matches that rule, the system uses the uploaded private key to decrypt the session.

The system can also decrypt outgoing traffic if you configure an internal CA object by uploading a CA certificate and private key. If you reference that object in a TLS/SSL rule with an action of **Decrypt - Resign** and traffic matches that rule, the system re-signs the server certificate passed to the client browser, then acts as a man-in-the-middle to decrypt the session.You can optionally replace the self-signed certificate key only and not the entire certificate, in which case users see a self-signed certificate key notice in the browser.

Controlling Traffic Based on Encrypted Session Characteristics

The system can control encrypted traffic based on the cipher suite or server certificate used to negotiate the session. You can configure one of several different reusable objects and reference the object in a TLS/SSL rule condition to match traffic. The following table describes the different types of reusable objects you can configure:

If you configure...	You can control encrypted traffic based on whether...
A cipher suite list containing one or more cipher suites	The cipher suite used to negotiate the encrypted session matches a cipher suite in the cipher suite list
A trusted CA object by uploading a CA certificate your organization trusts	The trusted CA trusts the server certificate used to encrypt the session, whether: • The CA issued the certificate directly • The CA issued a certificate to an intermediate CA that issued the server certificate
An external certificate object by uploading a server certificate	The server certificate used to encrypt the session matches the uploaded server certificate
A distinguished name object containing a certificate subject or issuer distinguished name	The subject or issuer common name, country, organization, or organizational unit on the certificate used to encrypt the session matches the configured distinguished name

Related Topics

Cipher Suite Lists, on page 443
Distinguished Name Objects, on page 444
PKI Objects, on page 447

Creating and Modifying TLS/SSL Rules

Procedure

Step 1 Log in to the Firepower Management Center.

Step 2 Click **Policies** > **Access Control** > **SSL**.

Step 3 Click **Edit** ( ) next to the SSL policy.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 You have the following choices:

- To add a new rule, click **Add Rule**.
- To edit an existing rule, click **Edit** ( ).

Step 5 Enter a **Name** for the rule.

Do not use accented characters (for example, Comunicación) in TLS/SSL rule rule names; doing so prevents the policy from being deployed to managed devices.

Step 6 Specify whether the rule is **Enabled**.

Step 7 Specify the rule position; see TLS/SSL Rule Order Evaluation, on page 1317.

Step 8 Click a rule **Action**; see Configuring TLS/SSL Rule Actions, on page 1361.

Step 9 Configure the rule's conditions; see TLS/SSL Rule Condition Types, on page 1358.

Step 10 Click **Save**.

If the following error displays, see Verify TLS/SSL Cipher Suites, on page 1397: **Traffic cannot match this rule; none of your selected cipher suites contain a signature algorithm that the resigning CA's signature algorithm**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Adding a TLS/SSL Rule to a Rule Category

Procedure

Step 1 In the SSL rule editor, from the **Insert** drop-down list, select **Into Category**, then select the category you want to use.

Step 2 Click **Save**.

Tip When you save the rule, it is placed last in that category.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Positioning a TLS/SSL Rule by Number

Procedure

Step 1 In the SSL rule editor, from the **Insert** drop-down list, select **above rule** or **below rule**, then type the appropriate rule number.

Step 2 Click **Save**.

Tip When you save the rule, it is placed where you specified.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

TLS/SSL Rule Search

You can search the list of TLS/SSL rules for matching values using an alphanumeric string, including spaces and printable, special characters. The search inspects the rule name and any rule condition you have added to the rule. For rule conditions, the search matches any name or value you can add for each condition type (zone, network, application, and so on). This includes individual object names or values, group object names, individual object names or values within a group, and literal values.

You can use complete or partial search strings. The column for matching values is highlighted for each matching rule. For example, if you search on all or part of the string `100Bao`, at a minimum, the Applications column is highlighted for each rule where you have added the 100Bao application. If you also have a rule named `100Bao`, both the Name and Applications columns are highlighted.

You can navigate to each previous or next matching rule. A status message displays the current match and the total number of matches.

Matches may occur on any page of a multi-page rule list. When the first match is not on the first page, the page where the first match occurs is displayed. Selecting the next match when you are at the last match takes you to the first match, and selecting the previous match when you are at the first match takes you to the last match.

Searching TLS/SSL Rules

Procedure

Step 1 In the SSL policy editor, click the **Search Rules** prompt, type a search string, then press Enter.

Tip Columns for rules with matching values are highlighted, with differentiated highlighting for the indicated (first) match.

Step 2 Find the rules you are interested in:

- To navigate between matching rules, click **Next-Match** or **Previous-Match**.
- To refresh the page and clear the search string and any highlighting, click **Clear** ( ).

Enabling and Disabling TLS/SSL Rules

When you create a TLS/SSL rule, it is enabled by default. If you disable a rule, the system does not use it to evaluate network traffic and stops generating warnings and errors for that rule. When viewing the list of rules in an SSL policy, disabled rules are grayed out, although you can still modify them. Note that you can also enable or disable a TLS/SSL rule using the rule editor.

Procedure

Step 1 In the SSL policy editor, right-click a rule and choose a rule state.

Step 2 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Moving a TLS/SSL Rule

Procedure

Step 1 In the SSL policy editor, select the rules by clicking in a blank area for each rule.

Step 2 Right-click the rule and select **Cut**.

Step 3 Right-click a blank area for a rule next to where you want to paste the cut rules and select **Paste above** or **Paste below.**

Tip You cannot copy and paste TLS/SSL rules between two different SSL policies.

Step 4 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Adding a New TLS/SSL Rule Category

You can create custom categories between the Standard Rules and Root Rules categories to further organize your rules without having to create additional policies. You can rename and delete categories that you add. You cannot move these categories, but you can move rules into, within, and out of them.

Procedure

Step 1 In the policy editor, click **Add Category**.

Tip If your policy already contains rules, you can click a blank area in the row for an existing rule to set the position of the new category before you add it. You can also right-click an existing rule and select **Insert new category**.

Step 2 Type a **Name**.

Step 3 You have the following choices:

- Select **above Category** from the first **Insert** drop-down list, then select the category above which you want to position the rule from the second drop-down list.
- Select **below rule** from the drop-down list, then enter an existing rule number. This option is valid only when at least one rule exists in the policy.
- Select **above rule** from the drop-down list, then, enter an existing rule number. This option is valid only when at least one rule exists in the policy.

Step 4 Click **OK**.

Tip Rules in a category you delete are added to the category above.

Step 5 Click **Save**.

TLS/SSL Rule Conditions

An SSL rule's conditions identify the type of encrypted traffic the rule handles. Conditions can be simple or complex, and you can specify more than one condition type per rule. Only if traffic meets all the conditions in a rule does the rule apply to the traffic.

If you do not configure a particular condition for a rule, the system does not match traffic based on that criterion. For example, a rule with a certificate condition but no version condition evaluates traffic based on the server certificate used to negotiate the session, regardless of the session SSL or TLS version.

Every TLS/SSL rule has an associated action that determines the following for matching encrypted traffic:

- Handling: Most importantly, the rule action governs whether the system will monitor, trust, block, or decrypt encrypted traffic that matches the rule's conditions
- Logging: The rule action determines when and how you can log details about matching encrypted traffic.

Your TLS/SSL inspection configuration handles, inspects, and logs decrypted traffic:

- The SSL policy's undecryptable actions handle traffic that the system cannot decrypt.
- The policy's default action handles traffic that does not meet the condition of any non-Monitor TLS/SSL rule.

You can log a connection event when the system blocks or trusts an encrypted session. You can also force the system to log connections that it decrypts for further evaluation by access control rules, regardless of how the system later handles or inspects the traffic. Connection logs for encrypted sessions contain details about the encryption, such as the certificate used to encrypt that session. You can log only end-of-connection events, however:

- For blocked connections (Block, Block with reset), the system immediately ends the sessions and generates an event
- For trusted connections (Do not decrypt), the system generates an event when the session ends

TLS/SSL Rule Condition Types

When you add or edit an SSL rule, use the tabs on the left side of the lower portion of the rule editor to add and edit rule conditions.

Table 95: TLS/SSL Rule Condition Types

This Condition...	Matches Encrypted Traffic...	Details
Zones	Entering or leaving a device via an interface in a specific security zone	A security zone is a logical grouping of one or more interfaces according to your deployment and security policies. Interfaces in a zone may be located across multiple devices.
Networks	By its source or destination IP address, country, or continent	You can explicitly specify IP addresses. The geolocation feature also allows you to control traffic based on its source or destination country or continent.
VLAN Tags	Tagged by VLAN	The system uses the innermost VLAN tag to identify a packet by VLAN.
Ports	By its source or destination port	You can control encrypted traffic based on the TCP port.
Users	By the user involved in the session	You can control encrypted traffic based on the LDAP user logged into a host involved in an encrypted, monitored session. You can control traffic based on individual users or groups retrieved from a Microsoft Active Directory server.

This Condition...	Matches Encrypted Traffic...	Details
Applications	By the application detected in a session	You can control access to individual applications in encrypted sessions, or filter access according to basic characteristics: type, risk, business relevance, and categories.
Categories	By the URL requested in the session, based on the certificate subject distinguished name	You can limit the websites that users on your network can access based on the URL's general classification and risk level.
Distinguished Names	The URL the user enters in the browser matches the Common Name (CN), or the URL is contained in the certificate's Subject Alternative Name (SAN)	You can control encrypted traffic based on the CA that issued a server certificate, or the server certificate holder.
Certificates	By the server certificate used to negotiate the encrypted session	You can control encrypted traffic based on the server certificate passed to the user's browser in order to negotiate the encrypted session.
Certificate Status	By properties of the server certificate used to negotiate the encrypted session	You can control encrypted traffic based on a server certificate's status.
Cipher Suites	By the cipher suite used to negotiate the encrypted session	You can control encrypted traffic based on the cipher suite selected by the server to negotiate the encrypted session.
Versions	By the version of SSL or TLS used to encrypt the session	You can control encrypted traffic based on the version of SSL or TLS used to encrypt the session.

Related Topics

Network-Based TLS/SSL Rule Conditions, on page 1366
User-Based TLS/SSL Rule Conditions, on page 1372
Reputation-Based URL Blocking in Encrypted Traffic, on page 1378
Server Certificate-Based TLS/SSL Rule Conditions, on page 1379
ClientHello Message Handling, on page 1307

TLS/SSL Rule Actions

The following sections discuss the actions available with TLS/SSL rules.

Related Topics

TLS/SSL Rule Decrypt Actions, on page 1361
TLS/SSL Rule Blocking Actions, on page 1360
TLS/SSL Rule Do Not Decrypt Action, on page 1360
TLS/SSL Rule Monitor Action, on page 1359

TLS/SSL Rule Monitor Action

The **Monitor** action is not designed to permit or deny traffic. Rather, its primary purpose is to force connection logging, regardless of how matching traffic is eventually handled. Traffic is then matched against additional rules, if present, to determine whether to trust, block, or decrypt it. The first non-Monitor rule matched

determines traffic flow and any further inspection. If there are no additional matching rules, the system uses the default action.

Because the primary purpose of Monitor rules is to track network traffic, the system automatically logs end-of connection events for monitored traffic to the Firepower Management Center database, regardless of the logging configuration of the rule or default action that later handles the connection.

TLS/SSL Rule Do Not Decrypt Action

The **Do Not Decrypt** action passes encrypted traffic for evaluation by the access control policy's rules and default action. Because some access control rule conditions require unencrypted traffic, this traffic may match fewer rules. The system cannot perform deep inspection on encrypted traffic, such as intrusion or file inspection.

Typical reasons for a **Do Not Decrypt** rule action include:

- When decrypting TLS/SSL traffic is prohibited by law.
- Sites you know you can trust.
- Sites you can disrupt by inspecting traffic (such as Windows Update).
- To view the values of TLS/SSL fields using connection events. (You do not need to decrypt traffic to view connection event fields.) For more information, see Requirements for Populating Connection Event Fields, on page 2240.

For more information, see Default Handling Options for Undecryptable Traffic, on page 1339

TLS/SSL Rule Blocking Actions

The Firepower System provides the following TLS/SSL rule actions for traffic you do not want to pass through the system:

- **Block** to terminate the connection, resulting in an error in the client browser.

  The error message does not indicate the site was blocked due to policy. Instead, errors might indicate that there are no common encryption algorithms. It is not obvious from this message that you blocked the connection on purpose.

- **Block with reset** to terminate and reset the connection, resulting in an error in the client browser.

  The error indicates the connection was reset but does not indicate why.



Tip You cannot use the **Block** or **Block with reset** action in a passive or inline (tap mode) deployment because the device does not directly inspect the traffic. If you create a rule with the **Block** or **Block with reset** action that contains passive or inline (tap mode) interfaces within a security zone condition, the policy editor displays a warning () next to the rule.

Related Topics

About HTTP Response Pages, on page 1243

TLS/SSL Rule Decrypt Actions

The **Decrypt - Known Key** and **Decrypt - Resign** actions decrypt encrypted traffic. The system inspects decrypted traffic with access control. Access control rules handle decrypted and unencrypted traffic identically — you can inspect it for discovery data as well as detect and block intrusions, prohibited files, and malware. The system reencrypts allowed traffic before passing it to its destination.

We recommend you use a certificate from a trusted Certificate Authority (CA) to decrypt traffic. This prevents **Invalid Issuer** from being displayed in the SSL Certificate Status column in connection events.

For more information about adding trusted objects, see Trusted Certificate Authority Objects, on page 452.

Configuring TLS/SSL Rule Actions

Before you begin

See:

- TLS/SSL Rule Blocking Actions, on page 1360
- TLS/SSL Rule Do Not Decrypt Action, on page 1360
- TLS/SSL Rule Monitor Action, on page 1359

Procedure

Step 1 In the SSL policy editor, you have the following options:

- To add a new rule, click **Add Rule**.
- To edit an existing rule, click **Edit** ( ).

Step 2 Select a rule action from the **Action** drop-down list.

- To block encrypted traffic, select **Block**.
- To block encrypted traffic and reset the connection, select **Block with reset**.
- To decrypt incoming traffic, see Configuring a Decrypt - Known Key Action, on page 1362 for more information.
- To decrypt outgoing traffic, see Configuring a Decrypt - Resign Action, on page 1362 for more information.
- To log encrypted traffic, select **Monitor**.
- To not decrypt encrypted traffic, select **Do not decrypt**.

Step 3 Click **Add**.

What to do next

- Configure rule conditions, as described in Network-Based TLS/SSL Rule Conditions, on page 1366, User-Based TLS/SSL Rule Conditions, on page 1372, Reputation-Based TLS/SSL Rule Conditions, on page 1373, and Server Certificate-Based TLS/SSL Rule Conditions, on page 1379.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configuring a Decrypt - Resign Action

Before you begin

See TLS/SSL Decrypt - Resign Guidelines, on page 1347.

Procedure

Step 1 In the SSL rule editor, select **Decrypt - Resign** from the **Action** list.

Step 2 Select an internal CA certificate object from the list.

Step 3 Check **Replace Key Only** .

Always check the **Replace Key Only** check box when you set up a **Decrypt - Resign** rule action.

When a user browses to a web site that uses a *self-signed* certificate, the user sees a security warning in the web browser and is aware that they are communicating with an unsecure site.

When a user browses to a web site that uses a trusted certificate, the user does not see a security warning.

Step 4 Click **Add**.

Step 5 *Optional*. To use a Trusted CA certificate in your SSL policy so you can avoid **Invalid Issuer** in the SSL Certificate Status column in connection events, add the certificate to the policy:

a) In the SSL policy editor page, click **Trusted CA Certificates**.

b) Add the CA certificate corresponding to your known key to the SSL policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configuring a Decrypt - Known Key Action

Before you begin

See TLS/SSL Decrypt - Known Key Guidelines, on page 1349.

Procedure

Step 1 In the SSL rule editor, select **Decrypt - Known Key** from the **Action** drop-down list.

Step 2 Click the **Click to select decryption certs** field.

Step 3 Select one or more internal certificate objects in the **Available Certificates** list, then click **Add to Rule**.

Step 4 Click **OK**.

Step 5 Click **Add**.

Step 6 *Optional*. To use a Trusted CA certificate in your SSL policy so you can avoid **Invalid Issuer** in the SSL Certificate Status column in connection events, add the certificate to the policy:

a) In the SSL policy editor page, click **Trusted CA Certificates**.

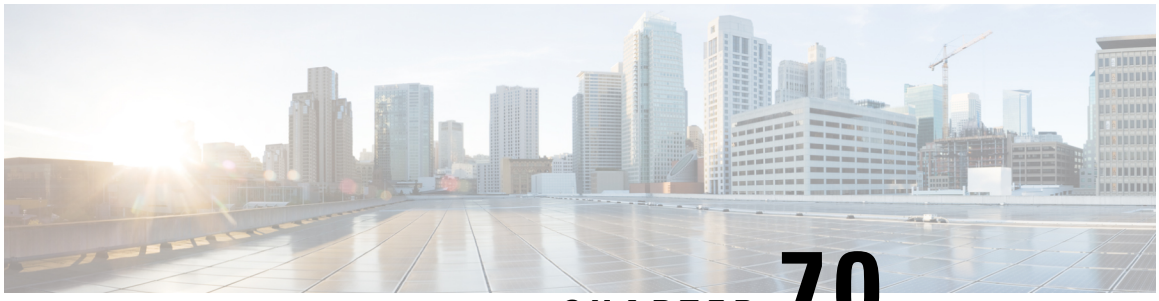
b) Add the CA certificate corresponding to your known key to the SSL policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

TLS/SSL Rules Management

The **Rules** page of the SSL policy editor allows you to add, edit, search, move, enable, disable, delete, and otherwise manage TLS/SSL rules in your policy.



CHAPTER 70

Decryption Tuning Using TLS/SSL Rules

The following topics provide an overview of how to configure TLS/SSL rule conditions:

- TLS/SSL Rule Conditions Overview, on page 1365
- Requirements and Prerequisites for Decryption Tuning, on page 1366
- Network-Based TLS/SSL Rule Conditions, on page 1366
- User-Based TLS/SSL Rule Conditions, on page 1372
- Reputation-Based TLS/SSL Rule Conditions, on page 1373
- Server Certificate-Based TLS/SSL Rule Conditions, on page 1379

TLS/SSL Rule Conditions Overview

A basic TLS/SSL rule applies its rule action to all encrypted traffic inspected by the device. To better control and decrypt encrypted traffic, you can configure rule conditions to handle and log specific types of traffic. Each TLS/SSL rule can contain 0, 1, or more rule conditions; a rule matches traffic only if the traffic matches every condition in that TLS/SSL rule.

Note When traffic matches a rule, the device applies the configured rule action to the traffic. When the connection ends, the device logs the traffic if configured to do so.

Each rule condition allows you to specify one or more properties of traffic you want to match against; these properties include details of:

- The flow of traffic, including the security zone through which it travels, IP address and port, country of origin or destination, and origin or destination VLAN.
- The user associated with a detected IP address.
- The traffic payload, including the application detected in the traffic.
- The connection encryption, including the TLS/SSL protocol version and cipher suite and server certificate used to encrypt the connection.
- The category and reputation of the URL specified in the server certificate's distinguished name..

Related Topics

Network-Based TLS/SSL Rule Conditions, on page 1366

User-Based TLS/SSL Rule Conditions, on page 1372
Reputation-Based URL Blocking in Encrypted Traffic, on page 1378
Server Certificate-Based TLS/SSL Rule Conditions, on page 1379
ClientHello Message Handling, on page 1307

Requirements and Prerequisites for Decryption Tuning

Model Support

Any except NGIPSv.

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Network-Based TLS/SSL Rule Conditions

TLS/SSL *rules* in *SSL policies* exert granular control over encrypted traffic logging and handling. Network-based conditions allow you to manage which encrypted traffic can traverse your network, using one or more of the following criteria:

- Zone conditions in TLS/SSL rules allow you to control encrypted traffic by its source and destination security zones. A *security zone* is a grouping of one or more interfaces, which might be located across multiple devices.
- Network conditions in TLS/SSL rules allow you to control and decrypt encrypted traffic by its source and destination IP address. You can either explicitly specify the source and destination IP addresses for the encrypted traffic you want to control, or use the geolocation feature, which associates IP addresses with geographical locations, to control encrypted traffic based on its source or destination country or continent.
- VLAN conditions in TLS/SSL rules allow you to control VLAN-tagged traffic. The system uses the innermost VLAN tag to identify a packet by VLAN.
- Port conditions in TLS/SSL rules allow you to control encrypted traffic by its source and destination TCP port.

You can combine network-based conditions with each other and with other types of conditions to create a TLS/SSL rule. These TLS/SSL rules can be simple or complex, matching and inspecting traffic using multiple conditions.

Related Topics

Firepower System IP Address Conventions, on page 15

Network Zone TLS/SSL Rule Conditions

You can add a maximum of 50 zones to each of the **Sources Zones** and **Destination Zones** in a single zone condition:

- To match encrypted traffic *leaving* the device from an interface in the zone, add that zone to the **Destination Zones**.

  Because devices deployed passively do not transmit traffic, you cannot use a zone comprised of passive interfaces in a **Destination Zone** condition.

- To match encrypted traffic *entering* the device from an interface in the zone, add that zone to the **Source Zones**.

If you add both source and destination zone conditions to a rule, matching traffic must originate from one of the specified source zones *and* egress through one of the destination zones.

Note that just as all interfaces in a zone must be of the same type (all inline, all passive, all switched, or all routed), all zones used in a zone condition for a TLS/SSL rule must be of the same type. That is, you cannot write a single rule that matches encrypted traffic to or from zones of different types.

Warning icons indicate invalid configurations, such as zones that contain no interfaces. For details, hover your pointer over the icon.

Controlling Encrypted Traffic by Network Zone

Procedure

Step 1 In the SSL rule editor, select the Zones tab.

Step 2 Find the zones you want to add from the **Available Zones**. To search for zones to add, click the **Search by name** prompt above the **Available Zones** list, then type a zone name. The list updates as you type to display matching zones.

Step 3 Click to select a zone. To select all zones, right-click and then select **Select All**.

Step 4 Click **Add to Source** or **Add to Destination**.

Tip You can also drag and drop selected zones.

Step 5 Save or continue editing the rule.

Example

For example, you could deploy additional identically configured devices—managed by the same Firepower Management Center—to protect similar resources in several different locations. Like the first device, each of these devices protects the assets in an **Internal** security zone.

Note You are not required to group all internal (or external) interfaces into a single zone. Choose the grouping that makes sense for your deployment and security policies.

In this deployment, you may decide that although you want these hosts to have unrestricted access to the Internet, you nevertheless want to protect them by decrypting and inspecting incoming encrypted traffic.

To accomplish this, configure a TLS/SSL rule with a zone condition where the **Destination Zone** is set to **Internal**. This simple SSL rule matches traffic that leaves the device from any interface in the Internal zone.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Interface Objects: Interface Groups and Security Zones, on page 410

Network or Geolocation TLS/SSL Rule Conditions

When you build a network-based TLS/SSL rule condition, you can manually specify IP address and geographical locations. Alternately, you can configure network conditions with network and geolocation *objects*, which are reusable and associate a name with one or more IP addresses, address blocks, countries, continents, and so on.

Note If you want to write rules to control traffic by geographical location, to ensure you are using up-to-date geolocation data to filter your traffic, Cisco **strongly** recommends you regularly update the geolocation database (GeoDB) on your Firepower Management Center.

You can add a maximum of 50 items to each of the **Source Networks** and **Destination Networks** in a single network condition, and you can mix network and geolocation-based configurations:

- To match encrypted traffic *from* an IP address or geographical location, configure the **Source Networks**.
- To match encrypted traffic *to* an IP address or geographical location, configure the **Destination Networks**.

If you add both source and destination network conditions to a rule, matching encrypted traffic must originate from one of the specified IP addresses *and* be destined for one of the destination IP addresses.

When building a network condition, warning icons indicate invalid configurations. For details, hover your pointer over the icon.

Related Topics

Firepower System IP Address Conventions, on page 15

Controlling Encrypted Traffic by Network or Geolocation

Before you begin

- Update the geolocation database (GeoDB) on your Firepower Management Center as described in Update the Geolocation Database, on page 141.

Procedure

Step 1 In the SSL rule editor, select the Networks tab.

Step 2 Find the networks you want to add from the **Available Networks**, as follows:

- Click the Networks tab to display network objects and groups to add; click the Geolocation tab to display geolocation objects.
- To add a network object on the fly, which you can then add to the condition, click the add icon ( ) above the Available Networks list.
- To search for network or geolocation objects to add, select the appropriate tab, click the **Search by name or value** prompt above the **Available Networks** list, then type an object name or the value of one of the object's components. The list updates as you type to display matching objects.

Step 3 To select an object, click it. To select all objects, right-click and then select **Select All**.

Step 4 Click **Add to Source** or **Add to Destination**.

Tip You can also drag and drop selected objects.

Step 5 Add any source or destination IP addresses or address blocks that you want to specify manually. Click the **Enter an IP address** prompt below the **Source Networks** or **Destination Networks** list; then type an IP address or address block and click **Add**.

Step 6 Save or continue editing the rule.

Example

The following graphic shows the network condition for a TLS/SSL rule that blocks encrypted connections originating from your internal network and attempting to access resources either in the Cayman Islands or an offshore holding corporation server at 182.16.0.3.



The example manually specifies the offshore holding corporation's server IP address, and uses a system-provided Cayman Islands geolocation object to represent Cayman Island IP addresses.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Network Objects, on page 402
Firepower System IP Address Conventions, on page 15

VLAN TLS/SSL Rule Conditions

When you build a VLAN-based TLS/SSL rule condition, you can manually specify a VLAN tag from 1 to 4094. Alternately, you can configure VLAN conditions with VLAN tag *objects*, which are reusable and associate a name with one or more VLAN tags.

Tip After you create a VLAN tag object, you can use it not only to build TLS/SSL rules, but also to represent VLAN tags in various other places in the system's web interface. You can create VLAN tag objects either using the object manager or on-the-fly while you are configuring access control rules.

You can add a maximum of 50 items to the **Selected VLAN Tags** in a single VLAN tag condition. When building a VLAN tag condition, warning icons indicate invalid configurations. For details, hover your pointer over the icon.

Controlling Encrypted VLAN Traffic

Procedure

Step 1 In the SSL rule editor, select the VLAN Tags tab.

Step 2 Find the VLANs you want to add from the **Available VLAN Tags**, as follows:

- To add a VLAN tag object on the fly, which you can then add to the condition, click the add icon ( ) above the Available VLAN Tags list.
- To search for VLAN tag objects and groups to add, click the **Search by name or value** prompt above the **Available VLAN Tags** list, then type either the name of the object, or the value of a VLAN tag in the object. The list updates as you type to display matching objects.

Step 3 To select an object, click it. To select all objects, right-click and then select **Select All**.

Step 4 Click **Add to Rule**.

Tip You can also drag and drop selected objects.

Step 5 Add any VLAN tags that you want to specify manually. Click the **Enter a VLAN Tag** prompt below the **Selected VLAN Tags** list; then type a VLAN tag or range and click **Add**. You can specify any VLAN tag from 1 to 4094; use a hyphen to specify a range of VLAN tags.

Step 6 Save or continue editing the rule.

Example

The following graphic shows a VLAN tag condition for an SSL rule that matches encrypted traffic on public-facing VLANs (represented by a VLAN tag object group), as well as the manually added VLAN 42.



What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

VLAN Tag Objects, on page 406

Port TLS/SSL Rule Conditions

When you build a port-based TLS/SSL rule condition, you can manually specify TCP ports. Alternately, you can configure port conditions with port *objects*, which are reusable and associate a name with one or more ports.

You can add a maximum of 50 items to each of the **Selected Source Ports** and **Selected Destination Ports** lists in a single network condition:

- To match encrypted traffic *from* a TCP port, configure the **Selected Source Ports**.
- To match encrypted traffic *to* a TCP port, configure the **Selected Destination Ports**.
- To match encrypted traffic both originating from TCP **Selected Source Ports***and* destined for TCP **Selected Destination Ports**, configure both.

You can only configure the **Selected Source Ports** and **Selected Destination Ports** lists with TCP ports. Port objects containing non-TCP ports are greyed out in the **Available Ports** list.

When building a port condition, warning icons indicate invalid configurations. For example, you can use the object manager to edit in-use port objects so that the rules that use those object groups become invalid. For details, hover your pointer over the icon.

Controlling Encrypted Traffic by Port

Procedure

Step 1 In the SSL rule editor, select the Ports tab.

Step 2 Find the TCP ports you want to add from the **Available Ports**, as follows:

- To add a TCP port object on the fly, which you can then add to the condition, click the add icon ( ) above the Available Ports list.
- To search for TCP-based port objects and groups to add, click the **Search by name or value** prompt above the **Available Ports** list, then type either the name of the object, or the value of a port in the object. The list updates as you type to display matching objects. For example, if you type 443, the Firepower Management Center displays the system-provided HTTPS port object.

Step 3 To select a TCP-based port object, click it. To select all TCP-based port objects, right-click and then select **Select All**. If the object includes non-TCP-based ports, you cannot add it to your port condition.

Step 4 Click **Add to Source** or **Add to Destination**.

Tip You can also drag and drop selected objects.

Step 5 Enter a **Port** under the **Selected Source Ports** or **Selected Destination Ports** list to manually specify source or destination ports. You can specify a single port with a value from 0 to 65535.

Step 6 Click **Add**.

Note The Firepower Management Center will not add a port to a rule condition that results in an invalid configuration.

Step 7 Save or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Port Objects, on page 404

User-Based TLS/SSL Rule Conditions

You can configure TLS/SSL rules to match traffic based on realm, group, or user. Realm, group, and user conditions in TLS/SSL rules allow you perform *user control* to manage which traffic can traverse your network by associating authoritative users with IP addresses.

For traffic to match a TLS/SSL rule with a user condition, the IP address of either the source or destination host in the monitored session must be associated with a logged in authoritative user. You can control traffic based on realms, individual users, or the groups those users belong to.

Controlling Encrypted Traffic Based on User

Before you begin

- Configure one or more authoritative user identity sources as described in User Identity Sources, on page 1891.
- Configure a realm as described in Create a Realm, on page 1949.

Procedure

Step 1 In the SSL rule editor, select Users.

Step 2 Search by name or value above the **Available Realms** list and select a realm.

Step 3 Search by name or value above the **Available Users** list and select a user or group.

Step 4 Click **Add to Rule**.

Tip You can also drag and drop selected users and groups.

Step 5 Save or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Reputation-Based TLS/SSL Rule Conditions

Reputation-based conditions in TLS/SSL rules allow you to manage which encrypted traffic can traverse your network, by contextualizing your network traffic and limiting it where appropriate. SSL rules govern the following types of reputation-based control:

- Application conditions allow you to perform *application control.* When the system analyzes encrypted IP traffic, it can identify and classify commonly used encrypted applications on your network prior to decrypting the encrypted session. The system uses this discovery-based *application awareness* feature to allow you to control encrypted application traffic on your network.

  I a single TLS/SSL rule, you can select individual applications, including custom applications. You can use system-provided *application filters*, which are named sets of applications organized according to its basic characteristics: type, risk, business relevance, and categories.

- URL conditions allow you to control web traffic based on a websites' assigned category and reputation.

Selected Applications and Filters in TLS/SSL Rules

Cisco frequently updates and adds additional detectors via system and vulnerability database (VDB) updates. You can also create your own detectors and assign characteristics (risk, relevance, and so on) to the applications

they detect. By using filters based on application characteristics, you can ensure that the system uses the most up-to-date detectors to monitor application traffic.

For traffic to match a TLS/SSL rule with an application condition, the traffic must match one of the filters or applications that you add to a **Selected Applications and Filters** list.



Note When you filter application traffic using access control rules, you can use application tags as a criterion. to filter. However, you cannot use application tags to filter encrypted traffic because there is no benefit. All applications that the system can detect in encrypted traffic are tagged **SSL Protocol**; applications without this tag can only be detected in unencrypted or decrypted traffic.

In a single application condition, you can add a maximum of 50 items to the **Selected Applications and Filters** list. Each of the following counts as an item:

- One or more filters from the **Application Filters** list, individually or in custom combination. This item represents set of applications, grouped by characteristic.
- A filter created by saving search of the applications in the **Available Applications** list. This item represents a set of applications, grouped by substring match.
- An individual application from the **Available Applications** list.

In the web interface, filters added to a condition are listed above and separately from individually added applications.

Note that when you deploy an SSL policy, for each rule with an application condition, the system generates a list of unique applications to match. In other words, you may use overlapping filters and individually specified applications to ensure complete coverage.

Application Filters in TLS/SSL Rules

When building an application condition in a TLS/SSL rule, use the **Application Filters** list to create a set of applications, grouped by characteristic, whose traffic you want to match.

For your convenience, the system characterizes each application by type, risk, business relevance, category, and tag. You can use these criteria as filters or create custom combinations of filters to perform application control.

Note that the mechanism for filtering applications in a TLS/SSL rule is the same as that for creating reusable, custom application filters using the object manager. You can also save many filters you create on-the-fly in access control rules as new, reusable filters. You cannot save a filter that includes another user-created filter because you cannot nest user-created filters.

Understanding How Filters Are Combined

When you select filters, singly or in combination, the **Available Applications** list updates to display only the applications that meet your criteria. You can select system-provided filters in combination, but not custom filters.

The system links multiple filters of the same filter type with an OR operation. For example, if you select the Medium and High filters under the Risks type, the resulting filter is:

```
Risk: Medium OR High
```

If the Medium filter contained 110 applications and the High filter contained 82 applications, the system displays all 192 applications in the **Available Applications** list.

The system links different types of filters with an AND operation. For example, if you select the Medium and High filters under the Risks type, and the Medium and High filters under the Business Relevance type, the resulting filter is:

```
Risk: Medium OR High
AND
Business Relevance: Medium OR High
```

In this case, the system displays only those applications that are included in both the Medium or High Risk type AND the Medium or High Business Relevance type.

Finding and Selecting Filters

To select filters, click the arrow next to a filter type to expand it, then select or clear the check box next to each filter whose applications you want to display or hide. You can also right-click a Cisco-provided filter type (**Risks**, **Business Relevance**, **Types**, or **Categories**) and select **Check All** or **Uncheck All**.

To search for filters, click the **Search by name** prompt above the **Available Filters** list, then type a name. The list updates as you type to display matching filters.

After you are done selecting filters, use the **Available Applications** list to add those filters to the rule.

Related Topics

Application Filters, on page 406

Available Applications in TLS/SSL Rules

When building an application condition in a TLS/SSL rule, use the **Available Applications** list to select the applications whose traffic you want to match.

Browsing the List of Applications

When you first start to build the condition the list is unconstrained, and displays every application the system detects, 100 at a time:

- To page through the applications, click the arrows underneath the list.
- To display a pop-up window with summary information about the application's characteristics, as well as Internet search links that you can follow, click **Information** () next to an application.

Finding Applications to Match

To help you find the applications you want to match, you can constrain the **Available Applications** list in the following ways:

- To search for applications, click the **Search by name** prompt above the list, then type a name. The list updates as you type to display matching applications.
- To constrain the applications by applying a filter, use the **Application Filters** list. The **Available Applications** list updates as you apply filters.

Once constrained, an **All apps matching the filter** option appears at the top of the **Available Applications** list.



Note If you select one or more filters in the Application Filters list and also search the **Available Applications** list, your selections and the search-filtered **Available Applications** list are combined using an AND operation. That is, the **All apps matching the filter** condition includes all the individual conditions currently displayed in the **Available Applications** list as well as the search string entered above the **Available Applications** list.

Selecting Single Applications to Match in a Condition

After you find an application you want to match, click to select it. To select all applications in the current constrained view, right-click and select **Select All**.

In a single application condition, you can match a maximum of 50 applications by selecting them individually; to add more than 50 you must either create multiple TLS/SSL rules or use filters to group applications.

Selecting All Applications Matching a Filter for a Condition

Once constrained by either searching or using the filters in the **Application Filters** list, the **All apps matching the filter** option appears at the top of the **Available Applications** list.

This option allows you to add the entire set of applications in the constrained **Available Applications** list to the **Selected Applications and Filters** list, at once. In contrast to adding applications individually, adding this set of applications counts as only one item against the maximum of 50, regardless of the number of individual application that comprise it.

When you build an application condition this way, the name of the filter you add to the **Selected Applications and Filters** list is a concatenation of the filter types represented in the filter plus the names of up to three filters for each type. More than three filters of the same type are followed by an ellipsis (...). For example, the following filter name includes two filters under the Risks type and four under Business Relevance:

```
Risks: Medium, High Business Relevance: Low, Medium, High,...
```

Filter types that are not represented in a filter you add with **All apps matching the filter** are not included in the name of the filter you add. The instructional text that is displayed when you hover your pointer over the filter name in the **Selected Applications and Filters** list indicates that these filter types are set to *any*; that is, these filter types do not constrain the filter, so any value is allowed for these.

You can add multiple instances of **All apps matching the filter** to an application condition, with each instance counting as a separate item in the **Selected Applications and Filters** list. For example, you could add all high risk applications as one item, clear your selections, then add all low business relevance applications as another item. This application condition matches applications that are high risk OR have low business relevance.

Application-Based TLS/SSL Rule Condition Requirements

For encrypted traffic to match a TLS/SSL rule with an application condition, the traffic must match one of the filters or applications that you add to a **Selected Applications and Filters** list.

You can add a maximum of 50 items per condition, and filters added to a condition are listed above and separately from individually added applications. When building an application condition, warning icons indicate invalid configurations. For details, hover your pointer over the icon.

Adding an Application Condition to a TLS/SSL Rule

For Classic device models, you must have the Control license to use this feature.

Procedure

Step 1 In the SSL rule editor, select the Applications tab.

Step 2 If you want to constrain the list of applications displayed in the **Available Applications** list, you must select one or more filters in the **Application Filters** list. For more information, see Application Filters in TLS/SSL Rules, on page 1374.

Step 3 Find and select the applications you want to add from the **Available Applications** list. You can search for and select individual applications, or, when the list is constrained, **All apps matching the filter**. For more information, see Available Applications in TLS/SSL Rules, on page 1375.

Step 4 Click **Add to Rule**.

Tip Click **Clear All Filters** to clear your existing selections. You can also drag and drop selected applications and filters.

Step 5 Save or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Limitations to Encrypted Application Control

Encrypted Application Identification

The system can identify unencrypted applications that become encrypted using StartTLS. This includes such applications as SMTPS, POPS, FTPS, TelnetS, and IMAPS. In addition, it can identify certain encrypted applications based on the Server Name Indication in the TLS ClientHello message, or the server certificate subject distinguished name value.

Speed of Application Identification

The system cannot perform application control on encrypted traffic before:

- An encrypted connection is established between a client and server, and
- The system identifies the application in the encrypted session

This identification occurs after the server certificate exchange. If traffic exchanged during the TLS/SSL handshake matches all other conditions in a TLS/SSL rule containing an application condition but the identification is not complete, the SSL policy allows the packet to pass. This behavior allows the handshake to complete so that applications can be identified. For your convenience, affected rules are marked with an **Information** ().

After the system completes its identification, the system applies the TLS/SSL rule action to the remaining session traffic that matches its application condition.

Automatically Enabling Application Detectors

At least one detector must be enabled for each application rule condition in the policy. If no detector is enabled for an application, the system automatically enables all system-provided detectors for the application; if none exist, the system enables the most recently modified user-defined detector for the application.

Related Topics

Activating and Deactivating Detectors, on page 1889

Reputation-Based URL Blocking in Encrypted Traffic

With a URL Filtering license, URL conditions in TLS/SSL rules can control access to encrypted websites, based on the category and reputation of the requested URLs. For detailed information, see URL Conditions (URL Filtering), on page 384.

Tip URL conditions in TLS/SSL rules do not support manual URL filtering. Instead, use a distinguished name condition matching on the subject common name.

Block Encrypted Traffic Based on URL Reputation

You must have the URL filtering license to use this feature.

Procedure

Step 1 In the SSL rule editor, select the Category tab.

Step 2 Find the categories of URL you want to add from the **Categories** list. To match encrypted web traffic regardless of category, select **Any** category. To search for categories to add, click the **Search by name or value** prompt above the **Categories** list, then type the category name. The list updates as you type to display matching categories.

Step 3 To select a category, click it.

Tip Although you can right-click and **Select All** categories, adding all categories this way exceeds the 50-item maximum for a TLS/SSL rule. Instead, use **Any**.

Step 4 If you want to qualify your category selections, you must click a reputation level from the **Reputations** list. You can only select one reputation level. If you do not specify a reputation level, the system defaults to **Any**, meaning all levels.

- If the rule blocks web access or decrypts traffic (the rule action is **Block**, **Block with reset**, **Decrypt - Known Key**, **Decrypt - Resign**, or **Monitor**) selecting a reputation level also selects all reputations more severe than that level. For example, if you configure a rule to block **Suspicious sites** (level 2), it also automatically blocks **High Risk** (level 1) sites.
- If the rule allows web access, subject to access control (the rule action is **Do not decrypt**), selecting a reputation level also selects all reputations less severe than that level. For example, if you configure a rule to allow **Benign sites** (level 4), it also automatically allows **Well known** (level 5) sites.

Note If you change the rule action for a rule, the system automatically changes the reputation levels in URL conditions according to the above points.

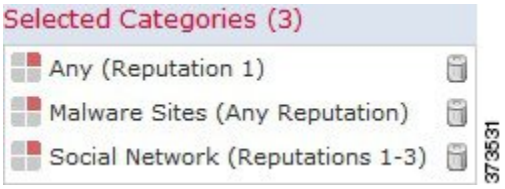
Step 5 Click **Add to Rule** to add the selected items to the **Selected Categories** list.

Tip You can also drag and drop selected items.

Step 6 Save or continue editing the rule.

Example

The following graphic shows the URL condition for an example access control rule that blocks: all malware sites, all high-risk sites, and all non-benign social networking sites.



The following table summarizes how you build the condition shown in the graphic above.

Table 96: Example: Building A URL Condition

To block...	Select this Category or URL Object...	And this Reputation...
malware sites, regardless of reputation	Malware Sites	Any
any URL with a high risk (level 1)	Any	1 - High Risk
social networking sites with a risk greater than benign (levels 1 through 3)	Social Network	3 - Benign sites with security risks

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Server Certificate-Based TLS/SSL Rule Conditions

TLS/SSL rules can handle and decrypt encrypted traffic based on server certificate characteristics. You can configure TLS/SSL rules based on the following server certificate attributes:

- Distinguished name conditions allow you to handle and inspect encrypted traffic based on the CA that issued a server certificate, or the certificate holder. Based on the issuer distinguished name, you can handle traffic based on the CA that issued a site's server certificate.

- Certificate conditions in TLS/SSL rules allow you to handle and inspect encrypted traffic based on the server certificate used to encrypt that traffic. You can configure a condition with one or more certificates; traffic matches the rule if the certificate matches any of the condition's certificates.
- Certificate status conditions in TLS/SSL rules allow you to handle and inspect encrypted traffic based on the status of the server certificate used to encrypt the traffic, including whether a certificate is valid, revoked, expired, not yet valid, self-signed, signed by a trusted CA, whether the Certificate Revocation List (CRL) is valid; whether the Server Name Indication (SNI) in the certificate matches the server in the request.
- Cipher suite conditions in TLS/SSL rules allow you to handle and inspect encrypted traffic based on the cipher suite used to negotiate the encrypted session.
- Session conditions in TLS/SSL rules allow you to inspect encrypted traffic based on the SSL or TLS version used to encrypt the traffic.

To detect multiple cipher suites in a rule, the certificate issuer, or the certificate holder, you can create reusable cipher suite list and distinguished name objects and add them to your rule. To detect the server certificate and certain certificate statuses, you must create external certificate and external CA objects for the rule.

Certificate Distinguished Name TLS/SSL Rule Conditions

When configuring the rule condition, you can manually specify a literal value, reference a distinguished name object, or reference a distinguished name group containing multiple objects.



Note You cannot configure a distinguished name condition if you also choose the **Decrypt - Known Key** action. Because that action requires you to choose a server certificate to decrypt traffic, the certificate already matches the traffic.

You can match against multiple subject and issuer distinguished names in a single certificate status rule condition; only one common or distinguished name needs to match to match the rule.

If you add a distinguished name manually, it can contain the common name attribute (**CN**). If you add a common name without `CN=`, the system prepends `CN=` before saving the object.

You can also add a distinguished name with one each of the following attributes, separated by commas: **C, CN, O, OU.**

In a single DN condition, you can add a maximum of 50 literal values and distinguished name objects to the **Subject DNs**, and 50 literal values and distinguished name objects to the **Issuer DNs**.

The system-provided DN object group, Cisco-Undecryptable-Sites, contains websites whose traffic the system cannot decrypt. You can add this group to a DN condition to block or not decrypt traffic to or from these websites, without wasting system resources attempting to decrypt that traffic. You can modify individual entries in the group. You cannot delete the group. System updates can modify the entries on this list, but the system preserves user changes.

The first time the system detects an encrypted session to a new server, DN data is not available for ClientHello processing, which can result in an undecrypted first session. After the initial session, the managed device caches data from the server Certificate message. For subsequent connections from the same client, the system can match the ClientHello message conclusively to rules with DN conditions and process the message to maximize decryption potential.

Controlling Encrypted Traffic by Certificate Distinguished Name

Procedure

Step 1 In the SSL rule editor, select DN.

Step 2 Find the distinguished names you want to add from the **Available DNs**, as follows:

- To add a distinguished name object on the fly, which you can then add to the condition, click **Add** ( ) above the **Available DNs** list.
- To search for distinguished name objects and groups to add, click the **Search by name or value** prompt above the **Available DNs** list, then type either the name of the object, or a value in the object. The list updates as you type to display matching objects.

Step 3 To select an object, click it. To select all objects, right-click and then select **Select All**.

Step 4 Click **Add to Subject** or **Add to Issuer**.

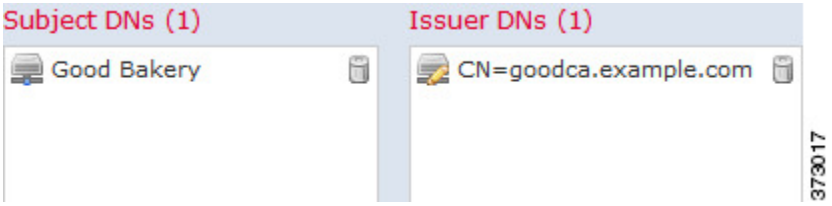
Tip You can also drag and drop selected objects.

Step 5 Add any literal common names or distinguished names that you want to specify manually. Click the **Enter DN or CN** prompt below the **Subject DNs** or **Issuer DNs** list; then type a common name or distinguished name and click **Add**.

Step 6 Add or continue editing the rule.

Example

The following figure shows a distinguished name rule condition searching for certificates issued to goodbakery.example.com or issued by goodca.example.com. Traffic encrypted with these certificates is allowed, subject to access control.



The following figure shows a distinguished name rule condition searching for certificates issued to badbakery.example.com and associated domains, or certificates issued by badca.example.com. Traffic encrypted with these certificates is decrypted using a re-signed certificate.



What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Distinguished Name Objects, on page 444

Certificate TLS/SSL Rule Conditions

When you build a certificate-based TLS/SSL rule condition, you can upload a server certificate; you save the certificate as an external certificate *object*, which is reusable and associates a name with a server certificate. Alternately, you can configure certificate conditions with existing external certificate objects and object groups.

You can search the **Available Certificates** field in the rule condition based for external certificate objects and object groups based on the following certificate distinguished name characteristics:

- Subject or issuer common name (CN)
- Subject or issuer organization (O)
- Subject or issuer organizational unit (OU)

You can choose to match against multiple certificates in a single certificate rule condition; if the certificate used to encrypt the traffic matches any of the uploaded certificates, the encrypted traffic matches the rule.

You can add a maximum of 50 external certificate objects and external certificate object groups to the **Selected Certificates** in a single certificate condition.

Note the following:

- You cannot configure a certificate condition if you also select the **Decrypt - Known Key** action. Because that action requires you to select a server certificate to decrypt traffic, the implication is that the certificate already matches the traffic.
- If you configure a certificate condition with an external certificate object, any cipher suites you add to a cipher suite condition, or internal CA objects you associate with the **Decrypt - Resign** action, must match the external certificate's signature algorithm type. For example, if your rule's certificate condition references an EC-based server certificate, any cipher suites you add, or CA certificates you associate with the **Decrypt - Resign** action, must also be EC-based. If you mismatch signature algorithm types in this case, the policy editor displays a warning next to the rule.
- The first time the system detects an encrypted session to a new server, certificate data is not available for ClientHello processing, which can result in an undecrypted first session. After the initial session, the managed device caches data from the server Certificate message. For subsequent connections from the same client, the system can match the ClientHello message conclusively to rules with certificate conditions and process the message to maximize decryption potential.

Controlling Encrypted Traffic by Certificate

Procedure

Step 1 In the SSL rule editor, select Certificate.

Step 2 Find the server certificates you want to add from the **Available Certificates**, as follows;

- To add an external certificate object on the fly, which you can then add to the condition, click **Add** ( ) above the **Available Certificates** list.
- To search for certificate objects and groups to add, click the **Search by name or value** prompt above the **Available Certificates** list, then type either the name of the object, or a value in the object. The list updates as you type to display matching objects.

Step 3 To select an object, click it. To select all objects, right-click and then select **Select All**.

Step 4 Click **Add to Rule**.

Tip You can also drag and drop selected objects.

Step 5 Add or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

External Certificate Objects, on page 454

Certificate Status TLS/SSL Rule Conditions

For each certificate status TLS/SSL rule condition you configure, you can match traffic against the presence or absence of a given status. You can select several statuses in one rule condition; if the certificate matches any of the selected statuses, the rule matches the traffic.

You can choose to match against the presence or absence of multiple certificate statuses in a single certificate status rule condition; the certificate needs to match only one of the criteria to match the rule.

You should consider, when setting this parameter, whether you're configuring a decrypt rule or a block rule. Typically, you should click **Yes** for a block rule and **No** for a decrypt rule. Examples:

- If you're configuring a **Decrypt - Resign** rule, the default behavior is to decrypt traffic with an expired certificate. To change that behavior, click **No** for **Expired** so traffic with an expired certificate is not decrypted and resigned.
- If you're configuring a **Block** rule, the default behavior is to allow traffic with an expired certificate. To change that behavior click **Yes** for **Expired** so traffic with an expired certificate is blocked.

The following table describes how the system evaluates encrypted traffic based on the encrypting server certificate's status.

Table 97: Certificate Status Rule Condition Criteria

Status Check	Status Set to Yes	Status Set to No
Revoked	The policy trusts the CA that issued the server certificate, and the CA certificate uploaded to the policy contains a CRL that revokes the server certificate.	The policy trusts the CA that issued the ser certificate, and the CA certificate uploaded policy does not contain a CRL that revoke certificate.
Self-signed	The detected server certificate contains the same subject and issuer distinguished name.	The detected server certificate contains dif subject and issuer distinguished names.
Valid	All of the following are true: • The policy trusts the CA that issued the certificate. • The signature is valid. • The issuer is valid. • None of the policy's trusted CAs revoked the certificate. • The current date is between the certificate Valid From and Valid To date.	At least one of the following is true: • The policy does not trust the CA that certificate. • The signature is invalid. • The issuer is invalid. • A trusted CA in the policy revoked th certificate. • The current date is before the certifica From date. • The current date is after the certificate date.
Invalid signature	The certificate's signature cannot be properly validated against the certificate's content.	The certificate's signature is properly validat the certificate's content.
Invalid issuer	The issuer CA certificate is not stored in the policy's list of trusted CA certificates.	The issuer CA certificate is stored in the po of trusted CA certificates.
Expired	The current date is after the certificate Valid To date.	The current date is before or on the certific To date.
Not yet valid	The current date is before the certificate Valid From date.	The current date is after or on the certificat From date.

Status Check	Status Set to Yes	Status Set to No
Invalid certificate	The certificate is not valid. At least one of the following is true: • Invalid or inconsistent certificate extension; that is, a certificate extension had an invalid value (for example, an incorrect encoding) or some value inconsistent with other extensions. • The certificate cannot be used for the specified purpose. • The Basic Constraints path length parameter has been exceeded. For more information, see RFC 5280, section 4.2.1.9. • The certificate's value for Not Before or Not After is invalid. These dates can be encoded as UTCTime or GeneralizedTime For more information, see RFC 5280 section 4.1.2.5. • The format of the name constraint is not recognized; for example, an email address format of a form not mentioned in RFC 5280, section 4.2.1.10. This could be caused by an improper extension or some new feature not currently supported. An unsupported name constraint type was encountered. OpenSSL currently supports only directory name, DNS name, email, and URI types. • The root certificate authority is not trusted for the specified purpose. • The root certificate authority rejects the specified purpose.	The certificate is valid. All of the follo • Valid certificate extension. • The certificate can be used for the purpose. • Valid Basic Constraints path lengt • Valid values for Not Before and N • Valid name constraint. • The root certificate is trusted for t purpose. • The root certificate accepts the spec

Status Check	Status Set to Yes	Status Set to No
Invalid CRL	The Certificate Revocation List (CRL) digital signature is not valid. At least one of the following is true: • The value of the CRL's Next Update or Last Update field is invalid. • The CRL is not yet valid. • The CRL has expired. • An error occurred when attempting to verify the CRL path. This error occurs only if extended CRL checking is enabled. • CRL could not be found. • The only CRLs that could be found did not match the scope of the certificate.	The CRL is valid. All of the following are • Next Update and Last Update fields a • The CRL's date is valid. • The path is valid. • The CRL was found. • The CRL matches the certificate's sco
Server mismatch	The server name does not match the server's Server Name Indication (SNI) name, which could indicate an attempt to spoof the server name.	The server name matches the SNI name of to which the client is requesting access.

Note that even though a certificate might match more than one status, the rule causes an action to be taken on the traffic only once.

Checking whether a CA issued or revoked a certificate requires uploading root and intermediate CA certificates and associated CRLs as objects. You then add these trusted CA objects to an SSL policy's list of trusted CA certificates.

Trusting External Certificate Authorities

You can trust CAs by adding root and intermediate CA certificates to your SSL policy, then use these trusted CAs to verify server certificates used to encrypt traffic.

If a trusted CA certificate contains an uploaded certificate revocation list (CRL), you can also verify whether a trusted CA revoked the encryption certificate.

Procedure

Step 1 In the SSL rule editor, select **Trusted CA Certificates**.

Step 2 Find the trusted CAs you want to add from the **Available Trusted CAs**, as follows:

- To add a trusted CA object on the fly, which you can then add to the condition, click **Add** ( ) above the **Available Trusted CAs** list.
- To search for trusted CA objects and groups to add, click the **Search by name or value** prompt above the **Available Trusted CAs** list, then type either the name of the object, or a value in the object. The list updates as you type to display matching objects.

Step 3 To select an object, click it. To select all objects, right-click and then select **Select All**.

Step 4 Click **Add to Rule**.

Tip You can also drag and drop selected objects.

Step 5 Add or continue editing the rule.

What to do next

- Add a certificate status TLS/SSL rule condition to your SSL rule. See Matching Traffic on Certificate Status, on page 1387 for more information.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Trusted Certificate Authority Objects, on page 452

Trusted External Certificate Authority Configuration

Verified server certificates include certificates signed by trusted CAs. After you add trusted CA certificates to the SSL policy, you can configure a TLS/SSL rule with certificate status conditions to match against this traffic.

Tip Upload all certificates within a root CA's chain of trust to the list of trusted CA certificates, including the root CA certificate and all intermediate CA certificates. Otherwise, it is more difficult to detect trusted certificates issued by intermediate CAs. Also, if you configure certificate status conditions to trust traffic based on the root issuer CA, all traffic within a trusted CA's chain of trust can be allowed without decryption, rather than unnecessarily decrypting it.

When you create an SSL policy, the system populates the Trusted CA Certificates tab with a default Trusted CA object group, Cisco Trusted Authorities.

You can modify individual entries in the group, and choose whether to include this group in your SSL policy. You cannot delete the group. System updates can modify the entries on this list, but user changes are preserved.

Matching Traffic on Certificate Status

Before you begin

- Add a trusted CA object or group to your SSL policy. See Trusting External Certificate Authorities, on page 1386 for more information.

Procedure

Step 1 In the Firepower Management Center, choose **Policies** > **Access Control** > **SSL**.

Step 2 Add a new policy or edit an existing policy.

Step 3 Add a new TLS/SSL rule or edit an existing rule.

Step 4 In the Add Rule or Editing Rule dialog box, choose **Cert Status**.

Step 5 For each certificate status, you have the following options:

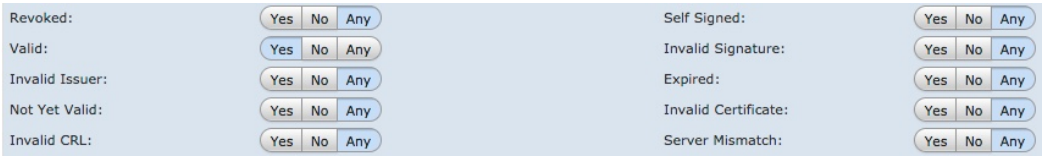
- Choose **Yes** to match against the presence of that certificate status.
- Choose **No** to match against the absence of that certificate status.
- Choose **Any** to skip the condition when matching the rule. In other words, choosing **Any** means the rule matches whether the certificate status is present or absent.

Step 6 Add or continue editing the rule.

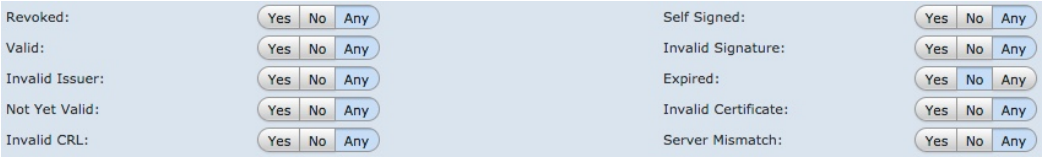
Example

The organization trusts the Verified Authority certificate authority. The organization does not trust the Spammer Authority certificate authority. The system administrator uploads the Verified Authority certificate and an intermediate CA certificate issued by Verified Authority to the system. Because Verified Authority revoked a certificate it previously issued, the system administrator uploads the CRL that Verified Authority provided.

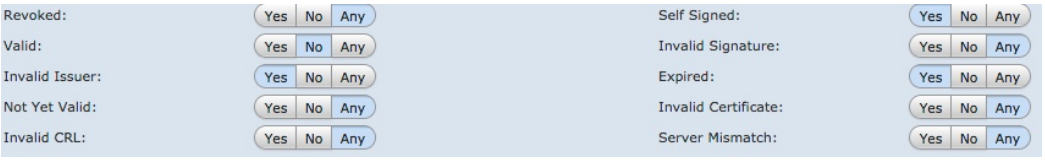
The following figure shows a certificate status rule condition checking for valid certificates, those issued by a Verified Authority, are not on the CRL, and still within the Valid From and Valid To date. Because of the configuration, traffic encrypted with these certificates is not decrypted and inspected with access control.



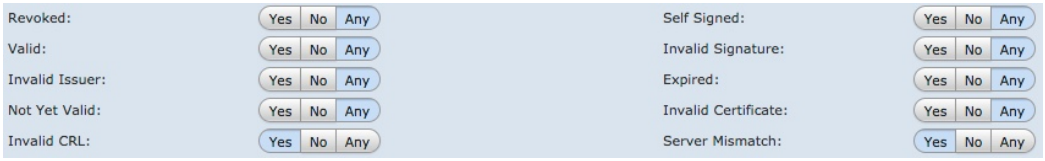
The following figure shows a certificate status rule condition checking for the absence of a status. In this case, because of the configuration, it matches against traffic encrypted with a certificate that has not expired and monitors that traffic.



The following graphic illustrates a certificate status rule condition that matches on the presence or absence of several statuses. Because of the configuration, if the rule matches incoming traffic encrypted with a certificate issued by an invalid user, self-signed, invalid, or expired, it decrypts the traffic with a known key.



The following graphic illustrates a certificate status rule condition that matches if the SNI of the request matches the server name or if the CRL is not valid. Because of the configuration, if the rule matches either condition, traffic is blocked.



What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Cipher Suite TLS/SSL Rule Conditions

The system provides predefined cipher suites you can add to a cipher suite rule condition. You can also add cipher suite list objects containing multiple cipher suites.

Note You cannot add new cipher suites. You can neither modify nor delete predefined cipher suites.

You can add a maximum of 50 cipher suites and cipher suite lists to the **Selected Cipher Suites** in a single cipher suite condition. The system supports adding the following cipher suites to a cipher suite condition:

- SSL_RSA_FIPS_WITH_3DES_EDE_CBC_SHA
- SSL_RSA_FIPS_WITH_DES_CBC_SHA
- TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA
- TLS_DHE_RSA_WITH_AES_128_CBC_SHA
- TLS_DHE_RSA_WITH_AES_128_CBC_SHA256
- TLS_DHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_DHE_RSA_WITH_AES_256_CBC_SHA
- TLS_DHE_RSA_WITH_AES_256_CBC_SHA256
- TLS_DHE_RSA_WITH_AES_256_GCM_SHA384
- TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA
- TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA256
- TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA
- TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA256
- TLS_DHE_RSA_WITH_DES_CBC_SHA
- TLS_DH_Anon_WITH_AES_128_GCM_SHA256
- TLS_DH_Anon_WITH_AES_256_GCM_SHA384

- TLS_DH_Anon_WITH_CAMELLIA_128_CBC_SHA
- TLS_DH_anon_WITH_CAMELLIA_128_CBC_SHA256
- TLS_DH_Anon_WITH_CAMELLIA_256_CBC_SHA
- TLS_DH_anon_WITH_CAMELLIA_256_CBC_SHA256
- TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384
- TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_ECDSA_WITH_NULL_SHA
- TLS_ECDHE_ECDSA_WITH_RC4_128_SHA
- TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384
- TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_RSA_WITH_NULL_SHA
- TLS_ECDHE_RSA_WITH_RC4_128_SHA
- TLS_RSA_WITH_3DES_EDE_CBC_SHA
- TLS_RSA_WITH_AES_128_CBC_SHA
- TLS_RSA_WITH_AES_128_CBC_SHA256
- TLS_RSA_WITH_AES_128_GCM_SHA256
- TLS_RSA_WITH_AES_256_CBC_SHA
- TLS_RSA_WITH_AES_256_CBC_SHA256
- TLS_RSA_WITH_AES_256_GCM_SHA384
- TLS_RSA_WITH_CAMELLIA_128_CBC_SHA
- TLS_RSA_WITH_CAMELLIA_128_CBC_SHA256
- TLS_RSA_WITH_CAMELLIA_256_CBC_SHA

- TLS_RSA_WITH_CAMELLIA_256_CBC_SHA256
- TLS_RSA_WITH_DES_CBC_SHA
- TLS_RSA_WITH_NULL_MD5
- TLS_RSA_WITH_NULL_SHA
- TLS_RSA_WITH_RC4_128_MD5
- TLS_RSA_WITH_RC4_128_SHA

Note the following:

- If you add cipher suites not supported for your deployment, you cannot deploy your configuration. For example, passive deployments do not support decrypting traffic with the any of the ephemeral Diffie-Hellman (DHE) or ephemeral elliptic curve Diffie-Hellman (ECDHE) cipher suites. Creating a rule with these cipher suites prevents you from deploying your access control policy.
- If you configure a cipher suite condition with a cipher suite, any external certificate objects you add to a certificate condition, or internal CA objects you associate with the **Decrypt - Resign** action, must match the cipher suite's signature algorithm type. For example, if your rule's cipher suite condition references an EC-based cipher suite, any server certificates you add, or CA certificates you associate with the **Decrypt - Resign** action, must also be EC-based. If you mismatch signature algorithm types in this case, the policy editor displays a warning icon next to the rule.
- You can add an anonymous cipher suite to the **Cipher Suite** condition in an SSL rule, but keep in mind:
  - The system automatically strips anonymous cipher suites during ClientHello processing. For the system to use the rule, you must also configure your TLS/SSL rules in an order that prevents ClientHello processing. For more information, see SSL Rule Order, on page 1197.
  - You cannot use either the **Decrypt - Resign** or **Decrypt - Known Key** action in the rule, because the system cannot decrypt traffic encrypted with an anonymous cipher suite.
- When specifying a cipher suite as a rule condition, consider that the rule matches on the negotiated cipher suite in the ServerHello message, rather than on the full list of cipher suites specified in the ClientHello message. During ClientHello processing, the managed device strips unsupported cipher suites from the ClientHello message. However, if this results in all specified cipher suites being stripped, the system retains the original list. If the system retains unsupported cipher suites, subsequent evaluation results in an undecrypted session.

Controlling Encrypted Traffic by Cipher Suite

Procedure

Step 1 In the SSL rule editor, select Cipher Suite.

Step 2 Find the cipher suites you want to add from the **Available Cipher Suites**, as follows;

- To add a cipher suite list on the fly, which you can then add to the condition, click **Add** () above the **Available Cipher Suites** list.

- To search for cipher suites and lists to add, click the **Search by name or value** prompt above the **Available Cipher Suites** list, then type either the name of the cipher suite, or a value in the cipher suite. The list updates as you type to display matching cipher suites.

Step 3 To select a cipher suite, click it. To select all cipher suites, right-click and then select **Select All**.

Step 4 Click **Add to Rule**.

Tip You can also drag and drop selected cipher suites.

Step 5 Add or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Cipher Suite Lists, on page 443

Encryption Protocol Version TLS/SSL Rule Conditions

You can choose to match against traffic encrypted with SSL version 3.0, or TLS version 1.0, 1.1, or 1.2. By default, all protocol versions are selected when you create a rule; if you select multiple versions, encrypted traffic that matches any of the selected versions matches the rule. You must select at least one protocol version when saving the rule condition.

You cannot select SSL v2.0 in a version rule condition; the system does not support decrypting traffic encrypted with SSL version 2.0. You can configure an undecryptable action to allow or block this traffic without further inspection.

Controlling Traffic by Encryption Protocol Version

Procedure

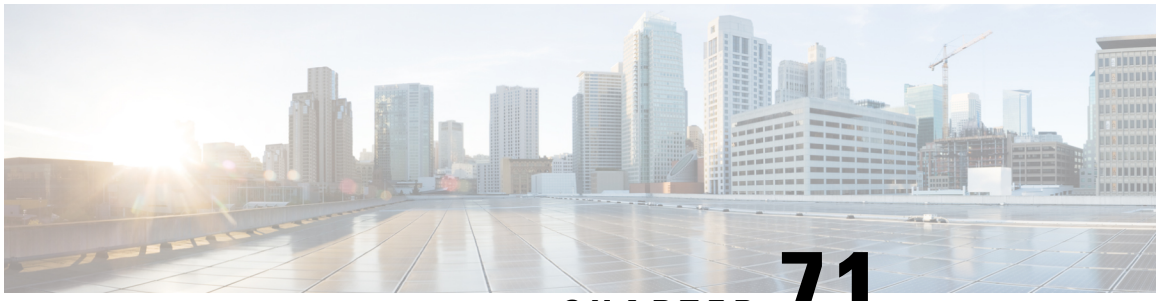
Step 1 In the SSL rule editor, select Version.

Step 2 Select the protocol versions you want to match against.

Step 3 Add or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.



CHAPTER 71

Troubleshoot TLS/SSL Rules

You can diagnose whether or not applications on your network use TLS/SSL pinning with connection events. If applications are using TLS/SSL pinning, see TLS/SSL Rule Guidelines and Limitations, on page 1345.

- About TLS/SSL Pinning, on page 1393
- Verify TLS/SSL Cipher Suites, on page 1397

About TLS/SSL Pinning

Some applications use a technique referred to as *TLS/SSL pinning* or *certificate pinning*, which embeds the fingerprint of the original server certificate in the application itself. As a result, if you configured a TLS/SSL rule with a **Decrypt - Resign** action, when the application receives a resigned certificate from a managed device, validation fails and the connection is aborted.

To confirm that TLS/SSL pinning is occurring, attempt to log in to a mobile application like Facebook. If a network connection error is displayed, log in using a web browser. (For example, you *cannot* log in to a Facebook mobile application but *can* log in to Facebook using Safari or Chrome.) You can use Firepower Management Center connection events as further proof of TLS/SSL pinning

Note TLS/SSL pinning is not limited to mobile applications.

Related Topics

Troubleshoot TLS/SSL Pinning, on page 1393

Troubleshoot TLS/SSL Pinning

You can view connection events to determine whether or not the devices are experiencing SSL pinning. You must add at least the **SSL Flow Flags** and **SSL Flow Messages** columns to the table view of connection events.

Before you begin

- Enable logging for your TLS/SSL rules as discussed in Logging Decryptable Connections with SSL Rules, on page 2219.

- Log in to a mobile application like Facebook; if a network connection error displays, log in to Facebook using Chrome or Safari. If you *can* log in using a web browser but not the native application, SSL pinning is likely occurring.

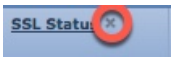
Procedure

Step 1 If you haven't done so already, log in to the Firepower Management Center.

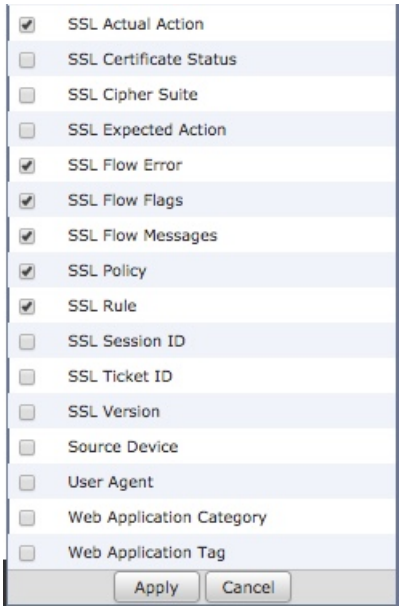
Step 2 Click **Analysis** > **Connections** > **Events**.

Step 3 Click **Table View of Connection Events**.

Step 4 Click **x** on any column in the connection events table to add additional columns for at least **SSL Flow Flags** and **SSL Flow Messages**.



The following example shows adding the **SSL Actual Action**, **SSL Flow Error**, **SSL Flow Flags**, **SSL Flow Messages**, **SSL Policy**, and **SSL Rule** columns to the table of connection events.



The columns are added in the order discussed in Connection and Security Intelligence Event Fields, on page 2225.

Step 5 Click **Apply**.

Step 6 The following paragraphs discuss how you can identify SSL pinning behavior.

Step 7 If you determine that applications in your network use SSL pinning, see TLS/SSL Rule Guidelines and Limitations, on page 1345.

What to do next

You can use TLS/SSL connection events to confirm TLS/SSL pinning is occurring by looking for any of the following:

- Applications that send an SSL ALERT Message as soon as the client receives the SERVER_HELLO, SERVER_CERTIFICATE, SERVER_HELLO_DONE message from the server, followed by a TCP Reset, exhibit the following symptoms. (The alert, `Unknown CA (48)`, can be viewed using a packet capture.)
  - The SSL Flow Flags column displays `ALERT_SEEN` but *not* `APP_DATA_C2S` or `APP_DATA_S2C`.
  - The SSL Flow Messages column typically displays: `CLIENT_HELLO`, `SERVER_HELLO`, `SERVER_CERTIFICATE`, `SERVER_KEY_EXCHANGE`, `SERVER_HELLO_DONE`.
  - `Success` is displayed in the SSL Flow Error column.
- Applications that send no alerts but instead send TCP Reset after the SSL handshake is finished exhibit the following symptoms:
  - The SSL Flow Flags column does *not* display `ALERT_SEEN`, `APP_DATA_C2S`, or `APP_DATA_S2C`.
  - The SSL Flow Messages column typically displays: `CLIENT_HELLO`, `SERVER_HELLO`, `SERVER_CERTIFICATE`, `SERVER_KEY_EXCHANGE`, `SERVER_HELLO_DONE`, `CLIENT_KEY_EXCHANGE`, `CLIENT_CHANGE_CIPHER_SPEC`, `CLIENT_FINISHED`,`SERVER_CHANGE_CIPHER_SPEC`, `SERVER_FINISHED`.
  - `Success` is displayed in the SSL Flow Error column.

Related Topics

Using Connection and Security Intelligence Event Tables, on page 2245
Connection and Security Intelligence Event Fields, on page 2225
Information Available in Connection Event Fields, on page 2242
Event Searches, on page 2177
Troubleshoot Unknown or Bad Certificates or Certificate Authorities, on page 1395

Troubleshoot Unknown or Bad Certificates or Certificate Authorities

You can view connection events to determine whether or not the devices are experiencing unknown certificate authorities, bad certificates, or unknown certificates. This procedure can also be used if a TLS/SSL certificate has been pinned. You must add at least the **SSL Flow Flags** and **SSL Flow Messages** columns to the table view of connection events.

Before you begin

- Set up a TLS/SSL decryption rule.
- Enable logging for your TLS/SSL rules as discussed in Logging Decryptable Connections with SSL Rules, on page 2219.

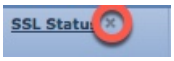
Procedure

Step 1 If you haven't done so already, log in to the Firepower Management Center.

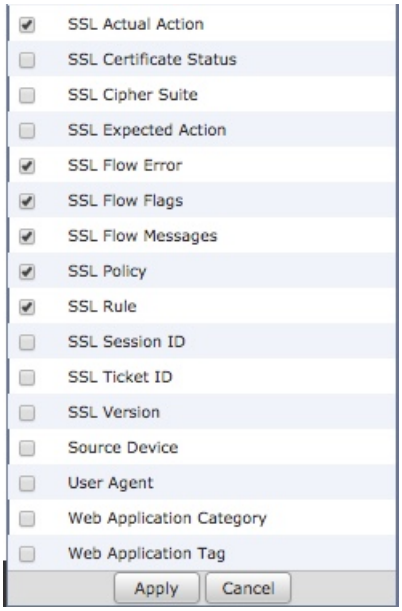
Step 2 Click **Analysis** > **Connections** > **Events**.

Step 3 Click **Table View of Connection Events**.

Step 4 Click **x** on any column in the connection events table to add additional columns for at least **SSL Flow Flags** and **SSL Flow Messages**.



The following example shows adding the **SSL Actual Action**, **SSL Flow Error**, **SSL Flow Flags**, **SSL Flow Messages**, **SSL Policy**, and **SSL Rule** columns to the table of connection events.



The columns are added in the order discussed in Connection and Security Intelligence Event Fields, on page 2225.

Step 5 Click **Apply**.

Step 6 The following table discusses how you can determine if a certificate or certificate authority is bad or missing.

SSL flow flag	Meaning
CLIENT_ALERT_SEEN_UNKNOWN_CA	Indicates a valid certificate chain or partial chain was received by an SSL client application, but the certificate was not accepted because the CA certificate could not be located or could not be matched with a known, trusted CA. This message always indicates an unrecoverable error.
CLIENT_ALERT_SEEN_BAD_CERTIFICATE	A certificate was corrupt, contained signatures that did not verify correctly, or had other problems.

SSL flow flag	Meaning
CLIENT_ALERT_SEEN_CERTIFICATE_UNKNOWN	Some other (unspecified) issue arose in processing the certificate, rendering it unacceptable.

Verify TLS/SSL Cipher Suites

Before you begin

This topic discusses actions you must take if you see the following error when saving a TLS/SSL rule that has cipher suite conditions:

```
Traffic cannot match this rule; none of your selected cipher suites contain a signature
algorithm that the resigning CA's signature algorithm
```

The error indicates that one or more of the cipher suites you chose for the TLS/SSL rule condition are incompatible with the certificate used in the TLS/SSL rule. To resolve the issue, you must have access to the certificate you're using.

Note The tasks in this topic assume knowledge of how TLS/SSL encryption works.

Procedure

Step 1 When you attempt to save an SSL rule with either **Decrypt - Resign** or **Decrypt - Known Key** with specified cipher suites, the following error is displayed:

Example:

```
Traffic cannot match this rule; none of your selected cipher suites contain a
signature algorithm that the resigning CA's signature algorithm
```

Step 2 Locate the certificate you're using to decrypt traffic and, if necessary, copy the certificate to a system that can run openssl commands.

Step 3 Run the following command to display the signature algorithm used by the certificate:

openssl x509 -in *CertificateName* **-text -noout**

The first few lines of output are displayed similar to the following:

```
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number: 4105 (0x1009)
    Signature Algorithm: ecdsa-with-SHA256
```

Step 4 The **Signature algorithm** tells you the following:

- The cryptographic function used (in the preceding example, **ECDSA** means Elliptic Curve Digital Signature Algorithm).

- The hash function used to create a digest of the encrypted message (in the preceding example, **SHA256**).

Step 5 Search a resource such as OpenSSL at University of Utah for cipher suites that match those values. The cipher suite must be in RFC format.
You can also search a variety of other sites, such as Server Side TLS at the Mozilla wiki or Appendix C of RFC 5246. Cipher Suites in TLS/SSL (Schannel SSP) in Microsoft documentation has a detailed explanation of cipher suites.

Step 6 If necessary, translate the OpenSSL name to an RFC name that the Firepower Management System uses. See the RFC mapping list on the on the `https://testssl.sh` site.

Step 7 The previous example, **ecdsa-with-SHA256**, can be found in the Modern Compatibility List on the Mozilla wiki.

a) Choose only cipher suites that have **ECDSA** and **SHA-256** in the name. These cipher suites follow:

```
ECDHE-ECDSA-AES128-GCM-SHA256
ECDHE-ECDSA-AES128-SHA256
```

b) Find the corresponding RFC cipher suite on RFC mapping list. These cipher suites follow:

```
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
```

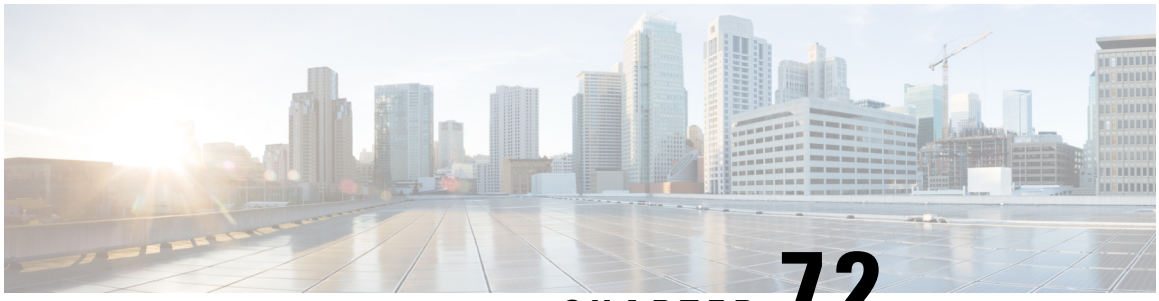
Step 8 Add the preceding cipher suites to your TLS/SSL rule.



PART XVII

Advanced Malware Protection (AMP) and File Control

- File Policies and Malware Protection, on page 1401
- File and Malware Inspection Performance and Storage Tuning, on page 1441



CHAPTER 72

File Policies and Malware Protection

The following topics provide an overview of file control, file policies, file rules, Advanced Malware Protection (AMP), cloud connections, and dynamic analysis connections.

- About File Policies and Advanced Malware Protection, on page 1401
- Requirements and Prerequisites for File Policies, on page 1402
- License Requirements for File and Malware Policies, on page 1403
- Best Practices for File Policies and Malware Detection , on page 1403
- How to Configure Malware Protection, on page 1406
- Cloud Connections for Malware Protection, on page 1410
- File Policies and File Rules, on page 1419
- Retrospective Disposition Changes, on page 1434
- (Optional) Malware Protection with AMP for Endpoints, on page 1434
- History for File Policies and Malware Protection, on page 1439

About File Policies and Advanced Malware Protection

To detect and block malware, use file policies. You can also use file policies to detect and control traffic by file type.

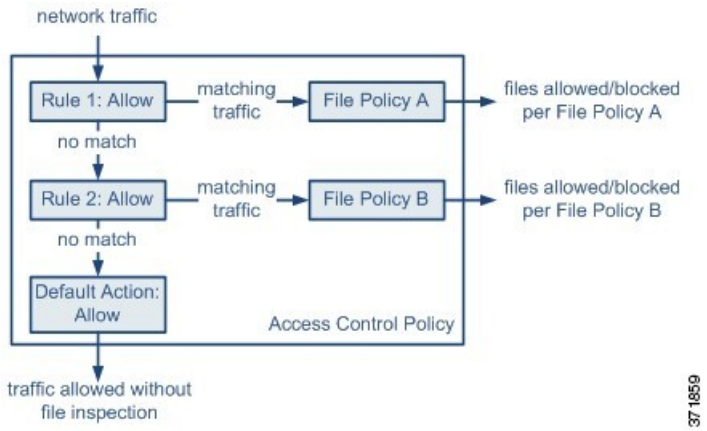
Advanced Malware Protection (AMP) for Firepower can detect, capture, track, analyze, log, and optionally block the transmission of malware in network traffic. In the Firepower Management Center web interface, this feature is called *AMP for Networks*, formerly called *AMP for Firepower*. Advanced Malware Protection identifies malware using managed devices deployed inline and threat data from the Cisco cloud.

You associate file policies with access control rules that handle network traffic as part of your overall access control configuration.

When the system detects malware on your network, it generates file and malware events. To analyze file and malware event data, see File/Malware Events and Network File Trajectory, on page 2295.

File Policies

A file policy is a set of configurations that the system uses to perform malware protection and file control, as part of your overall access control configuration. This association ensures that before the system passes a file in traffic that matches an access control rule's conditions, it first inspects the file. Consider the following diagram of a simple access control policy in an inline deployment.



The policy has two access control rules, both of which use the Allow action and are associated with file policies. The policy's default action is also to allow traffic, but without file policy inspection. In this scenario, traffic is handled as follows:

- Traffic that matches `Rule 1` is inspected by `File Policy A`.
- Traffic that does not match `Rule 1` is evaluated against `Rule 2`. Traffic that matches `Rule 2` is inspected by `File Policy B`.
- Traffic that does not match either rule is allowed; you cannot associate a file policy with the default action.

By associating different file policies with different access control rules, you have granular control over how you identify and block files transmitted on your network.

Requirements and Prerequisites for File Policies

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Access Admin

License Requirements for File and Malware Policies

To Do This	License Required	File Rule Action
Block or allow all files of a particular type (for example, block all .exe files)	Threat (for FTD devices) Protection (for Classic devices)	Allow, Block, Block with Reset
Selectively allow or block files based on a judgment that it contains or is likely to contain malware	Threat (for FTD devices) Protection (for Classic devices) Malware	Malware Cloud Lookup, Block Malware
Store files	Threat (for FTD devices) Protection (for Classic devices) Malware	Any file rule action with **Store Files** selected

For details about Malware licenses, see:

- Malware Licenses for Firepower Threat Defense Devices, on page 111
- Malware Licenses for Classic Devices, on page 128

Best Practices for File Policies and Malware Detection

In addition to the items described below, follow the steps in How to Configure Malware Protection, on page 1406 and referenced topics.

File Rule Best Practices

Note the following guidelines and limitations when configuring file rules:

- A rule configured to block files in a passive deployment does not block matching files. Because the connection continues to transmit the file, if you configure the rule to log the beginning of the connection, you may see multiple events logged for this connection.
- A policy can include multiple rules. When you create the rules, ensure that no rule is "shadowed" by a previous rule.
- The file types supported for dynamic analysis are a subset of the file types supported for other types of analysis. To view the file types supported for each type of analysis, navigate to the file rule configuration page, select the **Block Malware** action, and select the checkboxes of interest.

To ensure that the system examines all file types, create separate rules (within the same policy) for dynamic analysis and for other types of analysis.

- If a file rule is configured with a **Malware Cloud Lookup** or **Block Malware** action and the Firepower Management Center cannot establish connectivity with the AMP cloud, the system cannot perform any configured rule action options until connectivity is restored.
- Cisco recommends that you enable **Reset Connection** for the **Block Files** and **Block Malware** actions to prevent blocked application sessions from remaining open until the TCP connection resets. If you do not reset connections, the client session will remain open until the TCP connection resets itself.
- If you are monitoring high volumes of traffic, do **not** store all captured files, or submit all captured files for dynamic analysis. Doing so can negatively impact system performance.
- You cannot perform malware analysis on all file types detected by the system. After you select values from the **Application Protocol**, **Direction of Transfer**, and **Action** drop-down lists, the system constrains the list of file types.

File Detection Best Practices

Consider the following notes and limitations for file detection:

- If adaptive profiling is not enabled, access control rules cannot perform file control, including AMP.
- If a file matches a rule with an application protocol condition, file event generation occurs after the system successfully identifies a file's application protocol. Unidentified files do not generate file events.
- FTP transfers commands and data over different channels. In a passive or inline tap mode deployment, the traffic from an FTP data session and its control session may not be load-balanced to the same internal resource.
- If the total number of bytes for all file names for files in a POP3, POP, SMTP, or IMAP session exceeds 1024, file events from the session may not reflect the correct file names for files that were detected after the file name buffer filled.
- When transmitting text-based files over SMTP, some mail clients convert newlines to the CRLF newline character standard. Since Mac-based hosts use the carriage return (CR) character and Unix/Linux-based hosts use the line feed (LF) character, newline conversion by the mail client can modify the size of the file. Note that some mail clients default to newline conversion when processing an unrecognizable file type.
- To detect ISO files, set the "Limit the number of bytes inspected when doing file type detection" option to a value greater than 36870, as described in File and Malware Inspection Performance and Storage Options, on page 1441.
- .Exe files inside some .rar archives cannot be detected, including possibly rar5.

File Blocking Best Practices

Consider the following notes and limitations for file blocking:

- If an end-of-file marker is not detected for a file, regardless of transfer protocol, the file will not be blocked by a **Block Malware** rule or the custom detection list. The system waits to block the file until

the entire file has been received, as indicated by the end-of-file marker, and blocks the file after the marker is detected.

- If the end-of-file marker for an FTP file transfer is transmitted separately from the final data segment, the marker will be blocked and the FTP client will indicate that the file transfer failed, but the file will actually completely transfer to disk.
- File rules with **Block Files** and **Block Malware** actions block automatic resumption of file download via HTTP by blocking new sessions with the same file, URL, server, and client application detected for 24 hours after the initial file transfer attempt occurs.
- In rare cases, if traffic from an HTTP upload session is out of order, the system cannot reassemble the traffic correctly and therefore will not block it or generate a file event.
- If you transfer a file over NetBIOS-ssn (such as an SMB file transfer) that is blocked with a **Block Files** rule, you may see a file on the destination host. However, the file is unusable because it is blocked after the download starts, resulting in an incomplete file transfer.
- If you create file rules to detect or block files transferred over NetBIOS-ssn (such as an SMB file transfer), the system does not inspect files transferred in an established TCP or SMB session started before you deploy an access control policy invoking the file policy so those files will not be detected or blocked.
- If you configure Firepower Threat Defense high availability, and failover occurs while the original active device is identifying the file, the file type is not synced. Even if your file policy blocks that file type, the new active device downloads the file.

File Policy Best Practices

Note the following general guidelines and limitations when configuring file policies.

- You can associate a single file policy with an access control rule whose action is **Allow**, **Interactive Block**, or **Interactive Block with reset**.
- You **cannot** use a file policy to inspect traffic handled by the access control default action.
- For a new policy, the web interface indicates that the policy is not in use. If you are editing an in-use file policy, the web interface tells you how many access control policies use the file policy. In either case, you can click the text to jump to the Access Control Policies page.
- For file blocking to work, the NAP policy you apply to the access control policy must be operating in Protection mode, also known as Inline mode.
- Based on your configuration, you can either inspect a file the first time the system detects it, and wait for a cloud lookup result, or pass the file on this first detection without waiting for the cloud lookup result.
- By default, file inspection of encrypted payloads is disabled. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has file inspection configured.

How to Configure Malware Protection

This topic summarizes the steps you must take to set up your Firepower system to protect your network from malicious software.

Procedure

Step 1 Plan and Prepare for Malware Protection, on page 1406

Step 2 Configure File Policies, on page 1407

Step 3 Add File Policies to Your Access Control Configuration, on page 1408

Step 4 Configure network discovery policies to associate file and malware events with hosts on your network.

(Do not simply turn on network discovery; you must configure it to discover hosts on your network to build a network map of your organization.)

See Network Discovery Policies, on page 1917 and subtopics.

Step 5 Deploy policies to managed devices.

See Deploy Configuration Changes, on page 346.

Step 6 Test your system to be sure it is processing malicious files as you expect it to.

Step 7 Set Up Maintenance and Monitoring of Malware Protection, on page 1410

What to do next

- (Optional) To further enhance detection of malware in your network, deploy and integrate Cisco's AMP for Endpoints product. See (Optional) Malware Protection with AMP for Endpoints, on page 1434 and subtopics.
- Understand how to investigate file and malware events.

  See File/Malware Events and Network File Trajectory, on page 2295.

Plan and Prepare for Malware Protection

This procedure is the first set of steps in the complete process for configuring your system to provide malware protection.

Procedure

Step 1 Purchase and install licenses.

See License Requirements for File and Malware Policies, on page 1403 and Licensing the Firepower System, on page 105.

Step 2 Understand how file policies and malware protection fit into your access control plan.

See the chapter Understanding Access Control, on page 1185.

Step 3 Understand the file analysis and malware protection tools.

See File Rule Actions, on page 1426 and subtopics.

Consider also Advanced and Archive File Inspection Options, on page 1420.

Step 4 Determine whether you will use public clouds or private (on-premises) clouds for malware protection (file analysis and dynamic analysis.)

See Cloud Connections for Malware Protection, on page 1410 and subtopics.

Step 5 If you will use private (on-premises) clouds for malware protection: Purchase, deploy, and test those products.

For information, contact your Cisco sales representative or authorized reseller.

Step 6 Configure your firewall to allow communications with your chosen clouds.

See Security, Internet Access, and Communication Ports, on page 2425.

Step 7 Configure connections between Firepower and the malware protection clouds (public or private).

- For the AMP cloud, see Change AMP Options, on page 1415.
- If you deployed an on-premises Cisco Threat Grid appliance, see Connect to an On-Premises Dynamic Analysis Appliance, on page 1417. (Access to the public Threat Grid cloud does not require configuration.)

What to do next

Continue with the next step in the malware protection workflow:

See How to Configure Malware Protection, on page 1406.

Configure File Policies

Before you begin

Complete the tasks up to this point in the malware protection workflow:

See How to Configure Malware Protection, on page 1406.

Procedure

Step 1 Review file policy and file rule restrictions.

See Best Practices for File Policies and Malware Detection , on page 1403 and subtopics.

Step 2 Create a file policy.

See Create or Edit a File Policy, on page 1419.

Step 3 Create rules within your file policy.

See File Rules, on page 1424 and subtopics.

Step 4 Configure advanced options.

See Advanced and Archive File Inspection Options, on page 1420.

What to do next

Continue with the next step in the malware protection workflow:

See How to Configure Malware Protection, on page 1406.

Add File Policies to Your Access Control Configuration

An access control policy can have multiple access control rules associated with file policies. You can configure file inspection for any Allow or Interactive Block access control rule, which permits you to match different file and malware inspection profiles against different types of traffic on your network before it reaches its final destination.

Before you begin

Complete the tasks up to this point in the malware protection workflow:

See How to Configure Malware Protection, on page 1406.

Procedure

Step 1 Review guidelines for file policies in access control policies. (These are different from the file rule and file policy guidelines that you looked at previously.)

Review File and Intrusion Inspection Order, on page 1189.

Step 2 Associate the file policy with an access control policy.

See Configuring an Access Control Rule to Perform Malware Protection, on page 1409

Step 3 Assign the access control policy to managed devices.

See Setting Target Devices for an Access Control Policy, on page 1210.

What to do next

Continue with the next step in the malware protection workflow:

See How to Configure Malware Protection, on page 1406.

Configuring an Access Control Rule to Perform Malware Protection

Caution Selecting **Detect Files** or **Block Files**, enabling or disabling **Store files** in a **Detect Files** or **Block Files** rule, or adding the first or removing the last file rule that combines the **Malware Cloud Lookup** or **Block Malware** file rule action with an analyis option (**Spero Analysis or MSEXE**, **Dynamic Analysis**, or **Local Malware Analysis**) or a store files option (**Malware, Unknown**, **Clean**, or **Custom**), restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Note Inline normalization is enabled automatically when a file policy is included in an access control rule. For more information, see The Inline Normalization Preprocessor, on page 1762.

Before you begin

- Adaptive profiling **must** be enabled (its default state) as described in Configuring Adaptive Profiles, on page 1814 for access control rules to perform file control, including AMP.
- You must be an Admin, Access Admin, or Network Admin user to perform this task.

Procedure

Step 1 In the access control rule editor (from **Policies > Access Control)**, choose an **Action** of **Allow**, **Interactive Block**, or **Interactive Block with reset**.

Step 2 Click **Inspection**.

Step 3 Choose a **File Policy** to inspect traffic that matches the access control rule, or choose **None** to disable file inspection for matching traffic.

Step 4 (Optional) Disable logging of file or malware events for matching connections by clicking **Logging** and unchecking **Log Files**.

Note Cisco recommends that you leave file and malware event logging enabled.

Step 5 Save the rule.

Step 6 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Create or Edit a File Policy, on page 1419
Snort® Restart Scenarios, on page 348

Set Up Maintenance and Monitoring of Malware Protection

Ongoing maintenance is essential for protecting your network.

Before you begin

Configure your system to protect your network from malware.

See How to Configure Malware Protection, on page 1406 and referenced procedures.

Procedure

Step 1 Ensure that your system always has the most current and effective protection.

See Maintain Your System: Update File Types Eligible for Dynamic Analysis, on page 1419.

Step 2 Configure alerts for malware-related events and health monitoring.

See Configuring AMP for Networks Alerting, on page 2082 and information in Health Monitoring, on page 275 about the following modules:

- Local Malware Analysis
- Security Intelligence
- Intrusion and File Event Rate
- AMP for Firepower Status
- AMP for Endpoints Status

What to do next

Review "What to do next items" in the malware protection workflow:

See How to Configure Malware Protection, on page 1406.

Cloud Connections for Malware Protection

Connections to public or private clouds are required in order to protect your network from malware.

AMP Clouds

The Advanced Malware Protection (AMP) cloud is a Cisco-hosted server that uses big data analytics and continuous analysis to provide intelligence that the system uses to detect and block malware on your network.

The AMP cloud provides dispositions for possible malware detected in network traffic by managed devices, as well as data updates for local malware analysis and file pre-classification.

If your organization has deployed AMP for Endpoints and configured Firepower to import its data, the system imports this data from the AMP cloud, including scan records, malware detections, quarantines, and indications of compromise (IOC).

Cisco offers the following options for obtaining data from the Cisco cloud about known malware threats:

- **AMP public cloud**

  Your Firepower Management Center communicates directly with the public Cisco cloud.

- **An AMP private cloud**

  An AMP private cloud is deployed on your network and acts as a compressed, on-premises AMP cloud, as well as an anonymized proxy to connect to the public AMP cloud. For details, see Cisco AMP Private Cloud, on page 1413.

  If you integrate with AMP for Endpoints, the AMP private cloud has some limitations. See AMP for Endpoints and AMP Private Cloud, on page 1436.

Dynamic Analysis Cloud

- **Cisco Threat Grid cloud**

  Public cloud that processes eligible files that you send for dynamic analysis, and provides threat scores and dynamic analysis reports. Firepower supports 200 samples/day for Cisco Threat Grid analysis.

- **On-premises Cisco Threat Grid appliance**

  If your organization's security policy does not allow the Firepower System to send files outside of your network, you can deploy an on-premises appliance. This appliance does not contact the public Cisco Threat Grid cloud.

  For more information, see Dynamic Analysis On-Premises Appliance (Cisco Threat Grid) , on page 1417.

Configure Connections to AMP and Threat Grid Clouds

- AMP Cloud Connection Configurations, on page 1411
- Dynamic Analysis Connections, on page 1416

AMP Cloud Connection Configurations

The following topics describe AMP cloud connection configurations for different scenarios:

- Choose an AMP Cloud, on page 1412
- Connecting to an AMP Private Cloud, on page 1413
- Integrate Firepower and AMP for Endpoints, on page 1437

The following topics are also relevant:

- Cisco AMP Private Cloud, on page 1413
- Requirements and Best Practices for AMP Cloud Connections, on page 1412
- Managing Connections to the AMP Cloud (Public or Private) , on page 1414

Requirements and Best Practices for AMP Cloud Connections

Requirements for AMP Cloud Connections

You must be an Admin user to set up the AMP cloud.

To ensure your FMC can communicate with the AMP cloud, see the topics under Security, Internet Access, and Communication Ports, on page 2425.

To use the legacy port for AMP communications, see Communication Port Requirements, on page 2427.

AMP and High Availability

Although they share file policies and related configurations, Firepower Management Centers in a high availability pair share neither cloud connections nor captured files, file events, and malware events. To ensure continuity of operations, and to ensure that detected files' malware dispositions are the same on both Firepower Management Centers, both Active and Standby Firepower Management Centers must have access to the cloud.

In high availability configurations, you must configure AMP cloud connections independently on the Active and Standby instances of the Firepower Management Center; these configurations are not synchronized.

These requirements apply to both public and private AMP clouds.

AMP Cloud Connections and Multitenancy

In a multidomain deployment, you configure the AMP for Networks connection at the Global level only. Each Firepower Management Center can have only one AMP for Networks connection.

Choose an AMP Cloud

By default, a connection to the United States (US) AMP public cloud is configured and enabled for your Firepower system. (This connection appears in the web interface as AMP for Networks and sometimes AMP for Firepower.) You cannot delete or disable an AMP for Networks cloud connection, but you can switch between different geographical AMP clouds, or configure an AMP private cloud connection.

Before you begin

- If you will use an AMP private cloud, see Connecting to an AMP Private Cloud, on page 1413 instead of this topic.
- Unless Firepower is integrated with AMP for Endpoints, you can configure only one AMP cloud connection. This connection is labeled **AMP for Networks** or **AMP for Firepower**.
- If you have deployed AMP for Endpoints and you want to add one or more AMP clouds to integrate that application with Firepower, see Integrate Firepower and AMP for Endpoints, on page 1437.
- See Requirements and Best Practices for AMP Cloud Connections, on page 1412.

Procedure

Step 1 Choose **AMP** > **AMP Management**.

Step 2 Click pencil to edit the existing cloud connection.

Step 3 From the **Cloud Name** drop-down list, choose the regional cloud nearest to your Firepower Management Center.

Step 4 Click **Save**.

What to do next

- If your deployment is a high-availability configuration, see Requirements and Best Practices for AMP Cloud Connections, on page 1412.
- (Optional) Change AMP Options, on page 1415.

Cisco AMP Private Cloud

The Firepower Management Center must connect to the AMP cloud for disposition queries for files detected in network traffic and receipt of retrospective malware events. This cloud can be public or private.

Your organization may have privacy or security concerns that make frequent or direct connections between your monitored network and the AMP cloud difficult or impossible. In these situations, you can set up a Cisco AMP Private Cloud, a proprietary Cisco product that acts as a compressed, on-premises version of the AMP cloud, as well as a secure mediator between your network and the AMP cloud. Connecting a Firepower Management Center to an AMP private cloud disables existing direct connections to the public AMP cloud.

All connections to the AMP cloud funnel through the AMP private cloud, which acts as an anonymized proxy to ensure the security and privacy of your monitored network. This includes disposition queries for files detected in network traffic, receiving of retrospective malware events, and so on. The AMP private cloud does not share any of your endpoint data over an external connection.

Note The AMP private cloud does **not** perform dynamic analysis, nor does it support anonymized retrieval of threat intelligence for other features that rely on Cisco Collective Security Intelligence (CSI), such as URL and Security Intelligence filtering.

For information about AMP private cloud (sometimes referred to as "AMPv"), see https://www.cisco.com/c/en/us/products/security/fireamp-private-cloud-virtual-appliance/index.html.

Connecting to an AMP Private Cloud

Before you begin

- Configure your Cisco AMP private cloud or clouds according to the directions in the documentation for that product. During configuration, note the private cloud host name. You will need this host name in order to to configure the connection on the Firepower Management Center.
- Make sure the Firepower Management Center can communicate with the AMP private cloud, and confirm that the private cloud has internet access so it can communicate with the public AMP cloud. See the topics under Security, Internet Access, and Communication Ports, on page 2425.
- Unless your deployment is integrated with AMP for Endpoints, each Firepower Management Center can have only one AMP cloud connection. This connection is labeled **AMP for Networks** or **AMP for Firepower**.

If you integrate with AMP for Endpoints, you can configure multiple AMP for Endpoints cloud connections.

Procedure

Step 1 Choose **AMP** > **AMP Management**.

Step 2 Click **Add AMP Cloud Connection**.

Step 3 From the **Cloud Name** drop-down list, choose **Private Cloud**.

Step 4 Enter a **Name**.

This information appears in malware events that are generated or transmitted by AMP private cloud.

Step 5 In the **Host** field, enter the private cloud host name that you configured when you set up the private cloud.

Step 6 Click **Browse** next to the **Certificate Upload Path** field to browse to the location of a valid TLS or SSL encryption certificate for the private cloud. For more information, see the AMP private cloud documentation.

Step 7 If you want to use this private cloud for both AMP for Networks and AMP for Endpoints, select the **Use for AMP for Firepower** check box.

If you configured a different private cloud to handle AMP for Networks communications, you can clear this check box; if this is your only AMP private cloud connection, you cannot.

In a multidomain deployment, this check box appears only in the Global domain. Each Firepower Management Center can have only one AMP for Networks connection.

Step 8 To communicate with the AMP private cloud using a proxy, check the **Use Proxy for Connection** check box.

Step 9 Click **Register**, confirm that you want to disable existing direct connections to the AMP cloud, and finally confirm that you want to continue to the AMP private cloud management console to complete registration.

Step 10 Log into the management console and complete the registration process. For further instructions, see the AMP private cloud documentation.

What to do next

In high availability configurations, you must configure AMP cloud connections independently on the Active and Standby instances of the Firepower Management Center; these configurations are not synchronized.

Managing Connections to the AMP Cloud (Public or Private)

Use the Firepower Management Center to manage connections to public and private AMP clouds used for AMP for Networks or AMP for Endpoints or both.

You can delete a connection to a public or private AMP cloud if you no longer want to receive malware-related information from the cloud. Note that deregistering a connection using the AMP for Endpoints or AMP private cloud management console does not remove the connection from the system. Deregistered connections display a failed state on the Firepower Management Center web interface.

You can also temporarily disable a connection. When you reenable a cloud connection, the cloud resumes sending data to the system, including queued data from the disabled period.

Caution For disabled connections, the public or private AMP cloud can store malware events, indications of compromise, and so on until you re-enable the connection. In rare cases—for example, with a very high event rate or a long-term disabled connection—the cloud may not be able to store all information generated while the connection is disabled.

In a multidomain deployment, the system displays connections created in the current domain, which you can manage. It also displays connections created in ancestor domains, which you cannot manage. To manage connections in a lower domain, switch to that domain. Each Firepower Management Center can have only one AMP for Networks connection, which belongs to the Global domain.

Procedure

Step 1 Select **AMP** > **AMP Management**.

Step 2 Manage your AMP cloud connections:

- Delete — Click **Delete** ( ), then confirm your choice.
- Enable or Disable — Click the slider, then confirm your choice.

What to do next

In high availability configurations, you must configure AMP cloud connections independently on the Active and Standby instances of the Firepower Management Center; these configurations are not synchronized.

Change AMP Options

Procedure

Step 1 Choose **System** > **Integration**.

Step 2 Click **Cisco CSI**.

Step 3 Select options:

Table 98: AMP for Networks Options

Option	Description
Enable Automatic Local Malware Detection Updates	The local malware detection engine statically analyzes and preclassifies files using signatures provided by Cisco. If you enable this option, the Firepower Management Center checks for signature updates once every 30 minutes.

Option	Description
Share URI from Malware Events with Cisco	The system can send information about the files detected in network traffic to the AMP cloud. This information includes URI information associated with detected files and their SHA-256 hash values. Although sharing is opt-in, transmitting this information to Cisco helps future efforts to identify and track malware.
Use Legacy Port 32137 for AMP for Networks	By default, Firepower uses port 443/HTTPS to communicate with the AMP public or private cloud to obtain file disposition data. This option allows the system to use port 32137. If you updated from a previous version of the system, this option may be enabled. This option will be greyed out if the FMC is configured with Proxy settings.

Step 4 Click **Save**.

Dynamic Analysis Connections

Requirements for Dynamic Analysis

You must be an Admin, Access Admin, or Network Admin user, and be in the global domain, to use dynamic analysis.

With the appropriate license, the Firepower system automatically has access to the Cisco Threat Grid public cloud.

Dynamic analysis requires that managed devices have direct or proxied access to the Cisco Threat Grid public cloud or an on-premises Cisco Threat Grid appliance on port 443.

See also Which Files Are Eligible for Dynamic Analysis?, on page 1430.

If you will connect to an on-premises Threat Grid appliance, see also the prerequisites in Connect to an On-Premises Dynamic Analysis Appliance, on page 1417.

Viewing the Default Dynamic Analysis Connection

By default, the Firepower Management Center can connect to the public Cisco Threat Grid cloud for file submission and report retrieval. You can neither configure nor delete this connection.

Procedure

Step 1 Choose **AMP** > **Dynamic Analysis Connections**.

Step 2 Click **Edit** ( ).

Note For information about **Associate** ( )**Associate** ( ) on the **AMP** > **Dynamic Analysis Connections** page, see Enabling Access to Dynamic Analysis Results in the Public Cloud, on page 1418.

Dynamic Analysis On-Premises Appliance (Cisco Threat Grid)

If your organization has privacy or security concerns around submitting files to the public Cisco Threat Grid cloud, you can deploy an on-premises Cisco Threat Grid appliance. Like the public cloud, the on-premises appliance runs eligible files in a sandbox environment, and returns a threat score and dynamic analysis report to the Firepower System. However, the on-premises appliance does not communicate with the public cloud, or any other system external to your network.

For more information about on-premises Cisco Threat Grid appliances, see https://www.cisco.com/c/en/us/products/security/threat-grid/index.html.

Connect to an On-Premises Dynamic Analysis Appliance

If you install an on-premises Cisco Threat Grid appliance on your network, you can configure a dynamic analysis connection to submit files and retrieve reports from the appliance. When configuring the on-premises appliance dynamic analysis connection, you register the Firepower Management Center to the on-premises appliance.

Before you begin

- Set up your on-premises Cisco Threat Grid appliance; see the *Cisco Threat Grid Appliance Setup and Configuration Guide*.

  Documentation for this appliance is available from https://www.cisco.com/c/en/us/support/security/amp-threat-grid-appliances/tsd-products-support-series-home.html.

- If your Cisco Threat Grid appliance uses a self-signed public-key certificate, download the certificate from the Threat Grid appliance; see the *Cisco Threat Grid Appliance Administrator's Guide* for information.

  If you use a certificate signed by a Certificate Authority (CA), the certificate must meet the following requirements:

  - The server key and signed certificate must be installed on the Threat Grid appliance. Follow the upload instructions in the *Threat Grid Administrator's Guide*.
  - If there is a multi-level signing chain of CAs, all required intermediate certificates and the root certificate must be contained in a single file that will be uploaded to the FMC.
  - All certificates must be PEM-encoded.
  - The file's newlines must be UNIX, not DOS.

- If you want to connect to the on-premises appliance using a proxy, configure the proxy; see Modify FMC Management Interfaces, on page 886.
- Managed devices must have direct or proxied access to the Cisco Threat Grid appliance on port 443.

Procedure

Step 1 Choose **AMP** > **Dynamic Analysis Connections**.

Step 2 Click **Add New Connection**.

Step 3 Enter a **Name**.

Step 4 Enter a **Host**.

Step 5 Next to **Certificate Upload**, click **Browse** to upload the certificate for the on-premises appliance.

If the Threat Grid appliance will present a self-signed certificate, upload the certificate you downloaded from that appliance.

If the Threat Grid appliance will present a CA-signed certificate, upload the file containing the certificate signing chain.

Step 6 If you want to use a configured proxy to establish the connection, select **Use Proxy When Available**.

Step 7 Click **Register**.

Step 8 Click **Yes** to display the on-premises Cisco Threat Grid appliance login page.

Step 9 Enter your username and password to the on-premises Cisco Threat Grid appliance.

Step 10 Click **Sign in**.

Step 11 You have the following options:

- If you previously registered the Firepower Management Center to the on-premises appliance, click **Return**.
- If you did not register the Firepower Management Center, click **Activate**.

Enabling Access to Dynamic Analysis Results in the Public Cloud

Cisco Threat Grid offers more detailed reporting on analyzed files than is available in the Firepower Management Center. If your organization has an account in the Cisco Threat Grid public cloud, you can access the Cisco Threat Grid portal directly to view additional details about files sent for analysis from your managed devices. However, for privacy reasons, file analysis details are available only to the organization that submitted the files. Therefore, before you can view this information, you must associate your Firepower Management Center with the files submitted by its managed devices.

Before you begin

You must have an account on the Cisco Threat Grid public cloud, and have your account credentials ready.

Procedure

Step 1 Select **AMP > Dynamic Analysis Connections**.

Step 2 Click **Associate** ( ) in the table row corresponding to the Cisco Threat Grid public cloud.

A Cisco Threat Grid portal window opens.

Step 3 Sign in to the Cisco Threat Grid public cloud.

Step 4 Click **Submit Query**.

Note Do not change the default value in the **Devices** field.

If you have difficulties with this process, contact your Cisco Threat Grid representative at Cisco TAC.

It may take up to 24 hours for this change to take effect.

What to do next

After the association is activated, see Viewing Dynamic Analysis Results in the Cisco Threat Grid Public Cloud, on page 2313.

Maintain Your System: Update File Types Eligible for Dynamic Analysis

The list of file types eligible for Dynamic Analysis is determined by the vulnerability database (VDB), which is updated periodically (but no more than once per day.) If you are an Admin user, you can update file types eligible for dynamic analysis.

To ensure that your system has the current list:

Procedure

Step 1 Do one of the following:

- (Recommended) See Vulnerability Database Update Automation, on page 191
- Regularly check for new VDB updates, and Manually Update the VDB, on page 140 when needed.

  If you choose this option, we recommend that you schedule regular reminders to do this.

Step 2 If your file policies specify individual file types instead of the **Dynamic Analysis Capable** file type category, update your file policies to use the newly supported file types.

Step 3 If the list of eligible file types changes, deploy to managed devices.

File Policies and File Rules

Create or Edit a File Policy

Before you begin

If you are configuring policies for malware protection, see all required procedures in Configure File Policies, on page 1407.

Procedure

Step 1 Select **Policies** > **Access Control** > **Malware & File** .

Step 2 Create a new policy, or edit an existing policy.

If you are editing an existing policy: If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Tip To make a copy of an existing file policy, click **Copy** ( ), then type a unique name for the new policy in the dialog box that appears. You can then modify the copy.

Step 3 Add one or more rules to the file policy as described in Creating File Rules, on page 1433.

Step 4 Optionally, select Advanced and configure advanced options as described in Advanced and Archive File Inspection Options, on page 1420.

Step 5 Save the file policy.

What to do next

- If you are configuring policies for malware protection, see other required procedures in Configure File Policies, on page 1407.
- Otherwise:
  - Add the file policy to an access control rule as described in Add File Policies to Your Access Control Configuration, on page 1408.
  - Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Advanced and Archive File Inspection Options

The Advanced Settings in the file policy editor has the following general options:

- **First Time File Analysis**—Select this option to analyze first-seen files while AMP cloud disposition is pending. The file must match a rule configured to perform a malware cloud lookup and Spero, local malware, or dynamic analysis. If you deselect this option, files detected for the first time are marked with an Unknown disposition
- **Enable Custom Detection List**—Block files on the custom detection list.
- **Enable Clean List**—If enabled, this policy will allow files that are on the clean list.
- **Override AMP Cloud Disposition Based upon Threat Score**—Select an option:
  - If you select **Disabled,** the system will not override the disposition provided by the AMP Cloud.
  - If you set a threshold threat score, files with an AMP cloud verdict of Unknown are considered malware if their Dynamic Analysis score is equal to or worse than the threshold.
  - If you select a lower threshold value, you increase the number of files treated as malware. Depending on the action selected in your file policy, this can result in an increase of blocked files.
  - For numeric threat score ranges, see Threat Scores and Dynamic Analysis Summary Reports, on page 2312.

The Advanced Settings in the file policy editor has the following archive file inspection options:

- **Inspect Archives**—Enables inspection of the contents of archive files, for archive files as large as the **Maximum file size to store** advanced access control setting.

Caution Enabling or disabling **Inspect Archives** restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

- **Block Encrypted Archives**—Blocks archive files that have encrypted contents.
- **Block Uninspectable Archives**—Blocks archive files with contents that the system is unable to inspect for reasons other than encryption. This usually applies to corrupted files, or those that exceed your specified maximum archive depth.
- **Max Archive Depth**—Blocks nested archive files that exceed the specified depth. The top-level archive file is not considered in this count; depth begins at 1 with the first nested file .

Archive Files

Archive files are files that contain other files, such as `.zip` or `.rar` files.

If any individual file in an archive matches a file rule with a block action, the system blocks the entire archive, not just the individual file.

For details about options for archive file inspection, see Advanced and Archive File Inspection Options, on page 1420.

Archive Files That Can Be Inspected

- **File types**

  A complete list of inspectable archive file types appears in the FMC web interface on the file rule configuration page. To view that page, see Creating File Rules, on page 1433.

  Contained files that can be inspected appears in the same page.

- **File size**

  You can inspect archive files as large as the **Maximum file size to store** file policy advanced access control setting.

- **Nested archives**

  Archive files can contain other archive files, which can in turn contain archive files. The level at which a file is nested is its *archive file depth*. Note that the top-level archive file is not included in the depth count; depth begins at 1 with the first nested file.

  The system can inspect up to three levels of nested files beneath the outermost archive file (level 0). You can configure your file policy to block archive files that exceed that depth (or a lower maximum depth that you specify).

  If you choose not to block files that exceed the maximum archive file depth of 3, when archive files that contain some extractable contents and some contents nested at a depth of 3 or greater appear in monitored traffic, the system examines and reports data only for the files it was able to inspect.

All features applicable to uncompressed files (such as dynamic analysis and file storage) are available for nested files inside archive files.

- **Encrypted files**

  You can configure the system to block archives whose contents are encrypted or otherwise cannot be inspected.

- **Archives that are not inspected**

  If traffic that contains an archive file is on a Security Intelligence Block list or Do Not Block list, or if the top-level archive file's SHA-256 value is on the custom detection list, the system does not inspect the contents of the archive file.

  If a nested file is blocked, the entire archive is blocked; however, if a nested file is allowed, the archive is not automatically passed (depending on any other nested files and characteristics).

  .Exe files inside some .rar archives cannot be detected, including possibly rar5.

Archive File Dispositions

Archive file dispositions are based on the dispositions assigned to the files inside the archive. **All** archives that contain identified malware files receive a disposition of `Malware`. Archives without identified malware files receive a disposition of `Unknown` if they contain any unknown files, and a disposition of `Clean` if they contain only clean files.

Table 99: Archive File Disposition by Contents

Archive File Disposition	Number of Unknown Files	Number of Clean Files	Number of Malware Files
`Unknown`	1 or more	Any	0
`Clean`	0	1 or more	0
`Malware`	Any	Any	1 or more

Archive files, like other files, may have dispositions of `Custom Detection` or `Unavailable` if the conditions for those dispositions apply.

Viewing Archive Contents and Details

If your file policy is configured to inspect archive file contents, you can use the context menu in a table on pages under the Analysis > Files menu, and the network file trajectory viewer to view information about the files inside an archive when the archive file appears in a file event, malware event, or as a captured file.

All file contents of the archive are listed in table form, with a short summary of their relevant information: name, SHA-256 hash value, type, category, and archive depth. A network file trajectory icon appears by each file, which you can click to view further information about that specific file.

Override File Disposition Using Custom Lists

If a file has a disposition in the AMP cloud that you know to be incorrect, you can add the file's SHA-256 value to a file list that overrides the disposition from the cloud:

- To treat a file as if the AMP cloud assigned a clean disposition, add the file to the *clean list*.

- To treat a file as if the AMP cloud assigned a malware disposition, add the file to the *custom detection list.*

On subsequent detection, the device either allows or blocks the file without reevaluating the file's disposition. You can use the clean list or custom detection list per file policy.

Note To calculate a file's SHA-256 value, you must configure a rule in the file policy to either perform a malware cloud lookup or block malware on matching files.

For complete information about using file lists in Firepower, see File Lists, on page 438.

Alternatively, if applicable, use Centralized File Lists from AMP for Endpoints, on page 1423.

Centralized File Lists from AMP for Endpoints

If your organization has deployed AMP for Endpoints, Firepower can use Block and Allow lists created in AMP for Endpoints when it queries the AMP cloud for file dispositions.

Requirements:

- Your organization must be using the AMP public cloud.
- Your organization has deployed AMP for Endpoints.
- You have registered your Firepower system to AMP for Endpoints using the procedure in Integrate Firepower and AMP for Endpoints, on page 1437.

To create and deploy these lists, see the documentation or online help for AMP for Endpoints.

Note File lists created in Firepower override file lists created in AMP for Endpoints.

Managing File Policies

The File Policies page displays a list of existing file policies along with their last-modified dates. You can use this page to manage your file policies.

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Note The system checks for updates to the list of file types eligible for dynamic analysis (no more than once a day). If the list of eligible file types changes, this constitutes a change in the file policy; any access control policy using the file policy is marked out-of-date if deployed to any devices. You must deploy policies before the updated file policy can take effect on the device. See Maintain Your System: Update File Types Eligible for Dynamic Analysis, on page 1419.

Procedure

Step 1 Select **Policies** > **Access Control** > **Malware & File** .

Step 2 Manage your file policies:

- Compare—Click **Compare Policies**; see Comparing Policies, on page 355.
- Create — To create a file policy, click **New File Policy** and proceed as described in Create or Edit a File Policy, on page 1419.
- Copy — To copy a file policy, click **Copy** ( ).

  If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Delete — If you want to delete a file policy, click **Delete** ( ), then click **Yes** and **OK** as prompted.

  If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Deploy—Click **Deploy**; see Deploy Configuration Changes, on page 346.
- Edit — If you want to modify an existing file policy, click **Edit** ( ).
- Report—Click **Report** ( ); see Generating Current Policy Reports, on page 356.

File Rules

A file policy, like its parent access control policy, contains rules that determine how the system handles files that match the conditions of each rule. You can configure separate file rules to take different actions for different file types, application protocols, or directions of transfer.

For example, when a file matches a rule, the rule can:

- allow or block files based on simple file type matching
- block files based on disposition (whether or not evaluation indicates that it is malicious)
- store files to the device (For information, see Captured Files and File Storage, on page 1431)
- submit stored (captured) files for local malware, Spero, or dynamic analysis

In addition, the file policy can:

- automatically treat a file as if it is clean or malware based on entries in the clean list or custom detection list
- treat a file as if it is malware if the file's threat score exceeds a configurable threshold
- inspect the contents of archive files (such as `.zip` or `.rar`)

- block archive files whose contents are encrypted, nested beyond a specified maximum archive depth, or otherwise uninspectable

File Rule Components

Table 100: File Rule Components

File Rule Component	Description
application protocol	The system can detect and inspect files transmitted via FTP, HTTP, SMTP, IMAP, POP3, and NetBIOS-ssn (SMB). **Any**, the default, detects files in HTTP, SMTP, IMAP, POP3, FTP, and NetBIOS-ssn (SMB) traffic. To improve performance, you can restrict file detection to only one of those application protocols on a per-file rule basis.
direction of transfer	You can inspect incoming FTP, HTTP, IMAP, POP3, and NetBIOS-ssn (SMB) traffic for downloaded files; you can inspect outgoing FTP, HTTP, SMTP, and NetBIOS-ssn (SMB) traffic for uploaded files. **Tip** Use **Any** to detect files over multiple application protocols, regardless of whether users are sending or receiving.
file categories and types	The system can detect various types of files. These file types are grouped into basic categories, including multimedia (swf, mp3), executables (exe, torrent), and PDFs. You can configure file rules that detect individual file types, or on entire categories of file types. For example, you could block all multimedia files, or just ShockWave Flash (swf) files. Or, you could configure the system to alert you when a user downloads a BitTorrent (torrent) file. Note that executables include file types that can run macros and scripts, since these can contain malware. For a list of file types the system can inspect, select **Policies > Access Control > Malware & File**, create a temporary new file policy, then click **Add Rule**. Select a file type category and the file types that the system can inspect appear in the **File Types** list. **Note** Frequently triggered file rules can affect system performance. For example, detecting multimedia files in HTTP traffic (YouTube, for example, transmits significant Flash content) could generate an overwhelming number of events.

File Rule Component	Description
file rule action	A file rule's action determines how the system handles traffic that matches the conditions of the rule. Depending on the selected action, you can configure whether the system stores the file or performs Spero, local malware, or dynamic analysis on a file. If you select a Block action, you can also configure whether the system also resets the blocked connection. For descriptions of these actions and options, see File Rule Actions, on page 1426. File rules are evaluated in rule-action, not numerical, order. For details, see File Rule Actions: Evaluation Order, on page 1432.

File Rule Actions

File rules give you granular control over which file types you want to log, block, or scan for malware. Each file rule has an associated action that determines how the system handles traffic that matches the conditions of the rule. To be effective, a file policy must contain one or more rules. You can use separate rules within a file policy to take different actions for different file types, application protocols, or directions of transfer.

File Rule Actions

- *Detect Files* rules allow you to log the detection of specific file types to the database, while still allowing their transmission.
- *Block Files* rules allow you to block specific file types. You can configure options to reset the connection when a file transfer is blocked, and store captured files to the managed device.
- *Malware Cloud Lookup* rules allow you to obtain and log the disposition of files traversing your network, while still allowing their transmission.
- *Block Malware* rules allow you to calculate the SHA-256 hash value of specific file types, query the AMP cloud to determine if files traversing your network contain malware, then block files that represent threats.

File Rule Action Options

Depending on the action you select, you have different options:

File Rule Action Option	Block Files capable?	Block Malware capable?	Detect Files capable?	Malware Cloud Lookup capable?
Spero Analysis* for MSEXE	no	yes, you can submit executable files	no	yes, you can submit executable files
Dynamic Analysis*	no	yes, you can submit executable files with Unknown file dispositions	no	yes, you can submit executable files with Unknown file dispositions
Capacity Handling	no	yes	no	yes

File Rule Action Option	Block Files capable?	Block Malware capable?	Detect Files capable?	Malware Cloud Lookup capable?
Local Malware Analysis*	no	yes	no	yes
Reset Connection	yes (recommended)	yes (recommended)	no	no
Store files	yes, you can store all matching file types	yes, you can store file types matching the file dispositions you select	yes, you can store all matching file types	yes, you can store file types matching the file dispositions you select

* For complete information about these options, see Malware Protection Options (in File Rule Actions), on page 1427 and its subtopics.



Caution Selecting **Detect Files** or **Block Files**, enabling or disabling **Store files** in a **Detect Files** or **Block Files** rule, or adding the first or removing the last file rule that combines the **Malware Cloud Lookup** or **Block Malware** file rule action with an analysis option (**Spero Analysis or MSEXE**, **Dynamic Analysis**, or **Local Malware Analysis**) or a store files option (**Malware, Unknown**, **Clean**, or **Custom**), restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Malware Protection Options (in File Rule Actions)

The Firepower system applies several methods of file inspection and analysis to determine whether a file contains malware.

Depending on the options you enable in a file rule, the system inspects files using the following tools, in order:

1. Spero Analysis, on page 1429 and AMP Cloud Lookup, on page 1429
2. Local Malware Analysis, on page 1429
3. Dynamic Analysis, on page 1430

For a comparison of these tools, see Comparison of Malware Protection Options, on page 1427.

(You can also, if you choose, block all files based on their file type. For more information, see Block All Files by Type, on page 1432.)

See also information about Cisco's AMP for Endpoints product at (Optional) Malware Protection with AMP for Endpoints, on page 1434 and subtopics.

Comparison of Malware Protection Options

The following table details the benefits and drawbacks of each type of file analysis, as well as the way each malware protection method determines a file's disposition.

Analysis Type	Benefit	Limitations	Malware Identification
Spero analysis	Structural analysis of executable files, submits Spero signature to the AMP Cloud for analysis	Less thorough than local malware analysis or dynamic analysis, only for executable files	Disposition changes from Unknown to Malware only on positive identification of malware.
Local malware analysis	Consumes fewer resources than dynamic analysis, and returns results more quickly, especially if the detected malware is common	Less thorough results than dynamic analysis	Disposition changes from Unknown to Malware only on positive identification of malware.
Dynamic analysis	Thorough analysis of unknown files using Cisco Threat Grid	Eligible files are uploaded to the public cloud or an on-premises appliance. It takes some time to complete analysis	Threat score determines maliciousness of a file. Disposition can be based on the threat score threshold configured in the file policy.
Spero analysis and local malware analysis	Consumes fewer resources than configuring local malware analysis and dynamic analysis, while still using AMP cloud resources to identify malware	Less thorough than dynamic analysis, Spero analysis only for executable files	Disposition changes from Unknown to Malware only on positive identification of malware.
Spero analysis and dynamic analysis	Uses full capabilities of AMP cloud in submitting files and Spero signatures	Results obtained less quickly than if using local malware analysis	Threat score changes based on dynamic analysis results for files preclassified as possible malware. Disposition changes based on configured threat score threshold in the file policy, and from Unknown to Malware if the Spero analysis identifies malware.
Local malware analysis and dynamic analysis	Thorough results in using both types of file analysis	Consumes more resources than either alone	Threat score changes based on dynamic analysis results for files preclassified as possible malware. Disposition changes from Unknown to Malware if local malware analysis identifies malware, or based on configured threat score threshold in the file policy.

Analysis Type	Benefit	Limitations	Malware Identification
Spero analysis, local malware analysis and dynamic analysis	Most thorough results	Consumes most resources in running all three types of file analysis	Threat score changes based on dynamic analysis results for files preclassified as possible malware. Disposition changes from Unknown to Malware if Spero analysis or local malware analysis identifies malware, or based on configured threat score threshold in the file policy.
(Block transmission of all files of a specified file type)	Does not require a Malware license (This option is not technically a malware protection option.)	Legitimate files will also be blocked	(No analysis is performed.)

Note Preclassification does not itself determine a file's disposition; it is merely one of the factors that determine whether a file is eligible for Dynamic Analysis.

Spero Analysis

Spero analysis examines structural characteristics such as metadata and header information in executable files. After generating a Spero signature based on this information, if the file is an eligible executable file, the device submits it to the Spero heuristic engine in the AMP cloud. Based on the Spero signature, the Spero engine determines whether the file is malware. You can also configure rules to submit files for Spero analysis without also submitting them to the AMP cloud.

Note that you cannot manually submit files for Spero analysis.

AMP Cloud Lookup

For files that are eligible for assessment using Advanced Malware Protection, the Firepower Management Center performs a *malware cloud lookup*, querying the AMP cloud for the file's disposition based on its SHA-256 hash value.

To improve performance, the system caches dispositions returned by the cloud and uses the cached disposition for known files rather than querying the AMP cloud. For more information about this cache, see Cached Disposition Longevity, on page 1430.

Local Malware Analysis

Local malware analysis allows a managed device to locally inspect executables, PDFs, office documents, and other types of files for the most common types of malware, using a detection rule set provided by the Cisco Talos Intelligence Group (Talos). Because local analysis does not query the AMP cloud, and does not run the file, local malware analysis saves time and system resources.

If the system identifies malware through local malware analysis, it updates the existing file disposition from Unknown to Malware. The system then generates a new malware event. If the system does not identify malware, it does not update the file disposition from Unknown to Clean. After the system runs local malware

analysis, it caches file information such as SHA-256 hash value, timestamp, and disposition, so that if detected again within a certain period of time, the system can identify malware without additional analysis. For more information about the cache, see Cached Disposition Longevity, on page 1430.

Local malware analysis does not require establishing communications with the Cisco Threat Grid cloud. However, you must configure communications with the cloud to submit files for dynamic analysis, and to download updates to the local malware analysis ruleset.

Cached Disposition Longevity

Dispositions returned from an AMP cloud query, associated threat scores, and dispositions assigned by local malware analysis, have a time-to-live (TTL) value. After a disposition has been held for the duration specified in the TTL value without update, the system purges the cached information. Dispositions and associated threat scores have the following TTL values:

- Clean — 4 hours
- Unknown — 1 hour
- Malware — 1 hour

If a query against the cache identifies a cached disposition that timed out, the system re-queries the local malware analysis database and the AMP cloud for a new disposition.

Dynamic Analysis

You can configure your file policy to automatically submit files for dynamic analysis using Cisco Threat Grid (formerly AMP Threat Grid), Cisco's file analysis and threat intelligence platform.

Devices submit eligible files to Cisco Threat Grid (either the public cloud or to an on-premises appliance, whichever you have specified) regardless of whether the device stores the file.

Cisco Threat Grid runs the file in a sandbox environment, analyzes the file's behavior to determine whether the file is malicious, and returns a threat score that indicates the likelihood that a file contains malware. From the threat score, you can view a dynamic analysis summary report with the reasons for the assigned threat score. You can also look in Cisco Threat Grid to view detailed reports for files that your organization submitted, as well as scrubbed reports with limited data for files that your organization did not submit.

For more information about Cisco Cisco Threat Grid, see https://www.cisco.com/c/en/us/products/security/threat-grid/index.html

To configure your system to perform dynamic analysis, see the topics under Dynamic Analysis Connections, on page 1416.

Which Files Are Eligible for Dynamic Analysis?

A file's eligibility for dynamic analysis depends on:

- the file type
- the file size
- the file rule's action

Additionally:

- The system submits only files that match the file rules you configure.

- The file must have a malware cloud lookup disposition of Unknown or Unavailable at the time the file is sent for analysis.
- The system must preclassify the file as potential malware.

Dynamic Analysis and Capacity Handling

Capacity handling allows you to temporarily store files that are otherwise eligible for dynamic analysis if the system is temporarily unable to submit files to the cloud, either because the device cannot communicate with the cloud or because the maximum number of submissions has been reached. The system submits the stored files when the hindering condition has passed.

Some devices can store files on the device hard drive or in a malware storage pack. See also Malware Storage Pack, on page 1432.

Captured Files and File Storage

The file storage feature allows you to capture selected files detected in traffic, and automatically store a copy of the file temporarily to a device's hard drive, or, if installed, to the malware storage pack.

After your device captures the files, you can:

- Store captured files on the device's hard drive for later analysis.
- Download the stored file to a local computer for further manual analysis or archival purposes.
- Manually submit eligible captured files for AMP cloud lookup or dynamic analysis.

Note that once a device stores a file, it will not re-capture it if the file is detected in the future and the device still has that file stored.



Note When a file is detected for the first time on your network, you can generate a file event that represents the file's detection. However, if your file rule performs a malware cloud lookup, the system requires additional time to query the AMP cloud and return a disposition. Due to this delay, the system cannot store this file until the second time it is seen on your network, and the system can immediately determine the file's disposition.

Whether the system captures or stores a file, you can:

- Review information about the captured file from Analysis > Files > Captured Files, including whether the file was stored or submitted for dynamic analysis, file disposition, and threat score, allowing you to quickly review possible malware threats detected on your network.
- View the file's trajectory to determine how it traversed your network and which hosts have a copy.
- Add the file to the clean list or custom detection list to always treat the file as if it had a clean or malware disposition on future detection.

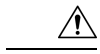
You configure file rules in a file policy to capture and store files of a specific type, or with a particular file disposition, if available. After you associate the file policy with an access control policy and deploy it to your devices, matching files in traffic are captured and stored. You can also limit the minimum and maximum file sizes to store.

Stored files are not included in system backups.

You can view captured file information under Analysis > Files > Captured Files, and download a copy for offline analysis.

Malware Storage Pack

Based on your file policy configuration, your device may store a substantial amount of file data to the hard drive. You can install a malware storage pack in the device; the system stores files to the malware storage pack, allowing more room on the primary hard drive to store events and configuration files. The system periodically deletes older files. If the device's primary hard drive does not have enough available space, and does not have an installed malware storage pack, you cannot store files.



Caution Do not attempt to install a hard drive that was not supplied by Cisco in your device. Installing an unsupported hard drive may damage the device. Malware storage pack kits are available for purchase **only** from Cisco, and are for use **only** with 8000 Series devices. Contact Support if you require assistance with the malware storage pack. See the *Firepower System Malware Storage Pack Guide* for more information.

Without a malware storage pack installed, when you configure a device to store files, it allocates a set portion of the primary hard drive's space to captured file storage. If you configure capacity handling to temporarily store files for dynamic analysis, the system uses the same hard drive allocation to store these files until it can resubmit them to the cloud.

When you install a malware storage pack in a device and configure file storage or capacity handling, the device allocates the entire malware storage pack for storing these files. The device cannot store any other information on the malware storage pack.

When the allocated space for captured file storage fills to capacity, the system deletes the oldest stored files until the allocated space reaches a system-defined threshold. Based on the number of files stored, you may see a substantial drop in disk usage after the system deletes files.

If a device has already stored files when you install a malware storage pack, the next time you restart the device, any captured files or capacity handling files stored on the primary hard drive are moved to the malware storage pack. Any future files the device stores are stored to the malware storage pack.

Block All Files by Type

If your organization wants to block not only the transmission of malware files, but all files of a specific type, regardless of whether the files contain malware, you can do so.

File control is supported for all file types where the system can detect malware, plus many additional file types. These file types are grouped into basic categories, such as multimedia (swf, mp3), executables (exe, torrent), and PDFs.

Blocking all files based on their type is not technically a malware protection feature; it does not require a Malware license and does not query the AMP cloud.

File Rule Actions: Evaluation Order

A file policy will likely contain multiple rules with different actions for different situations. If more than one rule can apply to a particular situation, the evaluation order described in this topic applies. In general, simple blocking takes precedence over malware inspection and blocking, which takes precedence over simple detection and logging.

The order of precedence of file-rule actions is:

- *Block Files*

- *Block Malware*
- *Malware Cloud Lookup*
- *Detect Files*

Creating File Rules

Caution Selecting **Detect Files** or **Block Files**, enabling or disabling **Store files** in a **Detect Files** or **Block Files** rule, or adding the first or removing the last file rule that combines the **Malware Cloud Lookup** or **Block Malware** file rule action with an analyis option (**Spero Analysis or MSEXE**, **Dynamic Analysis**, or **Local Malware Analysis**) or a store files option (**Malware, Unknown**, **Clean**, or **Custom**), restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Before you begin

If you are configuring rules for malware protection, see Configure File Policies, on page 1407.

Procedure

Step 1 In the file policy editor, click **Add File Rule**.

Step 2 Select an **Application Protocol** and **Direction of Transfer** as described in File Rule Components, on page 1425.

Step 3 Select one or more **File Types**.

The file types you see depend on the selected application protocol, direction of transfer, and action.

You can filter the list of file types in the following ways:

- Select one or more **File Type Categories**, then click **All types in selected Categories.**
- Search for a file type by its name or description. For example, type **`Windows`** in the **Search name and description** field to display a list of Microsoft Windows-specific files.

Tip Hover your pointer over a file type to view its description.

Step 4 Select a file rule **Action** as described in File Rule Actions, on page 1426, with consideration for File Rule Actions: Evaluation Order, on page 1432.

The actions available to you depend on the licenses you have installed. See License Requirements for File and Malware Policies, on page 1403.

Step 5 Depending on the action you selected, configure options:

- reset the connection after blocking the file
- store files that match the rule
- enable Spero analysis*

- enable local malware analysis*
- enable dynamic analysis* and capacity handling

* For information about these options, see File Rule Actions, on page 1426 and Malware Protection Options (in File Rule Actions), on page 1427 and its subtopics.

Step 6 Click **Add.**

Step 7 Click **Save** to save the policy.

What to do next

- If you are configuring policies for malware protection, return to Configure File Policies, on page 1407.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Access Control Rule Logging for Malware Protection

When the system detects a prohibited file (including malware) according to the settings in the file policy, it automatically logs an event to the Firepower Management Center database. If you do not want to log file or malware events, you can disable this logging on a per-access-control-rule basis.

The system also logs the end of the associated connection to the Firepower Management Center database, regardless of the logging configuration of the invoking access control rule.

Retrospective Disposition Changes

File dispositions can change. For example, as new information is discovered, the AMP cloud can determine that a file that was previously thought to be clean is now identified as malware, or the reverse—that a malware-identified file is actually clean. When the disposition changes for a file you queried in the past week, the AMP cloud notifies the system so it can automatically take action the next time it detects that file being transmitted. A changed disposition is called a *retrospective* disposition.

(Optional) Malware Protection with AMP for Endpoints

Cisco's AMP for Endpoints is a separate malware-protection product that can supplement malware protection provided by the Firepower system and be integrated with your Firepower deployment.

AMP for Endpoints is Cisco’s enterprise-class Advanced Malware Protection solution that runs as a lightweight connector on individual users' *endpoints* (computers and mobile devices) to discover, understand, and block advanced malware outbreaks, advanced persistent threats, and targeted attacks.

Benefits of AMP for Endpoints include:

- configure custom malware detection policies and profiles for your entire organization, as well as perform flash and full scans on all your users’ files
- perform malware analysis, including view heat maps, detailed file information, network file trajectory, and threat root causes

- configure multiple aspects of outbreak control, including automatic quarantines, application blocking to stop non-quarantined executables from running, and exclusion lists
- create custom protections, block execution of certain applications based on group policy, and create custom Allowed Applications lists
- use the AMP for Endpoints management console to help you mitigate the effect of malware. The management console provides a robust, flexible web interface where you control all aspects of your AMP for Endpoints deployment and manage all phases of an outbreak.

For detailed information about AMP for Endpoints, see:

- https://www.cisco.com/c/en/us/products/security/amp-for-endpoints/index.html.
- Online help in the AMP for Endpoints management console.
- AMP for Endpoints documentation available from: http://docs.amp.cisco.com.

Comparison of Malware Protection: Firepower vs. AMP for Endpoints

Table 101: Advanced Malware Protection Differences by Detecting Product

Feature	Firepower Malware Protection (AMP for Networks)	AMP for Endpoints
File type detection and blocking method (file control)	In network traffic, using access control and file policies	Not supported
Malware detection and blocking method	In network traffic, using access control and file policies	On individual endpoints (end-user computers and mobile devices), using a connector that communicates with the AMP cloud
Network traffic inspected	Traffic passing through a managed device	None; connectors installed on endpoints directly inspect files
Malware intelligence data source	AMP cloud (public or private)	AMP cloud (public or private)
Malware detection robustness	Limited file types	All file types
Malware analysis choices	FMC-based, plus analysis in the AMP cloud	FMC-based, plus additional options on the AMP for Endpoints management console
Malware mitigation	Malware blocking in network traffic, FMC-initiated remediations	AMP for Endpoints-based quarantine and outbreak control options, FMC-initiated remediations
Events generated	File events, captured files, malware events, and retrospective malware events	Malware events
Information in malware events	Basic malware event information, plus connection data (IP address, port, and application protocol)	In-depth malware event information; no connection data

Feature	Firepower Malware Protection (AMP for Networks)	AMP for Endpoints
Network file trajectory	FMC-based	FMC and the AMP for Endpoints management console each have a network file trajectory. Both are useful.
Required licenses or subscriptions	Licenses required to perform file control and AMP for Networks	AMP for Endpoints subscription. No license is required to bring AMP for Endpoints data into FMC.

About Integrating Firepower with AMP for Endpoints

If your organization has deployed AMP for Endpoints, you can optionally integrate that product with your Firepower deployment.

Integration with AMP for Endpoints does not require a dedicated Firepower license.

Benefits of Integrating Firepower and AMP for Endpoints

Integrating your AMP for Endpoints deployment with your Firepower system offers the following benefits:

- Centralized Blocked Applications and Allowed Applications lists configured in AMP for Endpoints can determine verdicts for file SHAs sent from Firepower to the AMP cloud for disposition.

  See Centralized File Lists from AMP for Endpoints, on page 1423.

- The system can import malware events detected by AMP for Endpoints into Firepower Management Center so you can manage these events along with malware events generated by the Firepower system. Imported data for these events includes scans, malware detections, quarantines, blocked executions, and cloud recalls, as well as indications of compromise (IOCs) that FMC displays for hosts that it monitors.

  For more information, see Malware Event Analysis with AMP for Endpoints, on page 2298.

- You can view file trajectory and other details in the AMP for Endpoints console.

  For details, see Work with Event Data in the AMP for Endpoints Console, on page 2323.



Important If you use a Cisco AMP Private Cloud, see limitations at AMP for Endpoints and AMP Private Cloud, on page 1436.

AMP for Endpoints and AMP Private Cloud

If you configure a Cisco AMP private cloud to collect AMP endpoint data on your network, all AMP for Endpoints connectors send data to the private cloud, which forwards that data to the Firepower Management Center. The private cloud does not share any of your endpoint data over an external connection.

If your organization has deployed an AMP private cloud, all connections to the AMP cloud funnel through the private cloud, which acts as an anonymized proxy to ensure the security and privacy of your monitored network. This includes importing AMP for Endpoints data. The private cloud does not share any of your endpoint data over an external connection.

The following integration features are not available if you use an AMP private cloud:

- Use of Blocked Applications and Allowed Applications lists configured in AMP for Endpoints. (These lists are used to block or allow files.)
- Visibility in AMP for Endpoints of malware events generated from Firepower.

You can configure multiple private clouds to support the capacity you require.

Integrate Firepower and AMP for Endpoints

If your organization has deployed Cisco's AMP for Endpoints product, you can integrate that application with Firepower to achieve the benefits described in Benefits of Integrating Firepower and AMP for Endpoints, on page 1436.

When you integrate with AMP for Endpoints, you must configure an AMP for Endpoints connection even if you already have an AMP for Networks (AMP for Firepower) connection configured. You can configure multiple AMP for Endpoints cloud connections.

Caution In a multidomain deployment, configure AMP for Endpoints connections at the leaf level only, especially if your leaf domains have overlapping IP space. If multiple subdomains have hosts with the same IP-MAC address pair, the system could save malware events that are generated by AMP for Endpoints to the wrong leaf domain, or associate IOCs with the wrong hosts.

However, you can configure AMP for Endpoints connections at any domain level, provided you use a separate AMP for Endpoints account for each connection. For example, each client of an MSSP might have its own AMP for Endpoints deployment.

Note An AMP for Endpoints connection that has not registered successfully does not affect AMP for Networks.

Before you begin

- You must be an Admin user to perform this task.
- If your deployment uses Cisco AMP Private Cloud, see limitations at AMP for Endpoints and AMP Private Cloud, on page 1436.
- AMP for Endpoints must be set up and working properly on your network.
- The Firepower Management Center must have direct access to the Internet.
- Make sure your FMC and AMP for Endpoints can communicate with each other. See the topics under Security, Internet Access, and Communication Ports, on page 2425.
- If you are connecting to the AMP cloud after either restoring your Firepower Management Center to factory defaults or reverting to a previous version, use the AMP for Endpoints management console to remove the previous connection.
- You will need your AMP for Endpoints credentials to log in to the AMP for Endpoints console during this procedure.

Procedure

Step 1 Choose **AMP** > **AMP Management**.

Step 2 Click **Add AMP Cloud Connection**.

Step 3 From the **Cloud Name** drop-down list, choose the cloud you want to use:

- The AMP cloud closest to the geographical location of your Firepower Management Center.
- For AMP private cloud (AMPv), choose **Private Cloud** and proceed as described in Cisco AMP Private Cloud, on page 1413.

Step 4 If you want to use this cloud for both AMP for Networks and AMP for Endpoints, select the **Use for AMP for Firepower** check box.

If you configured a different cloud to handle AMP for Networks (AMP for Firepower) communications, you can clear this check box; if this is your only AMP cloud connection, you cannot.

In a multidomain deployment, this check box appears only in the Global domain. Each Firepower Management Center can have only one AMP for Networks connection.

Step 5 Click **Register**.

A spinning state icon indicates that a connection is pending, for example, after you configure a connection on the Firepower Management Center, but before you authorize it using the AMP for Endpoints management console. A **Denied** () indicates that the cloud denied the connection or the connection failed for another reason.

Step 6 Confirm that you want to continue to the AMP for Endpoints management console, then log into the management console.

Step 7 Using the management console, authorize the AMP cloud to send AMP for Endpoints data to the Firepower Management Center.

Step 8 If you want to restrict the data that the FMC receives, select specific groups within your organization for which you want to receive information.

By default, the AMP cloud sends data for all groups. To manage groups, choose **Management** > **Groups** on the AMP for Endpoints management console. For detailed information, see the management console online help.

Step 9 Click **Allow** to enable the connection and start the transfer of data.

Clicking **Deny** returns you to the Firepower Management Center, where the connection is marked as denied. If you navigate away from the Applications page on the AMP for Endpoints management console, and neither deny nor allow the connection, the connection is marked as pending on the Firepower Management Center's web interface. The health monitor does **not** alert you of a failed connection in either of these situations. If you want to connect to the AMP cloud later, delete the failed or pending connection, then recreate it.

Incomplete registration of an AMP for Endpoints connection does not disable the AMP for Networks connection.

Step 10 To verify that the connection is correctly configured:

a) On the **AMP** > **AMP Management** page, click the Cloud Name that includes **`AMP for Endpoints`** in the **Cisco AMP Solution Type** column.
b) In the AMP for Endpoints console window that displays, choose **Accounts** > **Applications**.
c) Verify that your Firepower Management Center is on the list.

d) In the AMP for Endpoints console window, choose **Manage > Computers**.
e) Verify that your Firepower Management Center is on the list.

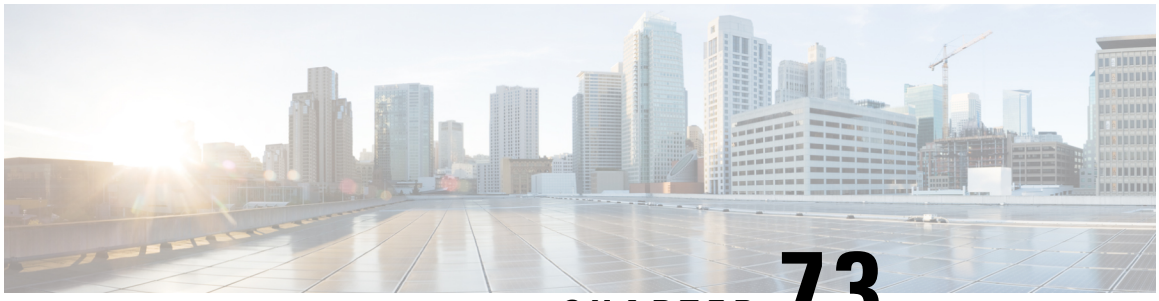
What to do next

- In the AMP for Endpoints console window, configure settings as needed. For example, define group membership for your management center and assign policies. For information, see the AMP for Endpoints online help or other documentation.
- In high availability configurations, you must configure AMP cloud connections independently on the Active and Standby instances of the Firepower Management Center; these configurations are not synchronized.
- The default health policy warns you if the Firepower Management Center cannot connect to the AMP for Endpoints portal after an initial successful connection, or if the connection is deregistered using the AMP portal.

  Verify that the **AMP for Endpoints Status** monitor is enabled under **System > Health > Policy**.

History for File Policies and Malware Protection

Feature	Version	Details
Chapter restructure	any version that is republished	Restructured this chapter's content to reduce confusion. Some content was moved to or from the chapter for File/Malware Events and Network File Trajectory on page 2295.
Moved URL Filtering information to the new URL Filtering chapter	6.3	Moved information about configuring cloud communications for URL Filtering to the new URL Filtering chapter. Made related changes to the structure of the Cisco CSI topics in the chapter.



CHAPTER 73

File and Malware Inspection Performance and Storage Tuning

The following topics describe how to configure file and malware inspection performance and storage:

- File and Malware Inspection Performance and Storage Options, on page 1441
- Tuning File and Malware Inspection Performance and Storage, on page 1443

File and Malware Inspection Performance and Storage Options

Increasing the file sizes can affect the performance of the system.



Caution Configuring a non-default value under Files and Malware Settings restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Table 102: Advanced Access Control File and AMP for Networks Options

Field	Description	Guidelines and Restrictions
Limit the number of bytes inspected when doing file type detection	Specifies the number of bytes inspected when performing file type detection.	0 - 4294967295 (4GB) 0 removes the restriction. The default value is the maximum segment size of a TCP packet (1460 bytes). In most cases, the system can identify common file types using the first packet. To detect ISO files, enter a value greater than 36870.

Field	Description	Guidelines and Restrictions
Allow file if cloud lookup for Block Malware takes longer than (seconds)	Specifies how long the system will hold the last byte of a file that matches a **Block Malware** rule and that does not have a cached disposition, while malware cloud lookup occurs. If the time elapses without the system obtaining a disposition, the file passes. Dispositions of Unavailable are not cached.	0 - 30 seconds Do *not* set this option to 0 without contacting Support. Cisco recommends that you use the default value to avoid blocking traffic because of connection failures.
Do not calculate SHA-256 hash values for files larger than (in bytes)	Prevents the system from storing files larger than a certain size, performing a malware cloud lookup on the files, or blocking the files if added to the custom detection list.	0 - 4294967295 (4GB) 0 removes the restriction. This value must be greater than or equal to **Maximum file size to store (bytes)** and **Maximum file size for dynamic analysis testing (bytes)**.

Minimum file size to store (bytes)	These settings specify:
	• The file size that the system can inspect using the following detectors:
	  • Spero analysis
	  • Sandboxing and preclassification
	  • Local malware analysis/ClamAV
	  • Archive inspection
	• The file size that the system can store using a file rule.
		0 - 10485760 (10MB)
		0 disables file storage.
		Must be less than or equal to **Maximum file size to store (bytes)** and **Do not calculate SHA-256 hash values for files larger than (in bytes)**.

Maximum file size to store (bytes)		0 - 10485760 (10MB) 0 disables file storage. Must be greater than or equal to **Minimum file size to store (bytes)**, and less than or equal to **Do not calculate SHA-256 hash values for files larger than (in bytes)**.
Minimum file size for dynamic analysis testing (bytes)	Specifies the minimum file size the system can submit to the AMP cloud for dynamic analysis.	0 -10485760 (10MB) Must be less than or equal to **Maximum file size for dynamic analysis testing (bytes)** and **Do not calculate SHA-256 hash values for files larger than (in bytes)**. The file size for dynamic analysis must be within the limits defined by the minimum and maximum settings for file analysis. The system checks the AMP cloud for updates to the minimum file size you can submit (no more than once a day). If the new minimum size is larger than your current value, your current value is updated to the new minimum, and your policy is marked out-of-date.

Field	Description	Guidelines and Restrictions
Maximum file size for dynamic analysis testing (bytes)	Specifies the maximum file size the system can submit to the AMP cloud for dynamic analysis.	0 -10485760 (10MB) Must be greater than or equal to **Minimum file size for dynamic analysis testing (bytes)**, and less than or equal to **Do not calculate SHA-256 hash values for files larger than (in bytes)**. The file size for dynamic analysis must be within the limits defined by the minimum and maximum settings for file analysis. The system checks the AMP cloud for updates to the maximum file size you can submit (no more than once a day). If the new maximum size is smaller than your current value, your current value is updated to the new maximum, and your policy is marked out-of-date.

Tuning File and Malware Inspection Performance and Storage

You must be an Admin, Access Admin, or Network Admin user to perform this task.

Caution Configuring a non-default value under Files and Malware Settings restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Procedure

Step 1 In the access control policy editor, click **Advanced Settings**.

Step 2 Click **Edit** ( ) next to **Files and Malware Settings**.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Set any of the options described in File and Malware Inspection Performance and Storage Options, on page 1441.

Step 4 Click **OK**.

Step 5 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

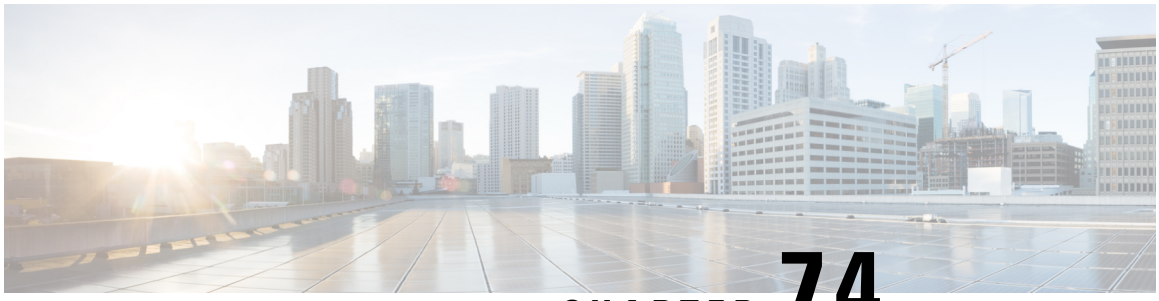
Snort® Restart Scenarios, on page 348



PART XVIII

Intrusion Detection and Prevention

- An Overview of Intrusion Detection and Prevention, on page 1447
- Layers in Intrusion and Network Analysis Policies, on page 1463
- Getting Started with Intrusion Policies, on page 1479
- Tuning Intrusion Policies Using Rules, on page 1489
- Tailoring Intrusion Protection to Your Network Assets, on page 1517
- Sensitive Data Detection, on page 1523
- Globally Limiting Intrusion Event Logging, on page 1537
- The Intrusion Rules Editor, on page 1543
- Intrusion Prevention Performance Tuning, on page 1655



CHAPTER 74

An Overview of Intrusion Detection and Prevention

The following topics provide an overview of network analysis and intrusion policies:

- Network Analysis and Intrusion Policy Basics, on page 1447
- How Policies Examine Traffic For Intrusions, on page 1448
- System-Provided and Custom Network Analysis and Intrusion Policies, on page 1453
- License Requirements for Network Analysis and Intrusion Policies, on page 1459
- Requirements and Prerequisites for Network Analysis and Intrusion Policies, on page 1459
- The Navigation Panel: Network Analysis and Intrusion Policies, on page 1459
- Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

Network Analysis and Intrusion Policy Basics

Network analysis and intrusion policies work together as part of the Firepower System's intrusion detection and prevention feature.

- The term *intrusion detection* generally refers to the process of passively monitoring and analyzing network traffic for potential intrusions and storing attack data for security analysis. This is sometimes referred to as "IDS."
- The term *intrusion prevention* includes the concept of intrusion detection, but adds the ability to block or alter malicious traffic as it travels across your network. This is sometimes referred to as "IPS."

In an intrusion prevention deployment, when the system examines packets:

- A **network analysis policy** governs how traffic is *decoded* and *preprocessed* so it can be further evaluated, especially for anomalous traffic that might signal an intrusion attempt.
- An **intrusion policy** uses *intrusion and preprocessor rules* (sometimes referred to collectively as *intrusion rules*) to examine the decoded packets for attacks based on patterns. Intrusion policies are paired with *variable sets*, which allow you to use named values to accurately reflect your network environment.

Both network analysis and intrusion policies are invoked by a parent access control policy, but at different times. As the system analyzes traffic, the network analysis (decoding and preprocessing) phase occurs before and separately from the intrusion prevention (additional preprocessing and intrusion rules) phase. Together, network analysis and intrusion policies provide broad and deep packet inspection. They can help you detect,

alert on, and protect against network traffic that could threaten the availability, integrity, and confidentiality of hosts and their data.

The Firepower System is delivered with several similarly named network analysis and intrusion policies (for example, Balanced Security and Connectivity) that complement and work with each other. By using system-provided policies, you can take advantage of the experience of the Cisco Talos Intelligence Group (Talos). For these policies, Talos sets intrusion and preprocessor rule states, as well as provides the initial configurations for preprocessors and other advanced settings.

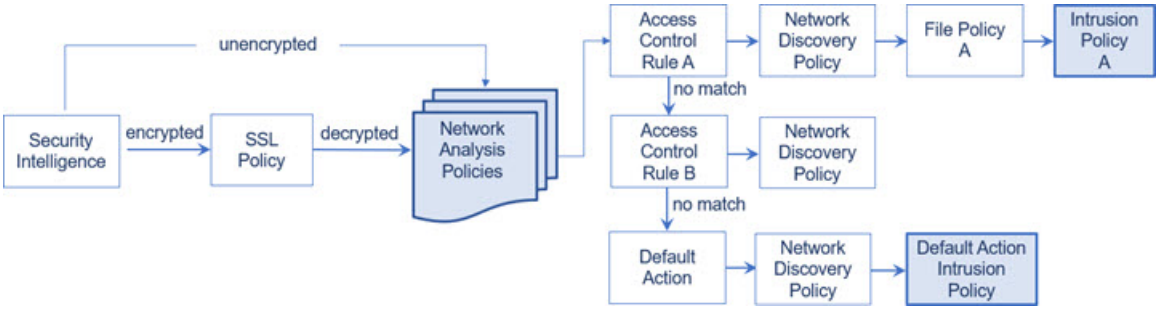
You can also create custom network analysis and intrusion policies. You can tune settings in custom policies to inspect traffic in the way that matters most to you so that you can improve both the performance of your managed devices and your ability to respond effectively to the events they generate.

You create, edit, save, and manage network analysis and intrusion policies using similar policy editors in the web interface. When you are editing either type of policy, a navigation panel appears on the left side of the web interface; the right side displays various configuration pages.

How Policies Examine Traffic For Intrusions

When the system analyzes traffic as part of your access control deployment, the network analysis (decoding and preprocessing) phase occurs before and separately from the intrusion prevention (intrusion rules and advanced settings) phase.

The following diagram shows, in a simplified fashion, the order of traffic analysis in an inline, intrusion prevention and AMP for Networks deployment. It illustrates how the access control policy invokes other policies to examine traffic, and in which order those policies are invoked. The network analysis and intrusion policy selection phases are highlighted.



In an inline deployment (that is, where relevant configurations are deployed to devices using routed, switched, or transparent interfaces, or inline interface pairs), the system can block traffic without further inspection at almost any step in the illustrated process. Security Intelligence, the SSL policy, network analysis policies, file policies, and intrusion policies can all either drop or modify traffic. Only the network discovery policy, which passively inspects packets, cannot affect the flow of traffic.

Similarly, at each step of the process, a packet could cause the system to generate an event. Intrusion and preprocessor events (sometimes referred to collectively as *intrusion events*) are indications that a packet or its contents may represent a security risk.

Tip The diagram does not reflect that access control rules handle encrypted traffic when your SSL inspection configuration allows it to pass, or if you do not configure SSL inspection. By default, the system disables intrusion and file inspection of encrypted payloads. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has intrusion and file inspection configured.

Note that for a single connection, although the system selects a network analysis policy before an access control rule as shown in the diagram, some preprocessing (notably application layer preprocessing) occurs after access control rule selection. This does **not** affect how you configure preprocessing in custom network analysis policies.

Decoding, Normalizing, and Preprocessing: Network Analysis Policies

Without decoding and preprocessing, the system could not appropriately evaluate traffic for intrusions because protocol differences would make pattern matching impossible. Network analysis policies govern these traffic-handling tasks:

- **after** traffic is filtered by Security Intelligence
- **after** encrypted traffic is decrypted by an optional SSL policy
- **before** traffic can be inspected by file or intrusion policies

A network analysis policy governs packet processing in phases. First the system decodes packets through the first three TCP/IP layers, then continues with normalizing, preprocessing, and detecting protocol anomalies:

- The packet decoder converts packet headers and payloads into a format that can be easily used by the preprocessors and later, intrusion rules. Each layer of the TCP/IP stack is decoded in turn, beginning with the data link layer and continuing through the network and transport layers. The packet decoder also detects various anomalous behaviors in packet headers.
- In inline deployments, the inline normalization preprocessor reformats (normalizes) traffic to minimize the chances of attackers evading detection. It prepares packets for examination by other preprocessors and intrusion rules, and helps ensure that the packets the system processes are the same as the packets received by the hosts on your network.

Note In a passive deployment, Cisco recommends that you enable adaptive profile updates at the access control policy level, instead of inline normalization at the network analysis level.

- Various network and transport layers preprocessors detect attacks that exploit IP fragmentation, perform checksum validation, and perform TCP and UDP session preprocessing.

  Note that some advanced transport and network preprocessor settings apply globally to all traffic handled by the target devices of an access control policy. You configure these in the access control policy rather than in a network analysis policy.

- Various application-layer protocol decoders normalize specific types of packet data into formats that the intrusion rules engine can analyze. Normalizing application-layer protocol encodings allows the system to effectively apply the same content-related intrusion rules to packets whose data is represented differently, and to obtain meaningful results.

- The Modbus and DNP3 SCADA preprocessors detect traffic anomalies and provide data to intrusion rules. Supervisory Control and Data Acquisition (SCADA) protocols monitor, control, and acquire data from industrial, infrastructure, and facility processes such as manufacturing, production, water treatment, electric power distribution, airport and shipping systems, and so on.
- Several preprocessors allow you to detect specific threats, such as Back Orifice, portscans, SYN floods and other rate-based attacks.

  Note that you configure the sensitive data preprocessor, which detects sensitive data such as credit card numbers and Social Security numbers in ASCII text, in intrusion policies.

In a newly created access control policy, one default network analysis policy governs preprocessing for *all* traffic for *all* intrusion policies invoked by the same parent access control policy. Initially, the system uses the Balanced Security and Connectivity network analysis policy as the default, but you can change it to another system-provided or custom network analysis policy. In a more complex deployment, advanced users can tailor traffic preprocessing options to specific security zones, networks, and VLANs by assigning different custom network analysis policies to preprocess matching traffic.

Access Control Rules: Intrusion Policy Selection

After initial preprocessing, access control rules (when present) evaluate traffic. In most cases, the first access control rule that a packet matches is the rule that handles that traffic; you can monitor, trust, block, or allow matching traffic.

When you allow traffic with an access control rule, the system can inspect the traffic for discovery data, malware, prohibited files, and intrusions, in that order. Traffic not matching any access control rule is handled by the access control policy's default action, which can also inspect for discovery data and intrusions.



Note All packets, **regardless** of which network analysis policy preprocesses them, are matched to configured access control rules—and thus are potentially subject to inspection by intrusion policies—in top-down order.

The diagram in How Policies Examine Traffic For Intrusions, on page 1448 shows the flow of traffic through a device in an inline, intrusion prevention and AMP for Networks deployment, as follows:

- Access Control Rule A allows matching traffic to proceed. The traffic is then inspected for discovery data by the network discovery policy, for prohibited files and malware by File Policy A, and then for intrusions by Intrusion Policy A.
- Access Control Rule B also allows matching traffic. However, in this scenario, the traffic is not inspected for intrusions (or files or malware), so there are no intrusion or file policies associated with the rule. Note that by default, traffic that you allow to proceed is inspected by the network discovery policy; you do not need to configure this.
- In this scenario, the access control policy's default action allows matching traffic. The traffic is then inspected by the network discovery policy, and then by an intrusion policy. You can (but do not have to) use a different intrusion policy when you associate intrusion policies with access control rules or the default action.

The example in the diagram does not include any blocking or trusting rules because the system does not inspect blocked or trusted traffic.

Intrusion Inspection: Intrusion Policies, Rules, and Variable Sets

You can use intrusion prevention as the system's last line of defense before traffic is allowed to proceed to its destination. Intrusion policies govern how the system inspects traffic for security violations and, in inline deployments, can block or alter malicious traffic. The main function of intrusion policies is to manage which intrusion and preprocessor rules are enabled and how they are configured.

Intrusion and Preprocessor Rules

An intrusion rule is a specified set of keywords and arguments that detects attempts to exploit vulnerabilities on your network; the system uses an intrusion rule to analyze network traffic to check if it matches the criteria in the rule. The system compares packets against the conditions specified in each rule and, if the packet data matches all the conditions specified in a rule, the rule triggers.

The system includes the following types of rules created by Cisco Talos Intelligence Group (Talos):

- *shared object intrusion rules*, which are compiled and cannot be modified (except for rule header information such as source and destination ports and IP addresses)
- *standard text intrusion rules*, which can be saved and modified as new custom instances of the rule.
- *preprocessor rules*, which are rules associated with preprocessors and packet decoder detection options in the network analysis policy. You cannot copy or edit preprocessor rules. Most preprocessor rules are disabled by default; you must enable them to use preprocessors to generate events and, in an inline deployment, drop offending packets.

When the system processes packets according to an intrusion policy, first a rule optimizer classifies all activated rules in subsets based on criteria such as: transport layer, application protocol, direction to or from the protected network, and so on. Then, the intrusion rules engine selects the appropriate rule subsets to apply to each packet. Finally, a multi-rule search engine performs three different types of searches to determine if the traffic matches the rule:

- The protocol field search looks for matches in particular fields in an application protocol.
- The generic content search looks for ASCII or binary byte matches in the packet payload.
- The packet anomaly search looks for packet headers and payloads that, rather than containing specific content, violate well-established protocols.

In a custom intrusion policy, you can tune detection by enabling and disabling rules, as well as by writing and adding your own standard text rules. You can also use Firepower recommendations to associate the operating systems, servers, and client application protocols detected on your network with rules specifically written to protect those assets.

Variable Sets

Whenever the system uses an intrusion policy to evaluate traffic, it uses an associated *variable set*. Most variables in a set represent values commonly used in intrusion rules to identify source and destination IP addresses and ports. You can also use variables in intrusion policies to represent IP addresses in rule suppressions and dynamic rule states.

The system provides a single default variable set, which is comprised of predefined default variables. Most system-provided shared object rules and standard text rules use these predefined default variables to define networks and port numbers. For example, the majority of the rules use the variable `$HOME_NET` to specify the protected network and the variable `$EXTERNAL_NET` to specify the unprotected (or outside) network. In addition,

specialized rules often use other predefined variables. For example, rules that detect exploits against web servers use the `$HTTP_SERVERS` and `$HTTP_PORTS` variables.

Tip Even if you use system-provided intrusion policies, Cisco **strongly** recommends that you modify key default variables in the default set. When you use variables that accurately reflect your network environment, processing is optimized and the system can monitor relevant systems for suspicious activity. Advanced users can create and use custom variable sets for pairing with one or more custom intrusion policies.

Related Topics

Predefined Default Variables, on page 414

Intrusion Event Generation

When the system identifies a possible intrusion, it generates an *intrusion or preprocessor event* (sometimes collectively called *intrusion events*). Managed devices transmit their events to the Firepower Management Center, where you can view the aggregated data and gain a greater understanding of the attacks against your network assets. In an inline deployment, managed devices can also drop or replace packets that you know to be harmful.

Each intrusion event in the database includes an event header and contains information about the event name and classification; the source and destination IP addresses; ports; the process that generated the event; and the date and time of the event, as well as contextual information about the source of the attack and its target. For packet-based events, the system also logs a copy of the decoded packet header and payload for the packet or packets that triggered the event.

The packet decoder, the preprocessors, and the intrusion rules engine can all cause the system to generate an event. For example:

- If the packet decoder (configured in the network analysis policy) receives an IP packet that is less than 20 bytes, which is the size of an IP datagram without any options or payload, the decoder interprets this as anomalous traffic. If, later, the accompanying decoder rule in the intrusion policy that examines the packet is enabled, the system generates a preprocessor event.
- If the IP defragmentation preprocessor encounters a series of overlapping IP fragments, the preprocessor interprets this as a possible attack and, when the accompanying preprocessor rule is enabled, the system generates a preprocessor event.
- Within the intrusion rules engine, most standard text rules and shared object rules are written so that they generate intrusion events when triggered by packets.

As the database accumulates intrusion events, you can begin your analysis of potential attacks. The system provides you with the tools you need to review intrusion events and evaluate whether they are important in the context of your network environment and your security policies.

System-Provided and Custom Network Analysis and Intrusion Policies

Creating a new access control policy is one of the first steps in managing traffic flow using the Firepower System. By default, a newly created access control policy invokes system-provided network analysis and intrusion policies to examine traffic.

The following diagram shows how a newly created access control policy in an inline, intrusion-prevention deployment initially handles traffic. The preprocessing and intrusion prevention phases are highlighted.



Note how:

- A default network analysis policy governs the preprocessing of *all* traffic handled by the access control policy. Initially, the system-provided *Balanced Security and Connectivity network analysis policy* is the default.
- The default action of the access control policy allows all non-malicious traffic, as determined by the system-provided *Balanced Security and Connectivity intrusion policy*. Because the default action allows traffic to pass, the discovery feature can examine it for host, application, and user data before the intrusion policy can examine and potentially block malicious traffic.
- The policy uses default Security Intelligence options (global Block and Do Not Block lists only), does not decrypt encrypted traffic with an SSL policy, and does not perform special handling and inspection of network traffic using access control rules.

A simple step you can take to tune your intrusion prevention deployment is to use a different set of system-provided network analysis and intrusion policies as your defaults. Cisco delivers several pairs of these policies with the Firepower System.

Or, you can tailor your intrusion prevention deployment by creating and using custom policies. You may find that the preprocessor options, intrusion rule, and other advanced settings configured in those policies do not address the security needs of your network. By tuning your network analysis and intrusion policies you can configure, at a very granular level, how the system processes and inspects the traffic on your network for intrusions.

System-Provided Network Analysis and Intrusion Policies

Cisco delivers several pairs of network analysis and intrusion policies with the Firepower System. By using system-provided network analysis and intrusion policies, you can take advantage of the experience of the Cisco Talos Intelligence Group (Talos). For these policies, Talos provides intrusion and preprocessor rule states as well as initial configurations for preprocessors and other advanced settings.

No system-provided policy covers every network profile, traffic mix, or defensive posture. Each covers common cases and network setups that provide a starting point for a well-tuned defensive policy. Although you can use system-provided policies as-is, Cisco strongly recommends that you use them as the base for custom policies that you tune to suit your network.



Tip Even if you use system-provided network analysis and intrusion policies, you should configure the system's intrusion variables to accurately reflect your network environment. At a minimum, modify key default variables in the default set.

As new vulnerabilities become known, Talos releases intrusion rule updates (also known as *Snort Rule Updates*). These rule updates can modify any system-provided network analysis or intrusion policy, and can provide new and updated intrusion rules and preprocessor rules, modified states for existing rules, and modified default policy settings. Rule updates may also delete rules from system-provided policies and provide new rule categories, as well as modify the default variable set.

If a rule update affects your deployment, the web interface marks affected intrusion and network analysis policies as out of date, as well as their parent access control policies. You must re-deploy an updated policy for its changes to take effect.

For your convenience, you can configure rule updates to automatically re-deploy affected intrusion policies, either alone or in combination with affected access control policies. This allows you to easily and automatically keep your deployment up-to-date to protect against recently discovered exploits and intrusions.

To ensure up-to-date preprocessing settings, you **must** re-deploy access control policies, which also deploys any associated SSL, network analysis, and file policies that are different from those currently running, and can also can update default values for advanced preprocessing and performance options.

Cisco delivers the following network analysis and intrusion policies with the Firepower System:

Balanced Security and Connectivity network analysis and intrusion policies

These policies are built for both speed and detection. Used together, they serve as a good starting point for most organizations and deployment types. The system uses the Balanced Security and Connectivity policies and settings as defaults in most cases.

Connectivity Over Security network analysis and intrusion policies

These policies are built for organizations where connectivity (being able to get to all resources) takes precedence over network infrastructure security. The intrusion policy enables far fewer rules than those enabled in the Security over Connectivity policy. Only the most critical rules that block traffic are enabled.

Security Over Connectivity network analysis and intrusion policies

These policies are built for organizations where network infrastructure security takes precedence over user convenience. The intrusion policy enables numerous network anomaly intrusion rules that could alert on or drop legitimate traffic.

Maximum Detection network analysis and intrusion policies

These policies are built for organizations where network infrastructure security is given even more emphasis than is given by the Security Over Connectivity policies, with the potential for even greater operational impact. For example, the intrusion policy enables rules in a large number of threat categories including malware, exploit kit, old and common vulnerabilities, and known in-the-wild exploits.

No Rules Active intrusion policy

In the No Rules Active intrusion policy, all intrusion rules, and all advanced settings except intrusion rule thresholds, are disabled. This policy provides a starting point if you want to create your own intrusion policy instead of basing it on the enabled rules in one of the other system-provided policies.



Note Depending on the system-provided base policy that is selected, the settings of the policy vary. To view the policy settings, click the **Edit** icon next to the policy and then click the **Manage Base Policy** link.

Benefits of Custom Network Analysis and Intrusion Policies

You may find that the preprocessor options, intrusion rules, and other advanced settings configured in the system-provided network analysis and intrusion policies do not fully address the security needs of your organization.

Building custom policies can improve the performance of the system in your environment and can provide a focused view of the malicious traffic and policy violations occurring on your network. By creating and tuning custom policies you can configure, at a very granular level, how the system processes and inspects the traffic on your network for intrusions.

All custom policies have a base policy, also called a base layer, which defines the default settings for all configurations in the policy. A layer is a building block that you can use to efficiently manage multiple network analysis or intrusion policies.

In most cases, you base custom policies on system-provided policies, but you can use another custom policy. However, all custom policies have a system-provided policy as the eventual base in a policy chain. Because rule updates can modify system-provided policies, importing a rule update may affect you even if you are using a custom policy as your base. If a rule update affects your deployment, the web interface marks affected policies as out of date.

Benefits of Custom Network Analysis Policies

By default, one network analysis policy preprocesses all unencrypted traffic handled by the access control policy. That means that all packets are decoded and preprocessed according to the same settings, regardless of the intrusion policy (and therefore intrusion rule set) that later examines them.

Initially, the system-provided Balanced Security and Connectivity network analysis policy is the default. A simple way to tune preprocessing is to create and use a custom network analysis policy as the default.

Tuning options available vary by preprocessor, but some of the ways you can tune preprocessors and decoders include:

- You can disable preprocessors that do not apply to the traffic you are monitoring. For example, the HTTP Inspect preprocessor normalizes HTTP traffic. If you are confident that your network does not include any web servers using Microsoft Internet Information Services (IIS), you can disable the preprocessor option that looks for IIS-specific traffic and thereby reduce system processing overhead.



Note If you disable a preprocessor in a custom network analysis policy, but the system needs to use that preprocessor to later evaluate packets against an enabled intrusion or preprocessor rule, the system automatically enables and uses the preprocessor although the preprocessor remains disabled in the network analysis policy web interface.

- Specify ports, where appropriate, to focus the activity of certain preprocessors. For example, you can identify additional ports to monitor for DNS server responses or encrypted SSL sessions, or ports on which you decode telnet, HTTP, and RPC traffic.

For advanced users with complex deployments, you can create multiple network analysis policies, each tailored to preprocess traffic differently. Then, you can configure the system to use those policies to govern the preprocessing of traffic using different security zones, networks, or VLANs. (Note that ASA FirePOWER modules cannot restrict preprocessing by VLAN.)



Note Tailoring preprocessing using custom network analysis policies—especially multiple network analysis policies—is an advanced task. Because preprocessing and intrusion inspection are so closely related, you **must** be careful to allow the network analysis and intrusion policies examining a single packet to complement each other.

Benefits of Custom Intrusion Policies

In a newly created access control policy initially configured to perform intrusion prevention, the default action allows all traffic, but first inspects it with the system-provided Balanced Security and Connectivity intrusion policy. Unless you add access control rules or change the default action, all traffic is inspected by that intrusion policy.

To customize your intrusion prevention deployment, you can create multiple intrusion policies, each tailored to inspect traffic differently. Then, configure an access control policy with rules that specify which policy inspects which traffic. Access control rules can be simple or complex, matching and inspecting traffic using multiple criteria including security zone, network or geographical location, VLAN, port, application, requested URL, or user.

The main function of intrusion policies is to manage which intrusion and preprocessor rules are enabled and how they are configured, as follows:

- Within each intrusion policy, you should verify that all rules applicable to your environment are enabled, and improve performance by disabling rules that are not applicable to your environment. In an inline deployment, you can specify which rules should drop or modify malicious packets.
- Firepower recommendations allow you to associate the operating systems, servers, and client application protocols detected on your network with rules specifically written to protect those assets.
- You can modify existing rules and write new standard text rules as needed to catch new exploits or to enforce your security policies.

Other customizations you might make to an intrusion policy include:

- The sensitive data preprocessor detects sensitive data such as credit card numbers and Social Security numbers in ASCII text. Note that other preprocessors that detect specific threats (back orifice attacks, several portscan types, and rate-based attacks that attempt to overwhelm your network with excessive traffic) are configured in network analysis policies.
- Global thresholds cause the system to generate events based on how many times traffic matching an intrusion rule originates from or is targeted to a specific address or address range within a specified time period. This helps prevent the system from being overwhelmed with a large number of events.
- Suppressing intrusion event notifications and setting thresholds for individual rules or entire intrusion policies can also can prevent the system from being overwhelmed with a large number of events.
- In addition to the various views of intrusion events within the web interface, you can enable logging to syslog facilities or send event data to an SNMP trap server. Per policy, you can specify intrusion event notification limits, set up intrusion event notification to external logging facilities, and configure external

responses to intrusion events. Note that in addition to these per-policy alerting configurations, you can globally enable or disable email alerting on intrusion events for each rule or rule group. Your email alert settings are used regardless of which intrusion policy processes a packet.

Limitations of Custom Policies

Because preprocessing and intrusion inspection are so closely related, you **must** be careful that your configuration allows the network analysis and intrusion policies processing and examining a single packet to complement each other.

By default, the system uses one network analysis policy to preprocess all traffic handled by managed devices using a single access control policy. The following diagram shows how a newly created access control policy in an inline, intrusion-prevention deployment initially handles traffic. The preprocessing and intrusion prevention phases are highlighted.



Notice how a default network analysis policy governs the preprocessing of *all* traffic handled by the access control policy. Initially, the system-provided Balanced Security and Connectivity network analysis policy is the default.

A simple way to tune preprocessing is to create and use a custom network analysis policy as the default. However, if you disable a preprocessor in a custom network analysis policy but the system needs to evaluate preprocessed packets against an enabled intrusion or preprocessor rule, the system automatically enables and uses the preprocessor although it remains disabled in the network analysis policy web interface.

Note In order to get the performance benefits of disabling a preprocessor, you **must** make sure that none of your intrusion policies have enabled rules that require that preprocessor.

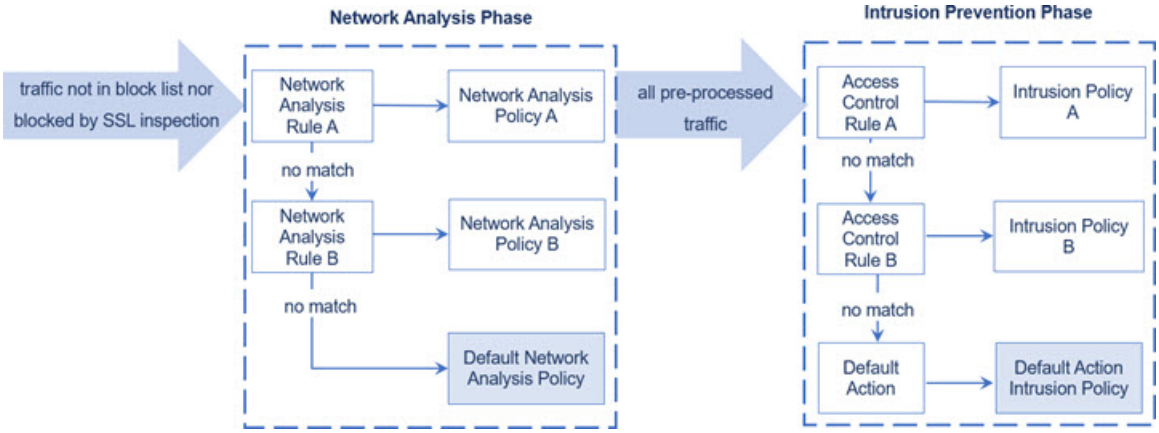
An additional challenge arises if you use multiple custom network analysis policies. For advanced users with complex deployments, you can tailor preprocessing to specific security zones, networks, and VLANs by assigning custom network analysis policies to preprocess matching traffic. (Note that ASA FirePOWER cannot restrict preprocessing by VLAN.) To accomplish this, you add custom *network analysis rules* to your access control policy. Each rule has an associated network analysis policy that governs the preprocessing of traffic that matches the rule.

Tip You configure network analysis rules as an advanced setting in an access control policy. Unlike other types of rules in the Firepower System, network analysis rules invoke—rather than being contained by—network analysis policies.

The system matches packets to any configured network analysis rules in top-down order by rule number. Traffic that does not match any network analysis rule is preprocessed by the default network analysis policy. While this allows you a great deal of flexibility in preprocessing traffic, keep in mind that all packets, **regardless** of which network analysis policy preprocessed them, are subsequently matched to access control rules—and thus to potential inspection by intrusion policies—in their own process. In other words, preprocessing a packet with a particular network analysis policy does **not** guarantee that the packet will be examined with any

particular intrusion policy. You **must** carefully configure your access control policy so it invokes the correct network analysis and intrusion policies to evaluate a particular packet.

The following diagram shows in focused detail how the network analysis policy (preprocessing) selection phase occurs before and separately from the intrusion prevention (rules) phase. For simplicity, the diagram eliminates the discovery and file/malware inspection phases. It also highlights the default network analysis and default-action intrusion policies.



In this scenario, an access control policy is configured with two network analysis rules and a default network analysis policy:

- Network Analysis Rule A preprocesses matching traffic with Network Analysis Policy A. Later, you want this traffic to be inspected by Intrusion Policy A.
- Network Analysis Rule B preprocesses matching traffic with Network Analysis Policy B. Later, you want this traffic to be inspected by Intrusion Policy B.
- All remaining traffic is preprocessed with the default network analysis policy. Later, you want this traffic to be inspected by the intrusion policy associated with the access control policy's default action.

After the system preprocesses traffic, it can examine the traffic for intrusions. The diagram shows an access control policy with two access control rules and a default action:

- Access Control Rule A allows matching traffic. The traffic is then inspected by Intrusion Policy A.
- Access Control Rule B allows matching traffic. The traffic is then inspected by Intrusion Policy B.
- The access control policy's default action allows matching traffic. The traffic is then inspected by the default action's intrusion policy.

Each packet's handling is governed by a network analysis policy and intrusion policy pair, but the system does **not** coordinate the pair for you. Consider a scenario where you misconfigure your access control policy so that Network Analysis Rule A and Access Control Rule A do not process the same traffic. For example, you could intend the paired policies to govern the handling of traffic on a particular security zone, but you mistakenly use different zones in the two rules' conditions. This could cause traffic to be incorrectly preprocessed. For this reason, tailoring preprocessing using network analysis rules and custom policies is an **advanced** task.

Note that for a single connection, although the system selects a network analysis policy before an access control rule, some preprocessing (notably application layer preprocessing) occurs after access control rule selection. This does **not** affect how you configure preprocessing in custom network analysis policies.

License Requirements for Network Analysis and Intrusion Policies

FTD License

Threat

Classic License

Protection

Requirements and Prerequisites for Network Analysis and Intrusion Policies

Model Support

Any.

Supported Domains

Any

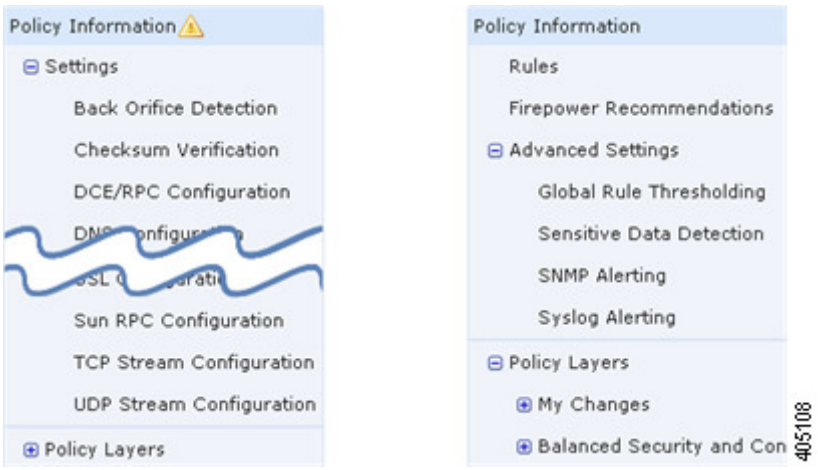
User Roles

- Admin
- Intrusion Admin

The Navigation Panel: Network Analysis and Intrusion Policies

Network analysis and intrusion policies use similar web interfaces to edit and save changes to their configurations.

A navigation panel appears on the left side of the web interface when you are editing either type of policy. The following graphic shows the navigation panel for the network analysis policy (left) and the intrusion policy (right).



A dividing line separates the navigation panel into links to policy settings you can configure with (below) or without (above) direct interaction with policy layers. To navigate to any settings page, click its name in the navigation panel. Dark shading of an item in the navigation panel highlights your current settings page. For example, in the illustration above the Policy Information page would be displayed to the right of the navigation panel.

Policy Information

The Policy Information page provides configuration options for commonly used settings. As shown in the illustration for the network analysis policy panel above, a **Policy Change icon** appears next to **Policy Information** in the navigation panel when the policy contains unsaved changes. The icon disappears when you save your changes.

Rules (intrusion policy only)

The Rules page in an intrusion policy allows you to configure rule states and other settings for shared object rules, standard text rules, and preprocessor rules.

Firepower Recommendations (intrusion policy only)

The Firepower Recommendations page in an intrusion policy allows you to associate the operating systems, servers, and client application protocols detected on your network with intrusion rules specifically written to protect those assets. This allows you to tailor your intrusion policy to the specific needs of your monitored network.

Settings (network analysis policy) and Advanced Settings (intrusion policy)

The Settings page in a network analysis policy allows you to enable or disable preprocessors and access preprocessor configuration pages. Expanding the **Settings** link displays sublinks to individual configuration pages for all enabled preprocessors in the policy.

The Advanced Settings page in an intrusion policy allows you to enable or disable advanced settings and access configuration pages for those advanced settings. Expanding the **Advanced Settings** link displays sublinks to individual configuration pages for all enabled advanced settings in the policy.

Policy Layers

The Policy Layers page displays a summary of the layers that comprise your network analysis or intrusion policy. Expanding the Policy Layers link displays sublinks to summary pages for the layers in your policy. Expanding each layer sublink displays further sublinks to the configuration pages for all rules, preprocessors, or advanced settings that are enabled in the layer.

Conflicts and Changes: Network Analysis and Intrusion Policies

When you edit a network analysis or intrusion policy, a **Policy Change icon** appears next to **Policy Information** in the navigation panel to indicate that the policy contains unsaved changes. You must save (or *commit*) your changes before the system recognizes them.



Note After you save, you must deploy the network analysis or intrusion policy for your changes to take effect. If you deploy a policy without saving, the system uses the most recently saved configuration.

Resolving Editing Conflicts

The Network Analysis Policy page (**Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**) and Intrusion Policy page (**Policies** > **Access Control** > **Intrusion**) display whether each policy has unsaved changes, as well as information about who is currently editing the policy. Cisco recommends that only one person edit a policy at a time. If you are performing simultaneous editing, the consequences are as follows:

- If you are editing a network analysis or intrusion policy at the same time another user is editing the same policy, and the other user saves their changes to the policy, you are warned when you commit the policy that you will overwrite the other user's changes.
- If you are editing the same network analysis or intrusion policy via multiple web interface instances as the same user, and you save your changes for one instance, you cannot save your changes for the other instance.

Resolving Configuration Dependencies

To perform their particular analysis, many preprocessors and intrusion rules require that traffic first be decoded or preprocessed in a certain way, or have other dependencies. When you save a network analysis or intrusion policy, the system either automatically enables required settings, or warns you that disabled settings will have no effect on traffic, as follows:

- You cannot save an intrusion policy if you added an SNMP rule alert but did not configure SNMP alerting. You must either configure SNMP alerting or disable the rule alert, then save again.
- You cannot save an intrusion policy if it includes enabled sensitive data rules but you have not enabled the sensitive data preprocessor. You must either allow the system to enable the preprocessor and save the policy, or disable the rules and save again.
- If you disable a required preprocessor in a network analysis policy, you can still save the policy. However, the system automatically uses the disabled preprocessor with its current settings, even though the preprocessor remains disabled in the web interface.

- If you disable inline mode in a network analysis policy but enable the Inline Normalization preprocessor, you can still save the policy. However, the system warns you that normalization settings will be ignored. Disabling inline mode also causes the system to ignore other settings that allow preprocessors to modify or block traffic, including checksum verification and rate-based attack prevention.

Committing, Discarding, and Caching Policy Changes

While editing a network analysis or intrusion policy, if you exit the policy editor without saving your changes, the system caches those changes. Your changes are cached even when you log out of the system or experience a system crash. The system cache can store unsaved changes for one network analysis and one intrusion policy per user; you must commit or discard your changes before editing another policy of the same type. The system discards the cached changes when you edit another policy without saving your changes to the first policy, or when you import an intrusion rule update.

You can commit or discard policy changes on the Policy Information page of either the network analysis or intrusion policy editor.

In the Firepower Management Center configuration, you can control:

- whether you are prompted (or required) to comment on your network analysis or intrusion policy changes when you commit them
- whether changes and comments are recorded in the audit log

Related Topics

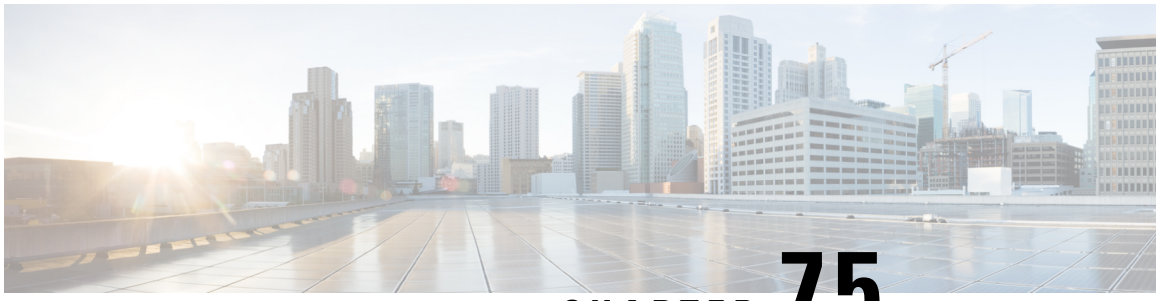
Configuring Network Analysis Policy Preferences
Configuring Intrusion Policy Preferences

Exiting a Network Analysis or Intrusion Policy

Procedure

If you want to exit the network analysis or intrusion policy advanced editor, you have the following choices:

- Cache — To exit the policy and cache changes, choose any menu or other path to another page. On exiting, click **Leave page** when prompted, or click **Stay on page** to remain in the advanced editor.
- Discard — To discard unsaved changes, click **Discard Changes** on the Policy Information page, then click **OK**.
- Save — To save changes to the policy, click **Commit Changes** on the Policy Information page. If prompted, enter a comment, and then click **OK**.



CHAPTER 75

Layers in Intrusion and Network Analysis Policies

The following topics explain how to use layers in intrusion and network analysis policies:

- Layer Basics, on page 1463
- License Requirements for Network Analysis and Intrusion Policy Layers, on page 1463
- Requirements and Prerequisites for Network Analysis and Intrusion Policy Layers, on page 1464
- The Layer Stack, on page 1464
- Layer Management, on page 1468

Layer Basics

Larger organizations with many managed devices may have many intrusion policies and network analysis policies to support the unique needs of different departments, business units or, in some instances, different companies. Configurations in both policy types are contained in building blocks called *layers,* which you can use to efficiently manage multiple policies.

Layers in intrusion and network analysis policies work in essentially the same way. You can create and edit either policy type without consciously using layers. You can modify your policy configurations and, if you have not added user layers to your policy, the system automatically includes your changes in a single configurable layer that is initially named *My Changes*. You can also add up to 200 layers where you can configure any combination of settings. You can copy, merge, move, and delete user layers and, most important, share individual user layers with other policies of the same type.

License Requirements for Network Analysis and Intrusion Policy Layers

FTD License

Threat

Classic License

Protection

Requirements and Prerequisites for Network Analysis and Intrusion Policy Layers

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Intrusion Admin

The Layer Stack

Layer stacks are composed of the following:

User Layers

User-configurable layers. You can copy, merge, move, or delete any user-configurable layer and set any user-configurable layer to be shared by other policies of the same type. This layer includes the automatically-generated layer initially named My Changes.

Built-in Layers

The read-only base policy layer. The policy in this layer can be either a system-provided policy or a custom policy you created.

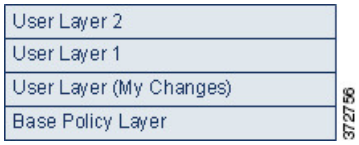
By default, a network analysis or intrusion policy includes a base policy layer and a My Changes layer. You can add user layers as necessary.

Each policy layer contains complete configurations for either all preprocessors in a network analysis policy or all intrusion rules and advanced settings in an intrusion policy. The lowest, base policy layer includes all the settings from the base policy you selected when you created the policy. A setting in a higher layer takes precedence over the same setting in a lower layer. Features not explicitly set in a layer *inherit* their settings from the next highest layer where they are explicitly set. The system *flattens* the layers, that is, it applies only the cumulative effect of all settings, when it handles network traffic.

Tip You can create an intrusion or network analysis policy based solely on the default settings in the base policy. In the case of an intrusion policy, you can also use Firepower rule state recommendations if you want to tailor your intrusion policy to the specific needs of your monitored network.

The following figure shows an example layer stack that, in addition to the base policy layer and the initial My Changes layer, also includes two additional user-configurable layers, *User Layer 1* and *User Layer 2*. Note

in the figure that each user-configurable layer that you add is initially positioned as the highest layer in the stack; thus, User Layer 2 in the figure was added last and is highest in the stack.



Regardless of whether you allow rule updates to modify your policy, changes in a rule update never override changes you make in a layer. This is because changes in a rule update are made in the base policy, which determines the default settings in your base policy layer; your changes are always made in a higher layer, so they override any changes that a rule update makes to your base policy.

The Base Layer

The base layer, also referred to as the base policy, of an intrusion or network analysis policy defines the default settings for all configurations in the policy, and is the lowest layer in the policy. When you create a new policy and change a setting without adding new layers, the change is stored in the My Changes layer, and overrides—but does not change—the setting in the base policy.

System-Provided Base Policies

The Firepower System provides several pairs of network analysis and intrusion policies. By using system-provided network analysis and intrusion policies, you can take advantage of the experience of the Cisco Talos Intelligence Group (Talos). For these policies, Talos sets intrusion and preprocessor rule states, as well as provides the initial configurations for preprocessors and other advanced settings. You can use these system-provided policies as-is, or you can use them as the base for custom policies.

If you use a system-provided policy as your base, importing rule updates may modify settings in your base policy. However, you can configure a custom policy so that the system does not automatically make these changes to its system-provided base policy. This allows you to update system-provided base policies manually, on a schedule independent of rule updates. In either case, changes that a rule update makes to your base policy do not change or override settings in your My Changes or any other layer.

System-provided intrusion and network analysis policies are similarly named but contain different configurations. For example, the Balanced Security and Connectivity network analysis policy and the Balanced Security and Connectivity intrusion policy work together and can both be updated in intrusion rule updates.

Custom Base Policies

You can use a custom policy as your base. You can tune settings in custom policies to inspect traffic in ways that matter most to you so you can improve both the performance of your managed devices and your ability to respond effectively to the events they generate.

If you change the custom policy that you use as the base for another policy, those changes are automatically used as the default settings of the policy that uses the base.

In addition, a rule update may affect your policy even if you use a custom base policy, because all policies have a system-provided policy as the eventual base in a policy chain. If the first custom policy in a chain (the one that uses the system-provided policy as its base) allows rule updates to modify its base policy, your policy may be affected.

Regardless of how changes are made to your base policy—whether by a rule update or when you modify a custom policy that you use as a base policy—they do not change or override settings in your My Changes or any other layer.

The Effect of Rule Updates on Base Policies

When you import rule updates, the system modifies system-provided intrusion, access control, and network analysis policies. Rule updates can include:

- modified network analysis preprocessor settings
- modified advanced settings in intrusion and access control policies
- new and updated intrusion rules
- modified states for existing rules
- new rule categories and default variables

Rule updates can also delete existing rules from system-provided policies.

Changes to default variables and rule categories are handled at the system level.

When you use a system-provided policy as your intrusion or network analysis base policy, you can allow rule updates to modify your base policy which, in this case, is a copy of the system-provided policy. If you allow rule updates to update your base policy, a new rule update makes the same changes in your base policy that it makes to the system-provided policy that you use as your base policy. If you have not modified the corresponding setting, a setting in your base policy determines the setting in your policy. However, rule updates do not override changes you make in your policy.

If you do not allow rule updates to modify your base policy, you can manually update your base policy after importing one or more rule updates.

Rule updates always delete intrusion rules that Talos deletes, regardless of the rule state in your intrusion policy or whether you allow rule updates to modify your base intrusion policy.

Until you re-deploy your changes to network traffic, rules in your currently deployed intrusion policies behave as follows:

- Disabled intrusion rules remain disabled.
- Rules set to **Generate Events** continue to generate events when triggered.
- Rules set to **Drop and Generate Events** continue to generate events and drop offending packets when triggered.

Rule updates do not modify a custom base policy unless both of the following conditions are met:

- You allow rule updates to modify the system-provided base policy of the parent policy, that is, the policy that originated the custom base policy.
- You have not made changes in the parent policy that override the corresponding settings in the parent's base policy.

When both conditions are met, changes in the rule update are passed to the child policy, that is, the policy using the custom base policy, when you save the parent policy.

For example, if a rule update enables a previously disabled intrusion rule, and you have not modified the rule's state in the parent intrusion policy, the modified rule state is passed to the base policy when you save the parent policy.

Likewise, if a rule update modifies a default preprocessor setting and you have not modified the setting in the parent network analysis policy, the modified setting is passed to the base policy when you save the parent policy.

Changing the Base Policy

You can choose a different system-provided or custom policy as your base policy.

You can chain up to five custom policies, with four of the five using one of the other four previously created policies as its base policy; the fifth must use a system-provided policy as its base.

Procedure

Step 1 While editing your policy, click **Policy Information** in the navigation panel.

Step 2 You can configure the following choices:

- Choose a base policy — Choose from the **Base Policy** drop-down list.
- Allow rule updates to modify the base policy — Click **Manage Base Policy**, then check the **Update when a new Rule Update is installed** check box.

Tip When you save your policy with the check box cleared and then import a rule update, an **Update Now** appears on the Base Policy summary page and the status message on the page updates to inform you that the policy is out of date. If you want to update your base policy with the changes in the most recently imported rule update, click **Update Now**.

Step 3 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

The Firepower Recommendations Layer

When you generate rule state recommendations in an intrusion policy, you can choose whether to automatically modify rule states based on the recommendations.

As seen in the following figure, using recommended rule states inserts a read-only, built-in Firepower Recommendations layer immediately above the base layer.



Note that this layer is unique to intrusion policies.

If you subsequently choose not to use recommended rule states, the system removes the Firepower Recommendations layer. You cannot manually delete this layer, but you can add and remove it by choosing to use or not use recommended rule states.

Adding the Firepower Recommendations layer adds a Firepower Recommendations link under Policy Layers in the navigation panel. This link leads you to a read-only view of the Firepower Recommendations layer page where you can access a recommendation-filtered view of the Rules page in read-only mode.

Using recommended rule states also adds a Rules sublink beneath the Firepower Recommendations link in the navigation panel. The Rules sublink provides access to a read-only display of the Rules page in the Firepower Recommendations layer. Note the following in this view:

- When there is no rule state icon in the state column, the state is inherited from the base policy.
- When there is no rule state icon in the Firepower Recommendation column in this or other Rules page views, there is no recommendation for this rule.

Related Topics

Tailoring Intrusion Protection to Your Network Assets, on page 1517

Layer Management

The Policy Layers page provides a single-page summary of the complete layer stack for your network analysis or intrusion policy. On this page you can add shared and unshared layers, copy, merge, move, and delete layers, access the summary page for each layer, and access configuration pages for enabled, disabled, and overridden configurations within each layer.

For each layer, you can view the following information:

- whether the layer is a built-in, shared user, or unshared user layer
- which layers contain the highest, that is the effective, preprocessor or advanced setting configurations, by feature name
- in an intrusion policy, the number of intrusion rules whose states are set in the layer, and the number of rules set to each rule state.

The Policy Layers page also provides a summary of the net effect of all enabled preprocessors (network analysis) or advanced settings (intrusion) and, for intrusion policies, intrusion rules.

The feature name in the summary for each layer indicates which configurations are enabled, disabled, overridden, or inherited in the layer, as follows:

When the feature is...	The feature name is...
enabled in the layer	written in plain text
disabled in the layer	struck out
overridden by the configuration in a higher layer	written in italic text
inherited from a lower layer	not present

You can add up to 200 layers to a network analysis or intrusion policy. When you add a layer, it appears as the highest layer in your policy. The initial state is Inherit for all features and, in an intrusion policy, no event filtering, dynamic state, or alerting rule actions are set.

You give a user-configurable layer a unique name when you add the layer to your policy. Later, you can change the name and, optionally, add or modify a description that is visible when you edit the layer.

You can copy a layer, move a layer up or down within the User Layers page area, or delete a user layer, including the initial My Changes layer. Note the following considerations:

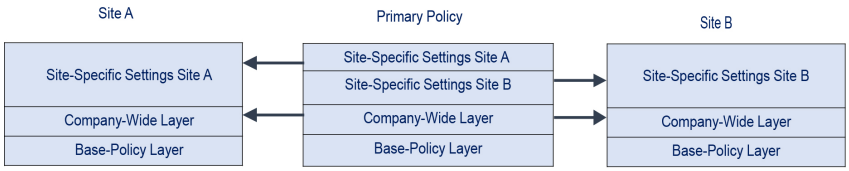
- When you copy a layer, the copy appears as the highest layer.
- Copying a shared layer creates a layer that is initially unshared and which you can then share if you choose.
- You cannot delete a shared layer; a layer with sharing enabled that you have not shared with another policy is not a shared layer.

You can merge a user-configurable layer with another user-configurable layer immediately beneath it. A merged layer retains all settings that were unique to either layer, and accepts the settings from the higher layer if both layers included settings for the same preprocessor, intrusion rule, or advanced setting. The merged layer retains the name of the lower layer. In the policy where you create a sharable layer that you can add to other policies, you can merge an unshared layer immediately above the sharable layer with the sharable layer, but you cannot merge the sharable layer with an unshared layer beneath it. In a policy where you add a shared layer that you created in another policy, you can merge the shared layer into an unshared layer immediately beneath it and the resulting layer is no longer shared; you cannot merge an unshared layer into a shared layer beneath it.

Shared Layers

A *shared layer* is a layer you add to your policy after creating the layer in another policy where you allow it to be shared. A *sharable layer* is a layer you allow to be shared.

The following figure shows an example primary policy where you create the company-wide layer and site-specific layers for sites A and B, and allow these to be shared. You then add these as shared layers to the policies for sites A and B.



The company-wide layer in the primary policy includes settings applicable to sites A and B. The site-specific layers include settings specific to each site. For example, in the case of a network analysis policy Site A might not have web servers on the monitored network and would not require the protection or processing overhead of the HTTP Inspect preprocessor, but both sites would likely require TCP stream preprocessing. You could enable TCP stream processing in the company-wide layer that you share with both sites, disable the HTTP Inspect preprocessor in the site-specific layer that you share with Site A, and enable the HTTP Inspect preprocessor in the site-specific layer that you share with Site B. By editing configurations in a higher layer in the site-specific policies, you could also further tune the policy for each site if necessary with any configuration adjustments.

It is unlikely that the flattened net settings in the example primary policy would be useful for monitoring traffic, but the time saved in configuring and updating the site-specific policies makes this a useful application of policy layers.

Many other layer configurations are possible. For example, you could define policy layers by company, by department, by network, or even by user. In the case of an intrusion policy, you could also include advanced settings in one layer and rule settings in another.

You can allow a user-configurable layer to be shared with other policies of the same type (intrusion or network analysis). When you modify a configuration within a sharable layer and then commit your changes, the system updates all policies that share the layer and provides you with a list of all affected policies. You can only change feature configurations in the policy where you created the layer.

You cannot disable sharing for a layer that you have added to another policy; you must first delete the layer from the other policy or delete the other policy.

You cannot add a shared layer to a policy when your base policy is a custom policy where the layer you want to share was created. To do so would give the policy a circular dependency.

In a multidomain deployment, you can add shared layers from ancestor policies to policies in descendant domains.

Managing Layers

Procedure

Step 1 While editing your policy, click **Policy Layers** in the navigation panel.

Step 2 You can take any of the following management actions on the Policy Layers page:

- Add a shared layer from another policy — Click add shared layer **Add** () next to User Layers, choose the layer from the **Add Shared Layer** drop-down list, then click **OK**.
- Add an unshared layer — Click add layer **Add** () next to User Layers, enter a **Name**, and click **OK**.
- Add or change the layer description — Click **Edit** () next to the layer, then add or change the **Description**.
- Allow a layer to be shared with another policy — Click **Edit** () next to the layer, then clear the **Sharing** check box.
- Change the layer name — Click **Edit** () next to the layer, then change the **Name**.
- Copy a layer — Click **Copy** () for the layer.

- Delete a layer — Click **Delete** ( ) for the layer, then click **OK**.
- Merge two layers — Click **Merge** ( ) for the upper of the two layers, then click **OK**.
- Move a layer — Click any open area in the layer summary and drag until the **Position Arrow** points to a line above or below a layer where you want to move the layer.

Step 3 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

Navigating Layers

Procedure

Step 1 While editing your policy, click **Policy Layers** in the navigation panel.

Step 2 You can take any of the following actions to navigate through your layers:

- Access a preprocessor or advanced settings page — If you want to access a layer-level preprocessor or advanced setting configuration page, click the feature name in the row for the layer. Configuration pages are read-only in the base policy and in shared layers.
- Access a rule page — If you want to access a layer-level rule configuration page filtered by rule state type, click **Drop and Generate Events**, **Generate Events**, or **Disabled** in the summary for the layer. No rules are displayed if the layer contains no rules set to the selected rule state.
- Display the Policy Information page — If you want to display the Policy Information page, click **Policy Summary** in the navigation panel.
- Display a layer summary page — If you want to display the summary page for a layer, click the layer name in the row for the layer or, alternately, click **Edit** ( ) next to a user layer. You can also click **View** ( ) to access the read-only summary page for a shared layer.

Step 3 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

Intrusion Rules in Layers

You can view individual layer settings on the Rules page for the layer, or view the net effect of all settings on the policy view of the Rules page. When you modify rule settings on the policy view of the Rules page, you are modifying the highest user-configurable layer in the policy. You can switch to another layer using the layer drop-down list on any Rules page.

The following table describes the effects of configuring the same type of setting in multiple layers.

Table 103: Layer Rule Settings

You can set...	Of this setting type...	To...
one	rule state	override a rule state set for the rule in a lower layer, and ignore all thresholds, suppressions, rate-based rule states, and alerts for that rule configured in lower layers. If you want a rule to inherit its state from the base policy or a lower layer, set the rule state to Inherit. Note that when you are working on the intrusion policy Rules page, you cannot set a rule state to Inherit because the intrusion policy Rules page is a composite view of the net effect of all rule settings.
one	threshold SNMP alert	override a setting of the same type for the rule in a lower layer. Note that setting a threshold overwrites any existing threshold for the rule in the layer.
one or more	suppression rate-based rule state	cumulatively combine settings of the same type for each selected rule down to the first layer where a rule state is set for the rule. Settings below the layer where a rule state is set are ignored.
one or more	comment	add a comment to a rule. Comments are rule-specific, not policy- or layer-specific. You can add one or more comments to a rule in any layer.

For example, if you set a rule state to Drop and Generate Events in one layer and to Disabled in a higher layer, the intrusion policy Rules page shows that the rule is disabled.

In another example, if you set a source-based suppression for a rule to 192.168.1.1 in one layer, and you also set a destination-based suppression for the rule to 192.168.1.2 in another layer, the Rules page shows that the cumulative effect is to suppress events for the source address 192.168.1.1 and the destination address 192.168.1.2. Note that suppression and rate-based rule state settings cumulatively combine settings of the same type for each selected rule down to the first layer where a rule state is set for the rule. Settings below the layer where a rule state is set are ignored.

Color-coding on each Rules page for a specific layer indicates whether the effective state is in a higher, lower, or the current layer, as follows:

- red—the effective state is in a higher layer
- yellow—the effective state is in a lower layer

- unshaded—the effective state is in the current layer

Because the intrusion policy Rules page is a composite view of the net effect of all rule settings, rule states are not color-coded on this page.

Configuring Intrusion Rules in Layers

In an intrusion policy, you can set the rule state, event filtering, dynamic state, alerting, and rule comments for a rule in any user-configurable layer. After accessing the layer where you want to make your changes, you add settings on the Rules page for the layer the same as you would on the intrusion policy Rules page.

Procedure

Step 1 While editing your intrusion policy, expand **Policy Layers** in the navigation panel.

Step 2 Expand the policy layer you want to modify.

Step 3 Click **Rules** immediately beneath the policy layer you want to modify.

Step 4 Modify any of the settings described in Tuning Intrusion Policies Using Rules, on page 1489.

Tip To delete an individual setting from an editable layer, double-click the rule message on the Rules page for the layer to display rule details. Click **Delete** next to the setting you want to delete, then click **OK** twice.

Step 5 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

Removing Rule Settings from Multiple Layers

You can simultaneously remove a specific type of event filter, dynamic state, or alerting from multiple layers in your intrusion policy. The system removes the selected setting and copies the remaining settings for the rule to the highest editable layer in the policy.

The system removes the setting type downward through each layer where it is set until it removes all the settings or encounters a layer where a rule state is set for the rule. In the latter case, it removes the setting from that layer and stops removing the setting type.

When the system encounters the setting type in a shared layer or in the base policy, and if the highest layer in the policy is editable, the system copies the remaining settings and rule state for the rule to that editable layer. Otherwise, if the highest layer in the policy is a shared layer, the system creates a new editable layer above the shared layer and copies the remaining settings and rule state for the rule to that editable layer.

Note Removing rule settings derived from a shared layer or the base policy causes any changes to this rule from lower layers or the base policy to be ignored. To stop ignoring changes from lower layers or the base policy, set the rule state to **Inherit** on the summary page for the topmost layer.

Procedure

Step 1 While editing your intrusion policy, click **Rules** immediately beneath **Policy Information** in the navigation panel.

Tip You can also choose **Policy** from the layer drop-down list on the Rules page for any layer, or click **Manage Rules** on the Policy Information page.

Step 2 Choose the rule or rules from which you want to remove multiple settings:

- Choose specific — If you want to choose specific rules, check the check box next to each rule.
- Choose all — If you want to choose all the rules in the current list, check the check box at the top of the column.

Step 3 Choose one of the following options:

- **Event Filtering** > **Remove Thresholds**
- **Event Filtering** > **Remove Suppressions**
- **Dynamic State** > **Remove Rate-Based Rule States**
- **Alerting** > **Remove SNMP Alerts**

Note Removing rule settings derived from a shared layer or the base policy causes any changes to this rule from lower layers or the base policy to be ignored. To stop ignoring changes from lower layers or the base policy, set the rule state to **Inherit** on the summary page for the topmost layer.

Step 4 Click **OK**.

Step 5 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

Accepting Rule Changes from a Custom Base Policy

When a custom network analysis or intrusion policy where you have not added layers uses another custom policy as its base policy, you must set a rule to inherit its rule state if:

- you delete an event filter, dynamic state, or SNMP alert that is set for the rule in the base policy, *and*
- you want the rule to accept subsequent changes that you make to it in the other custom policy that you use as your base policy

Procedure

Step 1 While editing your intrusion policy, expand **Policy Layers** in the navigation panel.

Step 2 Expand **My Changes**.

Step 3 Click the **Rules** link immediately beneath **My Changes**.

Step 4 Choose the rule or rules whose settings you want to accept. You have the following choices:

- Choose specific rules — If you want to choose specific rules, check the check box next to each rule.
- Choose all rules — If you want to choose all the rules in the current list, check the check box at the top of the column.

Step 5 Choose **Inherit** from the **Rule State** drop-down list.

Step 6 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

Preprocessors and Advanced Settings in Layers

You use similar mechanisms to configure preprocessors in a network analysis policy and advanced settings in an intrusion policy. You can enable and disable preprocessors on the network analysis Settings page and intrusion policy advanced settings on the intrusion policy Advanced Settings page. These pages also provide summaries of the effective states for all relevant features. For example, if the network analysis SSL preprocessor is disabled in one layer and enabled in a higher layer, the Settings page shows it as enabled. Changes made on these pages appear in the top layer of the policy. Note that the Back Orifice preprocessor has no user-configurable options.

You can also enable or disable preprocessors or advanced settings and access their configuration pages on the summary page for a user-configurable layer. On this page you can modify the layer name and description and configure whether to share the layer with other policies of the same type. You can switch to the summary page for another layer by selecting the layer name beneath **Policy Layers** in the navigation panel.

When you enable a preprocessor or advanced setting, a sublink to the configuration page for that feature appears beneath the layer name in the navigation panel, and an **Edit** () appears next to the feature on the summary page for the layer; these disappear when you disable the feature in the layer or set it to Inherit.

Setting the state (enabled or disabled) for a preprocessor or advanced setting overrides the state and configuration settings for that feature in lower layers. If you want a preprocessor or advanced setting to inherit its state and configuration from the base policy or a lower layer, set it to **Inherit**. Note that the Inherit selection is not available when you are working in the Settings or Advanced Settings page. Note also that if you inherit a feature that is currently enabled, the feature sublink in the navigation panel and the edit icon on the configuration page no longer appear.

The system uses the configuration in the highest layer where the feature is enabled. Unless you explicitly modify the configuration, the system uses the default configuration. For example, if you enable and modify the network analysis DCE/RPC preprocessor in one layer, and you also enable but do not modify it in a higher layer, the system uses the default configuration in the higher layer.

Color-coding on each layer summary page indicates whether the effective configuration is in a higher, lower, or the current layer, as follows:

- red—the effective configuration is in a higher layer
- yellow—the effective configuration is in a lower layer
- unshaded—the effective configuration is in the current layer

Because the Settings and Advanced Settings pages are composite views of all relevant settings, these page do not use color coding to indicate the positions of effective configurations.

Configuring Preprocessors and Advanced Settings in Layers

Procedure

Step 1 While editing your policy, expand **Policy Layers** in the navigation panel, then click the name of the layer you want to modify.

Step 2 You have the following choices:

- Change the layer **Name**.
- Add or change the **Description**.
- Check or clear the **Sharing** check box to specify whether a layer can be shared with another policy.
- To access the configuration page for an enabled preprocessor/advanced setting, click **Edit** () or the feature sublink.
- To disable a preprocessor/advanced setting in the current layer, click **Disabled** next to the feature.
- To enable a preprocessor/advanced setting in the current layer, click **Enabled** next to the feature.
- To inherit the preprocessor/advanced setting state and configuration from the settings in the highest layer below the current layer, click **Inherit**.

Step 3 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

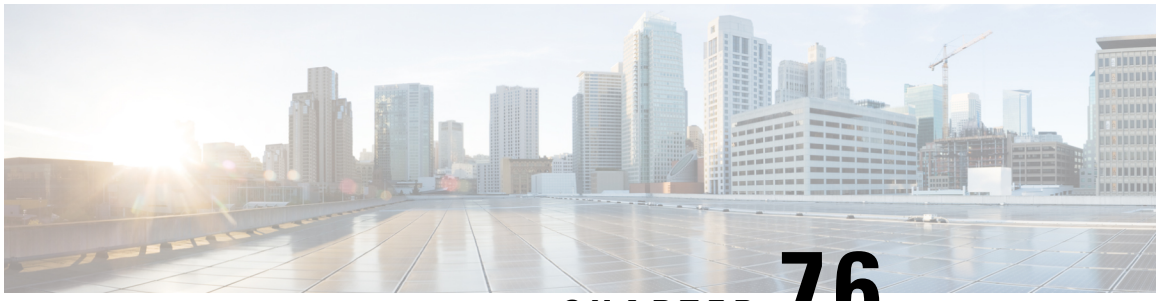
If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461



CHAPTER 76

Getting Started with Intrusion Policies

The following topics explain how to get started with intrusion policies:

- Intrusion Policy Basics, on page 1479
- License Requirements for Intrusion Policies, on page 1480
- Requirements and Prerequisites for Intrusion Policies, on page 1481
- Managing Intrusion Policies, on page 1481
- Custom Intrusion Policy Creation, on page 1482
- Editing Snort 2 Intrusion Policies, on page 1483
- Access Control Rule Configuration to Perform Intrusion Prevention, on page 1484
- Drop Behavior in an Inline Deployment, on page 1486
- Drop Behavior in a Dual System Deployment, on page 1487
- Intrusion Policy Advanced Settings, on page 1487
- Optimizing Performance for Intrusion Detection and Prevention, on page 1488

Intrusion Policy Basics

Intrusion policies are defined sets of intrusion detection and prevention configurations that inspect traffic for security violations and, in inline deployments, can block or alter malicious traffic. Intrusion policies are invoked by your access control policy and are the system's last line of defense before traffic is allowed to its destination.

At the heart of each intrusion policy are the intrusion rules. An enabled rule causes the system to generate intrusion events for (and optionally block) traffic matching the rule. Disabling a rule stops processing of the rule.

The Firepower System delivers several base intrusion policies, which enable you to take advantage of the experience of the Cisco Talos Intelligence Group (Talos). For these policies, Talos sets intrusion and preprocessor rule states (enabled or disabled), as well as provides the initial configurations for other advanced settings.

Tip System-provided intrusion and network analysis policies are similarly named but contain different configurations. For example, the Balanced Security and Connectivity network analysis policy and the Balanced Security and Connectivity intrusion policy work together and can both be updated in intrusion rule updates. However, the network analysis policy governs mostly preprocessing options, whereas the intrusion policy governs mostly intrusion rules.

If you create a custom intrusion policy, you can:

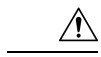
- Tune detection by enabling and disabling rules, as well as by writing and adding your own rules.
- Use Firepower recommendations to associate the operating systems, servers, and client application protocols detected on your network with rules specifically written to protect those assets.
- Configure various advanced settings such as external alerting, sensitive data preprocessing, and global rule thresholding.
- Use layers as building blocks to efficiently manage multiple intrusion policies.

In an inline deployment, an intrusion policy can block and modify traffic:

- *Drop rules* can drop matching packets and generate intrusion events. To configure an intrusion or preprocessor drop rule, set its state to Drop and Generate Events.
- Intrusion rules can use the `replace` keyword to replace malicious content.

For intrusion rules to affect traffic, you must correctly configure drop rules and rules that replace content, as well as correctly deploy managed devices inline, that is, with inline interface sets. Finally, you must enable the intrusion policy's *drop behavior*, or **Drop when Inline** setting.

When tailoring your intrusion policy, especially when enabling and adding rules, keep in mind that some intrusion rules require that traffic first be decoded or preprocessed in a certain way. Before an intrusion policy examines a packet, the packet is preprocessed according to configurations in a network analysis policy. If you disable a required preprocessor, the system automatically uses it with its current settings, although the preprocessor remains disabled in the network analysis policy web interface.



Caution Because preprocessing and intrusion inspection are so closely related, the network analysis and intrusion policies examining a single packet **must** complement each other. Tailoring preprocessing, especially using multiple custom network analysis policies, is an **advanced** task.

After you configure a custom intrusion policy, you can use it as part of your access control configuration by associating the intrusion policy with one or more access control rules or an access control policy's default action. This forces the system to use the intrusion policy to examine certain allowed traffic before the traffic passes to its final destination. A variable set that you pair with the intrusion policy allows you to accurately reflect your home and external networks and, as appropriate, the servers on your network.

Note that by default, the system disables intrusion inspection of encrypted payloads. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has intrusion inspection configured.

License Requirements for Intrusion Policies

FTD License

Threat

Classic License

Protection

Requirements and Prerequisites for Intrusion Policies

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Intrusion Admin

Managing Intrusion Policies

On the Intrusion Policy page (**Policies** > **Access Control** > **Intrusion**) you can view your current custom intrusion policies, along with the following information:

- the time and date the policy was last modified (in local time) and the user who modified it
- whether the **Drop when Inline** setting is enabled, which allows you to drop and modify traffic in an inline deployment
- which access control policies and devices are using the intrusion policy to inspect traffic
- whether a policy has unsaved changes, as well as information about who (if anyone) is currently editing the policy
- in a multidomain deployment, the domain where the policy was created

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Policies** > **Access Control** > **Intrusion**.

Step 2 Manage your intrusion policy:

- Compare—Click **Compare Policies**; see Comparing Policies, on page 355.
- Create — Click **Create Policy**; see Creating a Custom Intrusion Policy, on page 1482.
- Delete — Click **Delete** ( ) next to the policy you want to delete. The system prompts you to confirm and informs you if another user has unsaved changes in the policy. Click **OK** to confirm.

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

- Edit — Click **Edit** ( ) next to the policy you want to edit; see Editing Snort 2 Intrusion Policies, on page 1483.

  If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

- Export — If you want to export an intrusion policy to import on another Firepower Management Center, click **YouTube EDU** ( ); see Exporting Configurations, on page 175.

- Deploy—Click **Deploy**; see Deploy Configuration Changes, on page 346.

- Report—Click **Report** ( ) ; see Generating Current Policy Reports, on page 356.

Custom Intrusion Policy Creation

When you create a new intrusion policy you must give it a unique name, specify a base policy, and specify drop behavior.

The base policy defines the intrusion policy's default settings. Modifying a setting in the new policy overrides—but does not change—the settings in the base policy. You can use either a system-provided or custom policy as your base policy.

The intrusion policy's drop behavior, or **Drop when Inline** setting, determines how the system handles drop rules (intrusion or preprocessor rules whose rule state is set to Drop and Generate Events) and other intrusion policy configurations that affect traffic. You should enable drop behavior in inline deployments when you want to drop or replace malicious packets. Note that in passive deployments, the system cannot affect traffic flow regardless of the drop behavior.

Creating a Custom Intrusion Policy

Procedure

Step 1 Choose **Policies** > **Access Control** > **Intrusion**.

Step 2 Click **Create Policy**. If you have unsaved changes in another policy, click **Cancel** when prompted to return to the Intrusion Policy page.

Step 3 Enter a unique **Name** and, optionally, a **Description**.

Step 4 Choose the initial **Base Policy**.

You can use either a system-provided or another custom policy as your base policy.

Step 5 Set the system's drop behavior in an inline deployment as described in Setting Drop Behavior in an Inline Deployment, on page 1486.

Step 6 Create the policy:

- Click **Create Policy** to create the new policy and return to the Intrusion Policy page. The new policy has the same settings as its base policy.
- Click **Create and Edit Policy** to create the policy and open it for editing in the advanced intrusion policy editor; see Intrusion Policy Changes, on page 1484.

Related Topics

Intrusion Rules in Layers, on page 1472
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

Editing Snort 2 Intrusion Policies

Procedure

Step 1 Choose **Policies** > **Access Control** > **Intrusion**.

Step 2 Click **Edit** () next to the intrusion policy you want to configure.

If **View** () appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Edit your policy:

- Change the base policy—Choose a base policy from the **Base Policy** drop-down list; see Changing the Base Policy, on page 1467.
- Configure advanced settings—Click **Advanced Settings** in the navigation panel; see Intrusion Policy Advanced Settings, on page 1487.
- Configure Firepower recommended intrusion rules—Click **Firepower Recommendations** in the navigation panel; see Generating and Applying Firepower Recommendations, on page 1520.
- Drop behavior in an inline deployment—Check or clear **Drop when Inline**; see Setting Drop Behavior in an Inline Deployment, on page 1486.
- Filter rules by recommended rule state—After you generate recommendations, click **View** next to each recommendation type. Click **View Recommended Changes** to view all recommendations.
- Filter rules by current rule state—Click **View** next to each rule state type (generate events, drop and generate events); see Intrusion Rule Filters in an Intrusion Policy, on page 1497.
- Manage policy layers—Click **Policy Layers** in the navigation panel; see Layer Management, on page 1468.
- Manage intrusion rules—Click **Manage Rules**; see Viewing Intrusion Rules in an Intrusion Policy, on page 1491.
- View settings in base policy—Click **Manage Base Policy**; see The Base Layer, on page 1465.

Step 4 To save changes you made in this policy since the last policy commit, choose **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Generating and Applying Firepower Recommendations, on page 1520
Configuring Intrusion Rules in Layers, on page 1473
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

Intrusion Policy Changes

When you create a new intrusion policy, it has the same intrusion rule and advanced settings as its base policy.

The system caches one intrusion policy per user. While editing an intrusion policy, if you choose any menu or other path to another page, your changes stay in the system cache even if you leave the page.

Access Control Rule Configuration to Perform Intrusion Prevention

An access control policy can have multiple access control rules associated with intrusion policies. You can configure intrusion inspection for any Allow or Interactive Block access control rule, which permits you to match different intrusion inspection profiles against different types of traffic on your network before it reaches its final destination.

Whenever the system uses an intrusion policy to evaluate traffic, it uses an associated *variable set*. Variables in a set represent values commonly used in intrusion rules to identify source and destination IP addresses and ports. You can also use variables in intrusion policies to represent IP addresses in rule suppressions and dynamic rule states.



Tip Even if you use system-provided intrusion policies, Cisco **strongly** recommends you configure the system's intrusion variables to accurately reflect your network environment. At a minimum, modify default variables in the default set.

Understanding System-Provided and Custom Intrusion Policies

Cisco delivers several intrusion policies with the Firepower System. By using system-provided intrusion policies, you can take advantage of the experience of the Cisco Talos Intelligence Group (Talos). For these policies, Talos sets intrusion and preprocessor rule states, as well as provides the initial configurations for advanced settings. You can use system-provided policies as-is, or you can use them as the base for custom policies. Building custom policies can improve the performance of the system in your environment and provide a focused view of the malicious traffic and policy violations occurring on your network.

Connection and Intrusion Event Logging

When an intrusion policy invoked by an access control rule detects an intrusion and generates an intrusion event, it saves that event to the Firepower Management Center. The system also automatically logs the end of the connection where the intrusion occurred to the Firepower Management Center database, regardless of the logging configuration of the access control rule.

Related Topics

Predefined Default Variables, on page 414

Access Control Rule Configuration and Intrusion Policies

The number of unique intrusion policies you can use in a single access control policy depends on the model of the target devices; more powerful devices can handle more. Every unique **pair** of intrusion policy and variable set counts as one policy. Although you can associate a different intrusion policy-variable set pair with each Allow and Interactive Block rule (as well as with the default action), you cannot deploy an access control policy if the target devices have insufficient resources to perform inspection as configured.

Configuring an Access Control Rule to Perform Intrusion Prevention

You must be an Admin, Access Admin, or Network Admin to perform this task.



Caution Changing the total number of intrusion policies used by an access control policy restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information. You change the the total number of intrusion policies by adding an intrusion policy that is not currently used, or by removing the last instance of an intrusion policy. You can use an intrusion policy in an access control rule, as the default action, or as the default intrusion policy.

Procedure

Step 1 In the access control policy editor, create a new rule or edit an existing rule; see Access Control Rule Components, on page 1218.

Step 2 Ensure the rule action is set to **Allow**, **Interactive Block**, or **Interactive Block with reset**.

Step 3 Click .

Step 4 Choose a system-provided or custom **Intrusion Policy**, or choose **None** to disable intrusion inspection for traffic that matches the access control rule.

Step 5 If you want to change the variable set associated with the intrusion policy, choose a value from the **Variable Set** drop-down list.

Step 6 Click **Save** to save the rule.

Step 7 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Variable Sets, on page 412
Snort® Restart Scenarios, on page 348

Drop Behavior in an Inline Deployment

If you want to assess how your configuration would function in an inline deployment (that is, where relevant configurations are deployed to devices using routed, switched, or transparent interfaces, or inline interface pairs) without actually affecting traffic, you can disable drop behavior. In this case, the system generates intrusion events but does not drop packets that trigger the drop rules. When you are satisfied with the results, you can enable drop behavior.

Note that in passive deployments or inline deployments in tap mode, the system cannot affect traffic regardless of the drop behavior. In a passive deployment, rules set to **Drop and Generate Events** behave identically to rules set to **Generate Events**. The system generates intrusion events but cannot drop packets.

Note Suppose a file Block action causes a Block or Pending file policy verdict on a packet, and later, an IPS event is generated on the same packet. In that case, the IPS event is marked as Dropped instead of Would have dropped even if the IPS policy is in detection mode (IDS).

Note To block the transfer of malware over FTP, you must not only correctly configure AMP for Networks, but also enable **Drop when Inline** in your access control policy's default intrusion policy.

When you view intrusion events, workflows can include the *inline result*, which indicates whether traffic was actually dropped, or whether it only would have dropped.

Setting Drop Behavior in an Inline Deployment

Procedure

Step 1 Choose **Policies** > **Access Control** > **Intrusion**.

Step 2 Click **Snort 2 Version** next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Set the policy's drop behavior:

- Check the **Drop when Inline** check box to allow intrusion rules to affect traffic and generate events.
- Clear the **Drop when Inline** check box to prevent intrusion rules from affecting traffic while still generating events.

Step 4 Click **Commit Changes** to save changes you made in this policy since the last policy commit.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Drop Behavior in a Dual System Deployment

When there are two systems connected back to back in a network, it is normal to see the first system drop events and still record a drop or "would have dropped" event on the second system. The first system decides to drop the packets by the time it scans the last packet of the file, while the second system also investigates and identifies the traffic as "to be dropped".

For example, a 5 packet HTTP GET request whose first packet triggers a rule is blocked by the first system and only the last packet is dropped. The second system receives only 4 packets and the connection gets dropped, but when the second system finally flushes the partial GET request while it is pruning the session, it triggers the same rule with "would have dropped" as the inline result.

Intrusion Policy Advanced Settings

An intrusion policy's *advanced settings* require specific expertise to configure. The base policy for your intrusion policy determines which advanced settings are enabled by default and the default configuration for each.

When you choose **Advanced Settings** in the navigation panel of an intrusion policy, the policy lists its advanced settings by type. On the Advanced Settings page, you can enable or disable advanced settings in your intrusion policy, as well as access advanced setting configuration pages. An advanced setting must be enabled for you to configure it.

When you disable an advanced setting, the sublink and **Edit** link no longer appear, but your configurations are retained. Note that some intrusion policy configurations (sensitive data rules, SNMP alerts for intrusion rules) require enabled and correctly configured advanced settings.

Modifying the configuration of an advanced setting requires an understanding of the configuration you are modifying and its potential impact on your network.

Specific Threat Detection

The sensitive data preprocessor detects sensitive data such as credit card numbers and Social Security numbers in ASCII text.

Note that other preprocessors that detect specific threats (back orifice attacks, several portscan types, and rate-based attacks that attempt to overwhelm your network with excessive traffic) are configured in network analysis policies.

Intrusion Rule Thresholds

Global rule thresholding can prevent your system from being overwhelmed with a large number of events by allowing you to use thresholds to limit the number of times the system logs and displays intrusion events.

External Responses

In addition to the various views of intrusion events in the web interface, you can enable logging to system log (syslog) facilities or send event data to an SNMP trap server. Per policy, you can specify intrusion event

notification limits, set up intrusion event notification to external logging facilities, and configure external responses to intrusion events.

Note that in addition to these per-policy alerting configurations, you can globally enable or disable email alerting on intrusion events for each rule or rule group. Your email alert settings are used regardless of which intrusion policy processes a packet.

Related Topics

Sensitive Data Detection Basics, on page 1523
Global Rule Thresholding Basics, on page 1537

Optimizing Performance for Intrusion Detection and Prevention

If you want the Firepower System to perform intrusion detection and prevention but do not need to take advantage of discovery data, you can optimize performance by disabling new discovery as described below.

Before you begin

To perform this task, you must have one of the following user roles:

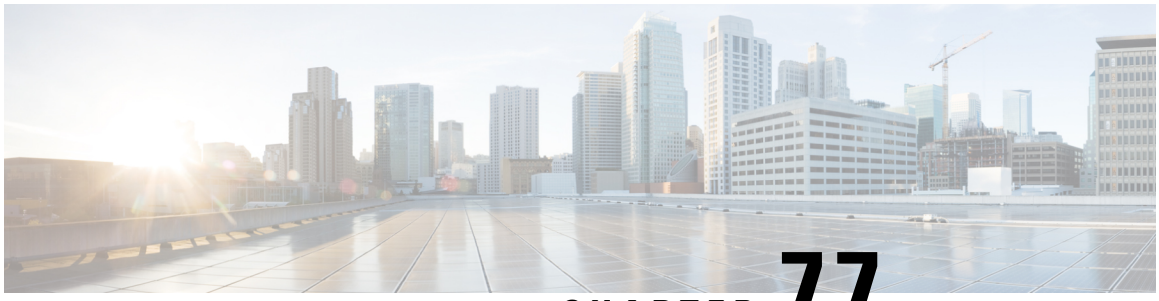
- Admin, Access Admin, or Network Admin for access control.
- Admin or Discovery Admin for network discovery.

Procedure

Step 1 Modify or delete rules associated with the access control policy deployed at the target device. None of the access control rules associated with that device can have user, application, or URL conditions; see Create and Edit Access Control Rules, on page 1220.

Step 2 Delete all rules from the network discovery policy for the target device; see Configuring Network Discovery Rules, on page 1920.

Step 3 Deploy the changed configuration to the target device; see Deploy Configuration Changes, on page 346.



CHAPTER 77

Tuning Intrusion Policies Using Rules

The following topics explain how to use rules to tune intrusion policies:

- Intrusion Rule Tuning Basics, on page 1489
- Intrusion Rule Types, on page 1489
- License Requirements for Intrusion Rules, on page 1490
- Requirements and Prerequisites for Intrusion Rules, on page 1491
- Viewing Intrusion Rules in an Intrusion Policy, on page 1491
- Intrusion Rule Filters in an Intrusion Policy, on page 1497
- Intrusion Rule States, on page 1503
- Intrusion Event Notification Filters in an Intrusion Policy, on page 1505
- Dynamic Intrusion Rule States, on page 1511
- Adding Intrusion Rule Comments, on page 1514

Intrusion Rule Tuning Basics

You can use the Rules page in an intrusion policy to configure rule states and other settings for shared object rules, standard text rules, and preprocessor rules.

You enable a rule by setting its rule state to Generate Events or to Drop and Generate Events. Enabling a rule causes the system to generate events on traffic matching the rule. Disabling a rule stops processing of the rule. You can also set your intrusion policy so that a rule set to Drop and Generate Events in an inline deployment generates events on, and drops, matching traffic. In a passive deployment, a rule set to Drop and Generate Events just generates events on matching traffic.

You can filter rules to display a subset of rules, enabling you to select the exact set of rules where you want to change rule states or rule settings.

When an intrusion rule or rule argument requires a disabled preprocessor, the system automatically uses it with its current configuration even though it remains disabled in the network analysis policy's web interface.

Intrusion Rule Types

An intrusion rule is a specified set of keywords and arguments that the system uses to detect attempts to exploit vulnerabilities in your network. As the system analyzes network traffic, it compares packets against the conditions specified in each rule, and triggers the rule if the data packet meets all the conditions specified in the rule.

An intrusion policy contains:

- *intrusion rules*, which are subdivided into *shared object rules* and *standard text rules*
- *preprocessor rules*, which are associated with a detection option of the packet decoder or with one of the preprocessors included with the Firepower System

The following table summarizes attributes of these rule types:

Table 104: Intrusion Rule Types

Type	Generator ID (GID)	Snort ID (SID)	Source	Can Copy?	Can Edit?
shared object rule	3	lower than 1000000	Cisco Talos Intelligence Group (Talos)	yes	limited
standard text rule	1	lower than 1000000	Talos	yes	limited
		1000000 or higher	Created or imported by user	yes	yes
preprocessor rule	decoder- or preprocessor-specific	lower than 1000000	Talos	no	no
		1000000 or higher	Generated by the system during option configuration	no	no

You cannot save changes to any rule created by Talos, but you can save a copy of a modified rule as a custom rule. You can modify either variables used in the rule or rule header information (such as source and destination ports and IP addresses). In a multidomain deployment, rules created by Talos belong to the Global domain. Administrators in descendant domains can save local copies of the rules, which they can then edit.

For the rules it creates, Talos assigns default rule states in each default intrusion policy. Most preprocessor rules are disabled by default and must be enabled if you want the system to generate events for preprocessor rules and, in an inline deployment, drop offending packets.

In a multidomain deployment, the system prepends a domain number to the SID of any custom rule created in or imported into a descendant domain. For example, a rule added in the Global domain would have a SID of 1000000 or greater, and rules added in descendant domains would have SIDs of [domain number]000000 or greater.

License Requirements for Intrusion Rules

FTD License

Threat

Classic License

Protection

Requirements and Prerequisites for Intrusion Rules

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Intrusion Admin

Viewing Intrusion Rules in an Intrusion Policy

You can adjust how rules are displayed in the intrusion policy, and can sort rules by several criteria. You can also display the details for a specific rule to see rule settings, rule documentation, and other rule specifics.

Procedure

Step 1 Choose **Policies** > **Access Control** > **Intrusion**.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Rules** under **Policy Information** in the navigation panel.

Step 4 While viewing the rules, you can:

- Filter the rules as described in Setting a Rule Filter in an Intrusion Policy, on page 1502.
- Sort the rules by clicking the title in the top of the column you want to sort by.
- View an intrusion rule's details as described in Viewing Intrusion Rule Details, on page 1493.
- View rules in different policy layers by choosing a layer from the **Policy** drop-down list.

Intrusion Rules Page Columns

The Intrusion Rules page uses the same icons in its menu bar and column headers. For example, the Rule State menu uses the same **Generate Events** as the Rule State column in the rule listing.

Table 105: Rules Page Columns

Heading	Description
GID	Integer that indicates the Generator ID (GID) for the rule.
SID	Integer that indicates the Snort ID (SID), which acts a unique identifier for the rule. For custom rules, the SID is 1000000 or higher. In a multidomain deployment, the system prepends a domain number to the SID of any custom rule created in or imported into a descendant domain. For example, a rule added in the Global domain would have a SID of 1000000 or greater, and rules added in descendant domains would have SIDs of [domain number]000000 or greater.
Message	Message included in events generated by this rule, which also acts as the name of the rule.
Generate Events	The rule state for the rule: • **Drop and Generate Events** • **Generate Events** • **Disabled** Note the icon for a disabled rule is a dimmed version of the icon for a rule that is set to generate events without dropping traffic. Also, clicking the rule state icon for a rule allows you to change the rule state.
Firepower Recommended rule state	Firepower recommended rule state for the rule.
Event Filter	Event filter, including event thresholds and event suppression, applied to the rule.
Dynamic state	Dynamic rule state for the rule, which goes into effect if specified rate anomalies occur.
Errors ( )	Alerts configured for the rule (currently SNMP alerts only).
Comment ( )	Comments added to the rule.

You can also use the layer drop-down list to switch to the Rules page for other layers in your policy. Note that, unless you add layers to your policy, the only editable views listed in the drop-down list are the policy Rules page and the Rules page for a policy layer that is originally named `My Changes`; note also that making changes in one of these views is the same as making the changes in the other. The drop-down list also lists the Rules page for the read-only base policy.

Intrusion Rule Details

You can view rule documentation, Firepower recommendations, and rule overhead from the Rule Detail view. You can also view and add rule-specific features.

Table 106: Rule Details

Item	Description
Summary	The rule summary. For rule-based events, this row appears when the rule documentation contains summary information.
Rule State	The current rule state for the rule. Also indicates the layer where the rule state is set.
Firepower Recommendation	If Firepower recommendations have been generated, an icon that represents the recommended rule state; see Intrusion Rules Page Columns, on page 1491. If the recommendation is to enable the rule, the system also indicates the network assets or configurations that triggered the recommendation.
Rule Overhead	The rule's potential impact on system performance and the likelihood that the rule might generate false positives. Local rules do not have an assigned overhead, unless they are mapped to a vulnerability.
Thresholds	Thresholds currently set for this rule, as well as the facility to add a threshold for the rule.
Suppressions	Suppression settings currently set for this rule, as well as the facility to add suppressions for the rule.
Dynamic State	Rate-based rule states currently set for this rule, as well as the facility to add dynamic rule states for the rule.
Alerts	SNMP alerts set for this rule, as well as the facility to add an alert for the rule.
Comments	Comments added to this rule, as well as the facility to add comments for the rule.
Documentation	The rule documentation for the current rule, supplied by the Cisco Talos Intelligence Group (Talos).

Viewing Intrusion Rule Details

Procedure

Step 1 Choose **Policies** > **Access Control** > **Intrusion**.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 On the navigation pane, click **Rules**.

Step 4 Click the rule whose rule details you want to view, then click **Show details** at the bottom of the page. Rule details appear, as described in Intrusion Rule Details, on page 1492.

Step 5 From the rule details, you can configure:

- Alerts—See Setting an SNMP Alert for an Intrusion Rule, on page 1496.

- Comments—See Adding a Comment to an Intrusion Rule, on page 1496.
- Dynamic rule states—See Setting a Dynamic Rule State from the Rule Details Page, on page 1495.
- Thresholds—See Setting a Threshold for an Intrusion Rule, on page 1494.
- Suppressions—See Setting Suppression for an Intrusion Rule, on page 1494.

Setting a Threshold for an Intrusion Rule

You can set a single threshold for a rule from the Rule Detail page. Adding a threshold overwrites any existing threshold for the rule.

Note that a **Revert** appears in a field when you enter an invalid value; click it to revert to the last valid value for that field or to clear the field if there was no previous value.

Procedure

Step 1 From an intrusion rule's details, click **Add** next to **Thresholds**.

Step 2 From the **Type** drop-down list, choose the type of threshold you want to set:

- Choose **Limit** to limit notification to the specified number of event instances per time period.
- Choose **Threshold** to provide notification for each specified number of event instances per time period.
- Choose **Both** to provide notification once per time period after a specified number of event instances.

Step 3 From the **Track By** drop-down list, choose **Source** or **Destination** to indicate whether you want the event instances tracked by source or destination IP address.

Step 4 In the **Count** field, enter the number of event instances you want to use as your threshold.

Step 5 In the **Seconds** field, enter a number that specifies the time period, in seconds, for which event instances are tracked.

Step 6 Click **OK**.

Tip The system displays an **Event Filter** next to the rule in the Event Filtering column. If you add multiple event filters to a rule, the system includes an indication of the number of event filters.

Setting Suppression for an Intrusion Rule

You can set one or more suppressions for a rule in your intrusion policy.

Note that a **Revert** appears in a field when you type an invalid value; click it to revert to the last valid value for that field or to clear the field if there was no previous value.

Procedure

Step 1 From an intrusion rule's details, click **Add** next to **Suppressions**.

Step 2 From the **Suppression Type** drop-down list, choose one of the following options:

- Choose **Rule** to completely suppress events for a selected rule.

- Choose **Source** to suppress events generated by packets originating from a specified source IP address.
- Choose **Destination** to suppress events generated by packets going to a specified destination IP address.

Step 3 If you chose **Source** or **Destination** for the suppression type, in the **Network** field enter the IP address, an address block, or a comma-separated list comprised of any combination of these.

If the intrusion policy is associated with the default action of an access control policy, you can also specify or list a network variable in the default action variable set.

The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.

Step 4 Click **OK**.

Tip The system displays an **Event Filter** next to the rule in the Event Filtering column next the suppressed rule. If you add multiple event filters to a rule, a number over the filter indicates the number of filters.

Setting a Dynamic Rule State from the Rule Details Page

You can set one or more dynamic rule states for a rule. The first dynamic rule state listed has the highest priority. When two dynamic rule states conflict, the action of the first is carried out.

Dynamic rule states are policy-specific.

Note that a **Revert** appears in a field when you enter an invalid value; click it to revert to the last valid value for that field or to clear the field if there was no previous value.

Procedure

Step 1 From an intrusion rule's details, click **Add** next to **Dynamic State**.

Step 2 From the **Track By** drop-down list, choose an option to indicate how you want the rule matches tracked:

- Choose **Source** to track the number of hits for that rule from a specific source or set of sources.
- Choose **Destination** to track the number of hits for that rule to a specific destination or set of destinations.
- Choose **Rule** to track all matches for that rule.

Step 3 If you set **Track By** to **Source** or **Destination**, enter the IP address of each host you want to track in the **Network** field.

The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.

Step 4 Next to **Rate**, specify the number of rule matches per time period to set the attack rate:

- In the **Count** field, specify the number of rule matches you want to use as your threshold.
- In the **Seconds** field, specify the number of seconds that make up the time period for which attacks are tracked.

Step 5 From the **New State** drop-down list, choose the new action to be taken when the conditions are met.

Step 6 Enter a value in the **Timeout** field.

After the timeout occurs, the rule reverts to its original state. Enter `0` to prevent the new action from timing out.

Step 7 Click **OK**.

Tip The system displays a dynamic state ( ) next to the rule in the Dynamic State column. If you add multiple dynamic rule state filters to a rule, a number over the filters indicates the number of filters.

Setting an SNMP Alert for an Intrusion Rule

You can set an SNMP alert for a rule from the Rule Detail page.

Procedure

From an intrusion rule's details, click **Add SNMP Alert** next to **Alerts**.

Tip The system displays an alert **Errors** ( ) next to the rule in the Alerting column. If you add multiple alerts to a rule, the system includes an indication of the number of alerts.

Adding a Comment to an Intrusion Rule

Procedure

Step 1 From an intrusion rule's details, click **Add** next to **Comments**.

Step 2 In the **Comment** field, enter the rule comment.

Step 3 Click **OK**.

Tip The system displays a **Comment** ( ) next to the rule in the Comments column. If you add multiple comments to a rule, a number over the comment indicates the number of comments.

Step 4 To delete a rule comment, click **Delete** in the rule comments section. You can only delete a comment if the comment is cached with uncommitted intrusion policy changes.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Intrusion Rule Filters in an Intrusion Policy

You can filter the rules you display on the Rules page by a single criteria, or a combination of one or more criteria.

Rule filter keywords help you find the rules for which you want to apply rule settings, such as rule states or event filters. You can filter by a keyword and simultaneously select the argument for the keyword by selecting the argument you want from the Rules page filter panel.

Intrusion Rule Filters Notes

The filter you construct is shown in the Filter text box. You can click keywords and keyword arguments in the filter panel to construct a filter. When you choose multiple keywords, the system combines them using AND logic to create a compound search filter. For example, if you choose **preprocessor** under **Category** and then choose **Rule Content** > **GID** and enter `116`, you get a filter of `Category: "preprocessor" GID:"116"`, which retrieves all rules that are preprocessor rules **and** have a GID of 116.

The Category, Microsoft Vulnerabilities, Microsoft Worms, Platform Specific, Preprocessor, and Priority filter groups allow you to submit more than one argument for a keyword, separated by commas. For example, you can choose **os-linux** and **os-windows** from **Category** to produce the filter `Category:"os-windows,os-linux"`, which retrieves any rules in the `os-linux` category or in the `os-windows` category.

To show the filter panel, click the **Show icon**.

To hide the filter panel, click the **Hide icon**.

Intrusion Policy Rule Filters Construction Guidelines

In most cases, when you are building a filter, you can use the filter panel to the left of the Rules page in the intrusion policy to choose the keywords/arguments you want to use.

Rule filters are grouped into rule filter groups in the filter panel. Many rule filter groups contain sub-criteria so that you can more easily find the specific rules you are looking for. Some rule filters have multiple levels that you can expand to drill down to individual rules.

Items in the filter panel sometimes represent filter type groups, sometimes represent keywords, and sometimes represent the argument to a keyword. Note the following:

- When you choose a filter type group heading that is not a keyword (Rule Configuration, Rule Content, Platform Specific, and Priority), it expands to list the available keywords.

  When you choose a keyword by clicking on a node in the criteria list, a pop-up window appears, where you supply the argument you want to filter by.

  If that keyword is already used in the filter, the argument you supply replaces the existing argument for that keyword.

  For example, if you click **Drop and Generate Events** under **Rule Configuration** > **Recommendation** in the filter panel, `Recommendation:"Drop and Generate Events"` is added to the filter text box. If you then click **Generate Events** under Rule **Configuration** > **Recommendation**, the filter changes to `Recommendation:"Generate Events"`.

- When you choose a filter type group heading that is a keyword (Category, Classifications, Microsoft Vulnerabilities, Microsoft Worms, Priority, and Rule Update), it lists the available arguments.

  When you choose an item from this type of group, the argument and the keyword it applies to are immediately added to the filter. If the keyword is already in the filter, it replaces the existing argument for the keyword that corresponds to that group.

  For example, if you click **os-linux** under **Category** in the filter panel, `Category:"os-linux"` is added to the filter text box. If you then click **os-windows** under **Category**, the filter changes to `Category:"os-windows"`.

- Reference under Rule Content is a keyword, and so are the specific reference ID types listed below it. When you choose any of the reference keywords, a pop-up window appears, where you supply an argument and the keyword is added to the existing filter. If the keyword is already in use in the filter, the new argument you supply replaces the existing argument.

  For example, if you click **Rule Content** > **Reference** > **CVE ID** in the filter panel, a pop-up window prompts you to supply the CVE ID. If you enter `2007`, then `CVE:"2007"` is added to the filter text box. In another example, if you click **Rule Content** > **Reference**in the filter panel, a pop-up window prompts you to supply the reference. If you enter `2007`, then `Reference:"2007"` is added to the filter text box.

- When you choose rule filter keywords from different groups, each filter keyword is added to the filter and any existing keywords are maintained (unless overridden by a new value for the same keyword).

  For example, if you click **os-linux** under **Category** in the filter panel, `Category:"os-linux"` is added to the filter text box. If you then click **MS00-006** under **Microsoft Vulnerabilities**, the filter changes to `Category:"os-linux" MicrosoftVulnerabilities:"MS00-006"`.

- When you choose multiple keywords, the system combines them using AND logic to create a compound search filter. For example, if you choose **preprocessor** under **Category** and then choose **Rule Content** > **GID** and enter `116`, you get a filter of `Category: "preprocessor" GID:"116"`, which retrieves all rules that are preprocessor rules **and** have a GID of 116.

- The Category, Microsoft Vulnerabilities, Microsoft Worms, Platform Specific, and Priority filter groups allow you to submit more than one argument for a keyword, separated by commas. For example, you can choose **os-linux** and **os-windows** from **Category** to produce the filter `Category:"os-windows,app-detect"`, which retrieves any rules in the `os-linux` category or in the `os-windows` category.

The same rule may be retrieved by more than one filter keyword/argument pair. For example, the DOS Cisco attempt rule (SID 1545) appears if rules are filtered by the **dos** category, and also if you filter by the **High** priority.



Note The Cisco Talos Intelligence Group (Talos) may use the rule update mechanism to add and remove rule filters.

Note that the rules on the Rules page may be either shared object rules (generator ID 3) or standard text rules (generator ID 1). The following table describes the different rule filters.

Table 107: Rule Filter Groups

Filter Group	Description	Multiple Argument Support?	Heading is...	Items in List are...
Rule Configuration	Finds rules according to the configuration of the rule.	No	A grouping	keywords
Rule Content	Finds rules according to the content of the rule.	No	A grouping	keywords
Category	Finds rules according to the rule categories used by the rule editor. Note that local rules appear in the local sub-group.	Yes	A keyword	arguments
Classifications	Finds rules according to the attack classification that appears in the packet display of an event generated by the rule.	No	A keyword	arguments
Microsoft Vulnerabilities	Finds rules according to Microsoft bulletin number.	Yes	A keyword	arguments
Microsoft Worms	Finds rules based on specific worms that affect Microsoft Windows hosts.	Yes	A keyword	arguments
Platform Specific	Finds rules according to their relevance to specific versions of operating systems. Note that a rule may affect more than one operating system or more than one version of an operating system. For example, enabling SID 2260 affects multiple versions of Mac OS X, IBM AIX, and other operating systems.	Yes	A keyword	arguments Note that if you pick one of the items from the sub-list, it adds a modifier to the argument.
Preprocessors	Finds rules for individual preprocessors. Note that you must enable preprocessor rules associated with a preprocessor option to generate events and, in an inline deployment, drop offending packets for the option when the preprocessor is enabled.	Yes	A grouping	sub-groupings
Priority	Finds rules according to high, medium, and low priorities. The classification assigned to a rule determines its priority. These groups are further grouped into rule categories. Note that local rules (that is, rules that you import or create) do not appear in the priority groups.	Yes	A keyword	arguments Note that if you pick one of the items from the sub-list, it adds a modifier to the argument.
Rule Update	Finds rules added or modified through a specific rule update. For each rule update, view all rules in the update, only new rules imported in the update, or only existing rules changed by the update.	No	A keyword	arguments

Intrusion Rule Configuration Filters

You can filter the rules listed in the Rules page by several rule configuration settings. For example, if you want to view the set of rules whose rule state does not match the recommended rule state, you can filter on rule state by selecting **Does not match recommendation**.

When you choose a keyword by clicking on a node in the criteria list, you can supply the argument you want to filter by. If that keyword is already used in the filter, the argument you supply replaces the existing argument for that keyword.

For example, if you click **Drop and Generate Events** under **Rule Configuration** > **Recommendation** in the filter panel, `Recommendation:"Drop and Generate Events"` is added to the filter text box. If you then click **Generate Events** under **Rule Configuration** > **Recommendation**, the filter changes to `Recommendation:"Generate Events"`.

Intrusion Rule Content Filters

You can filter the rules listed in the Rules page by several rule content items. For example, you can quickly retrieve a rule by searching for the rule's SID. You can also find all rules that inspect traffic going to a specific destination port.

When you select a keyword by clicking on a node in the criteria list, you can supply the argument you want to filter by. If that keyword is already used in the filter, the argument you supply replaces the existing argument for that keyword.

For example, if you click **SID** under **Rule Content** in the filter panel, a pop-up window appears, prompting you to supply a SID. If you type `1045`, then `SID:"1045"`is added to the filter text box. If you then click **SID** again and change the SID filter to `1044`, the filter changes to `SID:"1044"`.

Table 108: Rule Content Filters

This filter...	Finds rules that...
Message	contain the supplied string in the message field.
SID	have the specified SID.
GID	have the specified GID.
Reference	contain the supplied string in the reference field. You can also filter by a specific type of reference and supplied string.
Action	start with `alert` or `pass`.
Protocol	include the selected protocol.
Direction	are based on whether the rule includes the indicated directional setting.
Source IP	use the specified addresses or variables for the source IP address designation in the rule. You can filter by a valid IP address, a CIDR block/prefix length, or using variables such as `$HOME_NET` or `$EXTERNAL_NET`.
Destination IP	use the specified addresses or variables for the source IP address designation in the rule. You can filter by a valid IP address, a CIDR block/prefix length, or using variables such as `$HOME_NET` or `$EXTERNAL_NET`.

This filter...	Finds rules that...
Source port	include the specified source port. The port value must be an integer between 1 and 65535 or a port variable.
Destination port	include the specified destination port. The port value must be an integer between 1 and 65535 or a port variable.
Rule Overhead	have the selected rule overhead.
Metadata	have metadata containing the matching *key value* pair. For example, type `metadata:"service http"` to locate rules with metadata relating to the HTTP application protocol.

Intrusion Rule Categories

The Firepower System places rules in categories based on the type of traffic the rule detects. On the Rules page, you can filter by rule category, so you can set a rule attribute for all rules in a category. For example, if you do not have Linux hosts on your network, you could filter by the **os-linux** category, then disable all the rules showing to disable the entire **os-linux** category.

You can hover your pointer over a category name to display the number of rules in that category.



Note The Cisco Talos Intelligence Group (Talos) may use the rule update mechanism to add and remove rule categories.

Intrusion Rule Filter Components

You can edit your filter to modify the special keywords and their arguments that are supplied when you click on a filter in the filter panel. Custom filters on the Rules page function like those used in the rule editor, but you can also use any of the keywords supplied in the Rules page filter, using the syntax displayed when you select the filter through the filter panel. To determine a keyword for future use, click on the appropriate argument in the filter panel on the right. The filter keyword and argument syntax appear in the filter text box. Remember that comma-separated multiple arguments for a keyword are only supported for the Category and Priority filter types.

You can use keywords and arguments, character strings, and literal character strings in quotes, with spaces separating multiple filter conditions. A filter cannot include regular expressions, wild card characters, or any special operator such as a negation character (!), a greater than symbol (>), less than symbol (<), and so on. When you type in search terms without a keyword, without initial capitalization of the keyword, or without quotes around the argument, the search is treated as a string search and the category, message, and SID fields are searched for the specified terms.

Except for the `gid` and `sid` keywords, all arguments and strings are treated as partial strings. Arguments for `gid` and `sid` return only exact matches.

Each rule filter can include one or more keywords in the format:

keyword:"*argument*"

where keyword is one of the keywords in the intrusion rule filter groups and argument is enclosed in double quotes and is a single, case-insensitive, alphanumeric string to search for in the specific field or fields relevant to the keyword. Note that keywords should be typed with initial capitalization.

Arguments for all keywords except `gid` and `sid` are treated as partial strings. For example, the argument `123` returns `"12345"`, `"41235"`, `"45123",` and so on. The arguments for `gid` and `sid` return only exact matches; for example, `sid:3080` returns only `SID 3080`.

Each rule filter can also include one or more alphanumeric character strings. Character strings search the rule Message field, Snort ID (SID), and Generator ID (GID). For example, the string `123` returns the strings `"Lotus123"`, `"123mania"`, and so on in the rule message, and also returns `SID 6123`, `SID 12375`, and so on. You can search for a partial SID by filtering with one or more character strings.

All character strings are case-insensitive and are treated as partial strings. For example, any of the strings `ADMIN`, `admin`, or `Admin` return `"admin"`, `"CFADMIN"`, `"Administrator"` and so on.

You can enclose character strings in quotes to return exact matches. For example, the literal string `"overflow attempt"` in quotes returns only that exact string, whereas a filter comprised of the two strings `overflow` and `attempt` without quotes returns `"overflow attempt"`, `"overflow multipacket attempt"`, `"overflow with evasion attempt"`, and so on.

You can narrow filter results by entering any combination of keywords, character strings, or both, separated by spaces. The result includes any rule that matches all the filter conditions.

You can enter multiple filter conditions in any order. For example, each of the following filters returns the same rules:

- `url:at login attempt cve:200`
- `login attempt cve:200 url:at`
- `login cve:200 attempt url:at`

Intrusion Rule Filter Usage

You can select predefined filter keywords from the filter panel on the left side of the Rules page in the intrusion policy. When you select a filter, the page displays all matching rules, or indicates when no rules match.

You can add keywords to a filter to further constrain it. Any filter you enter searches the entire rules database and returns all matching rules. When you enter a filter while the page still displays the result of a previous filter, the page clears and returns the result of the new filter instead.

You can also type a filter using the same keyword and argument syntax supplied when you select a filter, or modify argument values in a filter after you select it. When you type in search terms without a keyword, without initial capitalization of the keyword, or without quotes around the argument, the search is treated as a string search and the category, message, and SID fields are searched for the specified terms.

Setting a Rule Filter in an Intrusion Policy

You can filter the rules on the Rules page to display a subset of rules. You can then use any of the page features, including choosing any of the features available in the context menu. This can be useful, for example, when you want to set a threshold for all the rules in a specific category. You can use the same features with rules in a filtered or unfiltered list. For example, you can apply new rule states to rules in a filtered or unfiltered list.

All filter keywords, keyword arguments, and character strings are case-insensitive. If you click an argument for a keyword already in the filter, it replaces the existing argument.

Procedure

Step 1 Choose **Policies** > **Access Control** > **Intrusion**.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Rules**.

Step 4 Construct a filter using any of the following methods, separately or in combination:

- Enter a value in the **Filter** text box, and press Enter.
- Expand any of the predefined keywords. For example, click **Rule Configuration**.
- Click a keyword, and specify an argument value if prompted. For example:
  - Under **Rule Configuration**, you could click **Rule State**, choose `Generate Events` from the drop-down-list, and click **OK**.
  - Under **Rule Configuration**, you could click **Comment**, enter the string of comment text to filter by, and click **OK**.
  - Under **Category**, you could click **app-detect**, which the system uses as the argument value.
- Expand a keyword, and click an argument value. For example, expand **Rule State** and click **Generate Events**.

Intrusion Rule States

Intrusion rule states allow you to enable or disable the rule within an individual intrusion policy, as well as specify which action the system takes if monitored conditions trigger the rule.

The Cisco Talos Intelligence Group (Talos) sets the default state of each intrusion and preprocessor rule in each default policy. For example, a rule may be enabled in the Security over Connectivity default policy and disabled in the Connectivity over Security default policy. Talos sometimes uses a rule update to change the default state of one or more rules in a default policy. If you allow rule updates to update your base policy, you also allow the rule update to change the default state of a rule in your policy when the default state changes in the default policy you used to create your policy (or in the default policy it is based on). Note, however, that if you have changed the rule state, the rule update does not override your change.

When you create an intrusion rule, it inherits the default states of the rules in the default policy you use to create your policy.

Intrusion Rule State Options

In an intrusion policy, you can set a rule's state to the following values:

Generate Events

You want the system to detect a specific intrusion attempt and generate an intrusion event when it finds matching traffic. When a malicious packet crosses your network and triggers the rule, the packet is sent to its destination and the system generates an intrusion event. The malicious packet reaches its target, but you are notified via the event logging.

Drop and Generate Events

You want the system to detect a specific intrusion attempt, drop the packet containing the attack, and generate an intrusion event when it finds matching traffic. The malicious packet never reaches its target, and you are notified via the event logging.

Note that rules set to this rule state generate events but do not drop packets in a passive deployment, including deployments where a 7000 or 8000 Series device inline interface set is in tap mode. For the system to drop packets, **Drop when Inline** must also be enabled (the default setting) in your intrusion policy and you must deploy your device inline.

Disable

You do not want the system to evaluate matching traffic.

Note Choosing either the **Generate Events** or **Drop and Generate Events** options enables the rule. Choosing **Disable** disables the rule.

Cisco **strongly** recommends that you **do not** enable all the intrusion rules in an intrusion policy. The performance of your managed device is likely to degrade if all rules are enabled. Instead, tune your rule set to match your network environment as closely as possible.

Setting Intrusion Rule States

Intrusion rule states are policy-specific.

Procedure

Step 1 Choose **Policies** > **Access Control** > **Intrusion**.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Tip This page indicates the total number of enabled rules, the total number of enabled rules set to Generate Events, and the total number set to Drop and Generate Events. Note also that in a passive deployment, rules set to Drop and Generate Events only generate events.

Step 3 Click **Rules** immediately under **Policy Information** in the navigation panel.

Step 4 Choose the rule or rules where you want to set the rule state.

Step 5 Choose one of the following:

- **Rule State** > **Generate Events**

- **Rule State** > **Drop and Generate Events**
- **Rule State** > **Disable**

Step 6 To save changes you made in this policy since the last policy commit, click **Policy Information** in the navigation panel, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Intrusion Event Notification Filters in an Intrusion Policy

The importance of an intrusion event can be based on frequency of occurrence, or on source or destination IP address. In some cases you may not care about an event until it has occurred a certain number of times. For example, you may not be concerned if someone attempts to log into a server until they fail a certain number of times. In other cases, you may only need to see a few occurrences to know there is a widespread problem. For example, if a DoS attack is launched against your web server, you may only need to see a few occurrences of an intrusion event to know that you need to address the situation. Seeing hundreds of the same event only overwhelms your system.

Intrusion Event Thresholds

You can set thresholds for individual rules, per intrusion policy, to limit the number of times the system logs and displays an intrusion event based on how many times the event is generated within a specified time period. This can prevent you from being overwhelmed with a large number of identical events. You can set thresholds per shared object rule, standard text rule, or preprocessor rule.

Intrusion Event Thresholds Configuration

To set a threshold, first specify the thresholding type.

Table 109: Thresholding Options

Option	Description
Limit	Logs and displays events for the specified number of packets (specified by the Count argument) that trigger the rule during the specified time period. For example, if you set the type to **Limit**, the **Count** to `10`, and the **Seconds** to `60`, and 14 packets trigger the rule, the system stops logging events for the rule after displaying the first 10 that occur within the same minute.

Option	Description
Threshold	Logs and displays a single event when the specified number of packets (specified by the Count argument) trigger the rule during the specified time period. Note that the counter for the time restarts after you hit the threshold count of events and the system logs that event. For example, you set the type to **Threshold**, **Count** to `10`, and **Seconds** to `60`, and the rule triggers 10 times by second 33. The system generates one event, then resets the Seconds and Count counters to 0. The rule then triggers another 10 times in the next 25 seconds. Because the counters reset to 0 at second 33, the system logs another event.
Both	Logs and displays an event once per specified time period, after the specified number (count) of packets trigger the rule. For example, if you set the type to **Both**, **Count** to two, and **Seconds** to `10`, the following event counts result: • If the rule is triggered once in 10 seconds, the system does not generate any events (the threshold is not met) • If the rule is triggered twice in 10 seconds, the system generates one event (the threshold is met when the rule triggers the second time) • If the rule is triggered four times in 10 seconds, the system generates one event (the threshold is met when the rule triggers the second time, and following events are ignored)

Next, specify tracking, which determines whether the event threshold is calculated per source or destination IP address.

Table 110: Thresholding IP Options

Option	Description
Source	Calculates event instance count per source IP address.
Destination	Calculates event instance count per destination IP address.

Finally, specify the number of instances and time period that define the threshold.

Table 111: Thresholding Instance/Time Options

Option	Description
Count	The number of event instances per specified time period per tracking IP address required to meet the threshold.
Seconds	The number of seconds that elapse before the count resets. If you set the threshold type to **limit**, the tracking to **Source IP**, the **count** to `10`, and the **seconds** to `10`, the system logs and displays the first 10 events that occur in 10 seconds from a given source port. If only 7 events occur in the first 10 seconds, the system logs and displays those; if 40 events occur in the first 10 seconds, the system logs and displays 10, then begins counting again when the 10-second time period elapses.

Note that you can use intrusion event thresholding alone or in any combination with rate-based attack prevention, the `detection_filter` keyword, and intrusion event suppression.

Tip You can also add thresholds from within the packet view of an intrusion event.

Related Topics

The detection_filter Keyword, on page 1638
Setting Threshold Options within the Packet View, on page 2278

Adding and Modifying Intrusion Event Thresholds

You can set a threshold for one or more specific rules in an intrusion policy. You can also separately or simultaneously modify existing threshold settings. You can set a single threshold for each. Adding a threshold overwrites any existing threshold for the rule.

You can also modify the global threshold that applies by default to all rules and preprocessor-generated events associated with the intrusion policy.

A **Revert** appears in a field when you enter an invalid value; click it to revert to the last valid value for that field or to clear the field if there was no previous value.

Tip A global or individual threshold on a managed device with multiple CPUs may result in a higher number of events than expected.

Procedure

Step 1 Choose **Policies** > **Access Control** > **Intrusion**.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Rules** immediately under **Policy Information** in the navigation pane.

Step 4 Choose the rule or rules where you want to set a threshold.

Step 5 Choose **Event Filtering** > **Threshold**.

Step 6 Choose a threshold type from the **Type** drop-down list.

Step 7 From the **Track By** drop-down list, choose whether you want the event instances tracked by **Source** or **Destination** IP address.

Step 8 Enter a value in the **Count** field.

Step 9 Enter a value in the **Seconds** field.

Step 10 Click **OK**.

Tip The system displays an **Event Filter** next to the rule in the Event Filtering column. If you add multiple event filters to a rule, a number over the filter indicates the number of event filters.

Step 11 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Global Rule Thresholding Basics, on page 1537

Viewing and Deleting Intrusion Event Thresholds

You may want to view or delete an existing threshold setting for a rule. You can use the Rules Details view to display the configured settings for a threshold to see if they are appropriate for your system. If they are not, you can add a new threshold to overwrite the existing values.

Note that you can also modify the global threshold that applies by default to all rules and preprocessor-generated events logged by the intrusion policy.

Procedure

Step 1 Choose **Policies** > **Access Control** > **Intrusion**.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Rules** immediately under **Policy Information** in the navigation pane.

Step 4 Choose the rule or rules with a configured threshold you want to view or delete.

Step 5 To remove the threshold for each selected rule, choose **Event Filtering** > **Remove Thresholds**.

Step 6 Click **OK**.

Step 7 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Global Rule Thresholding Basics, on page 1537

Intrusion Policy Suppression Configuration

You can suppress intrusion event notification when a specific IP address or range of IP addresses triggers a specific rule or preprocessor. This is useful for eliminating false positives. For example, if you have a mail server that transmits packets that look like a specific exploit, you might suppress event notification for that event when it is triggered by your mail server. The rule triggers for all packets, but you only see events for legitimate attacks.

Intrusion Policy Suppression Types

Note that you can use intrusion event suppression alone or in any combination with rate-based attack prevention, the `detection_filter` keyword, and intrusion event thresholding.

Tip You can add suppressions from within the packet view of an intrusion event. You can also access suppression settings by using the right-click context menu on the intrusion rules editor page (**Objects** > **Intrusion Rules**) and on any intrusion event page (if the event was triggered by an intrusion rule).

Related Topics

The detection_filter Keyword, on page 1638
Setting Threshold Options within the Packet View, on page 2278

Suppressing Intrusion Events for a Specific Rule

You can suppress intrusion event notification for a rule or rules in your intrusion policy. When notification is suppressed for a rule, the rule triggers but events are not generated. You can set one or more suppressions for a rule. The first suppression listed has the highest priority. When two suppressions conflict, the action of the first is carried out.

Note that a **Revert** appears in a field when you enter an invalid value; click it to revert to the last valid value for that field or to clear the field if there was no previous value.

Procedure

Step 1 Choose **Policies** > **Access Control** > **Intrusion**.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Rules** immediately under **Policy Information** in the navigation panel.

Step 4 Choose the rule or rules for which you want to configure suppression conditions.

Step 5 Choose **Event Filtering** > **Suppression**.

Step 6 Choose a **Suppression Type**.

Step 7 If you chose **Source** or **Destination** for the suppression type, in the **Network** field enter the IP address, address block, or variable you want to specify as the source or destination IP address, or a comma-separated list comprised of any combination of these.

The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.

Step 8 Click **OK**.

Tip The system displays an **Event Filter** next to the rule in the Event Filtering column next the suppressed rule. If you add multiple event filters to a rule, a number over the filter indicates the number of event filters.

Step 9 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Viewing and Deleting Suppression Conditions

You may want to view or delete an existing suppression condition. For example, you can suppress event notification for packets originating from a mail server IP address because the mail server normally transmits packets that look like exploits. If you then decommission that mail server and reassign the IP address to another host, you should delete the suppression conditions for that source IP address.

Procedure

Step 1 Choose **Policies** > **Access Control** > **Intrusion**.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Rules** immediately under **Policy Information** in the navigation panel.

Step 4 Choose the rule or rules for which you want to view or delete suppressions.

Step 5 You have the following choices:

- To remove all suppression for a rule, choose **Event Filtering** > **Remove Suppressions**.
- To remove a specific suppression setting, click the rule, then click **Show details**. Expand the suppression settings and click **Delete** next to the suppression settings you want to remove.

Step 6 Click **OK**.

Step 7 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Dynamic Intrusion Rule States

Rate-based attacks attempt to overwhelm a network or host by sending excessive traffic toward the network or host, causing it to slow down or deny legitimate requests. You can use rate-based prevention to change the action of a rule in response to excessive rule matches for specific rules.

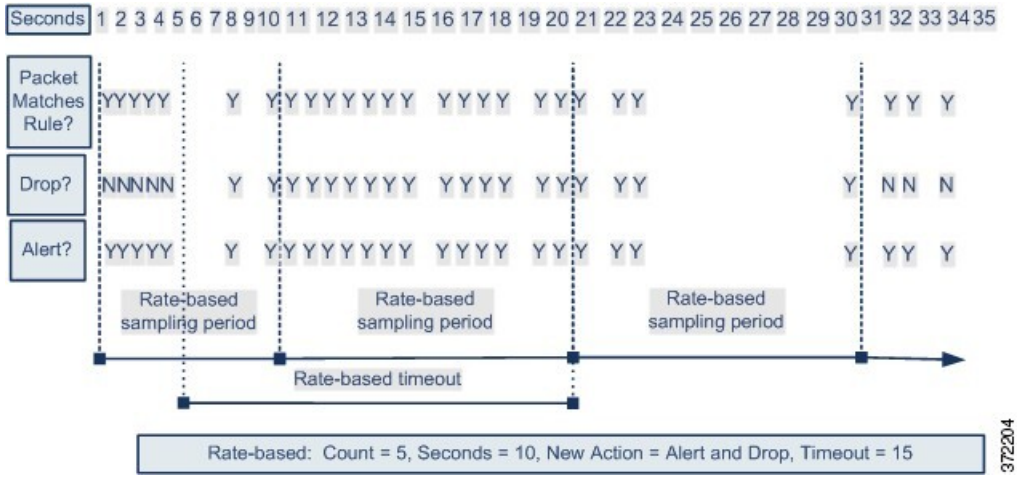
You can configure your intrusion policies to include a rate-based filter that detects when too many matches for a rule occur in a given time period. You can use this feature on managed devices deployed inline to block rate-based attacks for a specified time, then revert to a rule state where rule matches only generate events and do not drop traffic.

Rate-based attack prevention identifies abnormal traffic patterns and attempts to minimize the impact of that traffic on legitimate requests. You can identify excessive rule matches in traffic going to a particular destination IP address or addresses or coming from a particular source IP address or addresses. You can also respond to excessive matches for a particular rule across all detected traffic.

In some cases, you may not want to set a rule to the Drop and Generate Events state because you do not want to drop every packet that matches the rule, but you do want to drop packets matching the rule if a particular rate of matches occurs in a specified time. Dynamic rule states let you configure the rate that should trigger a change in the action for a rule, what the action should change to when the rate is met, and how long the new action should persist.

The following diagram shows an example where an attacker is attempting to access a host. Repeated attempts to find a password trigger a rule which has rate-based attack prevention configured. The rate-based settings change the rule attribute to Drop and Generate Events after rule matches occur five times in a 10-second span. The new rule attribute times out after 15 seconds.

After the timeout, note that packets are still dropped in the rate-based sampling period that follows. If the sampled rate is above the threshold in the current or previous sampling period, the new action continues. The new action reverts to Generate Events only after a sampling period completes where the sampled rate was below the threshold rate.



Dynamic Intrusion Rule State Configuration

In the intrusion policy, you can configure a rate-based filter for any intrusion or preprocessor rule. The rate-based filter contains three components:

- the rule matching rate, which you configure as a count of rule matches within a specific number of seconds
- a new action to be taken when the rate is exceeded, with three available actions: Generate Events, Drop and Generate Events, and Disable
- the duration of the action, which you configure as a timeout value

Note that when started, the new action occurs until the timeout is reached, even if the rate falls below the configured rate during that time period. When the timeout is reached, if the rate has fallen below the threshold, the action for the rule reverts to the action initially configured for the rule.

You can configure rate-based attack prevention in an inline deployment to block attacks, either temporarily or permanently. Without rate-based configuration, rules set to Generate Events do generate events, but the system does not drop packets for those rules. However, if the attack traffic matches rules that have rate-based criteria configured, the rate action may cause packet dropping to occur for the period of time that the rate action is active, even if those rules are not initially set to Drop and Generate Events.

Note Rate-based actions cannot enable disabled rules or drop traffic that matches disabled rules.

You can define multiple rate-based filters on the same rule. The first filter listed in the intrusion policy has the highest priority. Note that when two rate-based filter actions conflict, the action of the first rate-based filter is carried out.

Setting a Dynamic Rule State from the Rules Page

You can set one or more dynamic rule states for a rule. The first dynamic rule state listed has the highest priority. When two dynamic rule states conflict, the action of the first is carried out.

Dynamic rule states are policy-specific.

A **Revert** appears in a field when you enter an invalid value; click it to revert to the last valid value for that field or to clear the field if there was no previous value.

Note Dynamic rule states cannot enable disabled rules or drop traffic that matches disabled rules.

Procedure

Step 1 Choose **Policies** > **Access Control** > **Intrusion**.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Rules** immediately under **Policy Information** in the navigation pane.

Step 4 Choose the rule or rules where you want to add a dynamic rule state.

Step 5 Choose **Dynamic State** > **Add Rate-Based Rule State**.

Step 6 Choose a value from the **Track By** drop-down list.

Step 7 If you set **Track By** to **Source** or **Destination**, enter the address of each host you want to track in the **Network** field. You can specify a single IP address, address block, variable, or a comma-separated list comprised of any combination of these.

The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.

Step 8 Next to **Rate**, specify the number of rule matches per time period to set the attack rate:

- Enter a value in the **Count** field.
- Enter a value in the **Seconds** field.

Step 9 From the **New State** drop-down list, specify the new action to be taken when the conditions are met.

Step 10 Enter a value in the **Timeout** field.

After the timeout occurs, the rule reverts to its original state. Specify `0` or leave the **Timeout** field blank to prevent the new action from timing out.

Step 11 Click **OK**.

Tip The system displays a **Dynamic State** next to the rule in the Dynamic State column. If you add multiple dynamic rule state filters to a rule, a number over the filter indicates the number of filters.

Tip To delete all dynamic rule settings for a set of rules, choose the rules on the Rules page, then choose **Dynamic State** > **Remove Rate-Based States**. You can also delete individual rate-based rule state filters from the rule details for the rule by choosing the rule, clicking **Show details**, then clicking **Delete** by the rate-based filter you want to remove.

Step 12 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Adding Intrusion Rule Comments

You can add comments to rules in your intrusion policy. Comments added this way are policy-specific; that is, comments you add to a rule in one intrusion policy are not visible in other intrusion policies. Any comments you add can be seen in the Rule Details view on the Rules page for the intrusion policy.

After you commit the intrusion policy changes containing the comment, you can also view the comment by clicking **Rule Comment** on the rule Edit page.

Procedure

Step 1 Choose **Policies** > **Access Control** > **Intrusion**.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Rules** immediately under **Policy Information** in the navigation panel.

Step 4 Choose the rule or rules where you want to add a comment.

Step 5 Choose **Comments** > **Add Rule Comment**.

Step 6 In the **Comment** field, enter the rule comment.

Step 7 Click **OK**.

Tip The system displays a **Comment** ( ) next to the rule in the Comments column. If you add multiple comments to a rule, a number over the comment indicates the number of comments.

Step 8 Optionally, delete a rule comment by clicking **Delete** next to the comment.

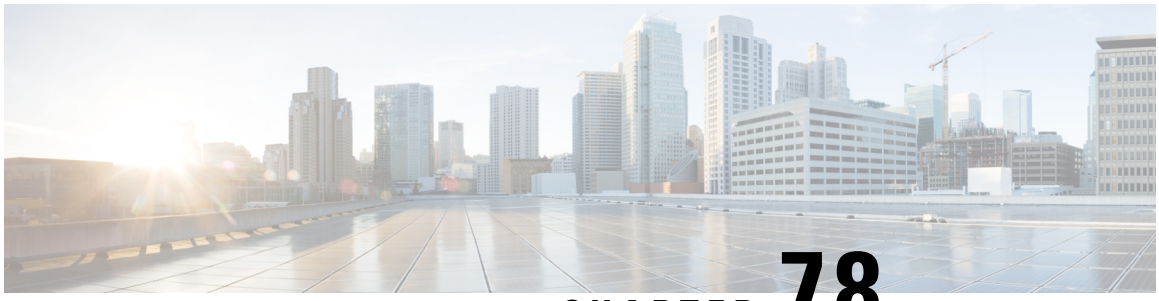
You can only delete a comment if the comment is cached with uncommitted intrusion policy changes. After intrusion policy changes are committed, the rule comment is permanent.

Step 9 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.



CHAPTER 78

Tailoring Intrusion Protection to Your Network Assets

The following topics describe how to use Firepower recommended rules:

- About Firepower Recommended Rules, on page 1517
- Default Settings for Firepower Recommendations, on page 1518
- Advanced Settings for Firepower Recommendations, on page 1519
- Generating and Applying Firepower Recommendations, on page 1520

About Firepower Recommended Rules

You can use Firepower intrusion rule recommendations to associate the operating systems, servers, and client application protocols detected on your network with rules specifically written to protect those assets. This allows you to tailor your intrusion policy to the specific needs of your monitored network.

The system makes an individual set of recommendations for each intrusion policy. It typically recommends rule state changes for standard text rules and shared object rules. However, it can also recommend changes for preprocessor and decoder rules.

When you generate rule state recommendations, you can use the default settings or configure advanced settings. Advanced settings allow you to:

- Redefine which hosts on your network the system monitors for vulnerabilities
- Influence which rules the system recommends based on rule overhead
- Specify whether to generate recommendations to disable rules

You can also choose either to use the recommendations immediately or to review the recommendations (and affected rules) before accepting them.

Choosing to use recommended rule states adds a read-only Firepower Recommendations layer to your intrusion policy, and subsequently choosing not to use recommended rule states removes the layer.

You can schedule a task to generate recommendations automatically based on the most recently saved configuration settings in your intrusion policy.

The system does not change rule states that you set manually:

- Manually setting the states of specified rules *before* you generate recommendations prevents the system from modifying the states of those rules in the future.
- Manually setting the states of specified rules *after* you generate recommendations overrides the recommended states of those rules.

Tip The intrusion policy report can include a list of rules with rule states that differ from the recommended state.

While displaying the recommendation-filtered Rules page, or after accessing the Rules page directly from the navigation panel or the Policy Information page, you can manually set rule states, sort rules, and take any of the other actions available on the Rules page, such as suppressing rules, setting rule thresholds, and so on.

Note The Cisco Talos Intelligence Group (Talos) determines the appropriate state of each rule in the system-provided policies. If you use a system-provided policy as your base policy, and you allow the system to set your rules to the Firepower recommended rule state, the rules in your intrusion policy match the settings recommended by Cisco for your network assets.

Recommended Rules and Multitenancy

The system builds a separate network map for each leaf domain. In a multidomain deployment, if you enable this feature in an intrusion policy in an ancestor domain, the system generates recommendations using data from all descendant leaf domains. This can enable intrusion rules tailored to assets that may not exist in all leaf domains, which can affect performance.

Default Settings for Firepower Recommendations

When you generate Firepower recommendations, the system searches your base policy for rules that protect against vulnerabilities associated with your network assets, and identifies the current state of rules in your base policy. The system then recommends rule states and, if you choose to, sets the rules to the recommended states.

The system performs the following basic analysis to generate recommendations:

Table 112: Rule State Recommendations Based on Vulnerabilities

Rule Protects Discovered Assets?	Base Policy Rule State	Recommend Rule State
Yes	Disabled	Generate Events
	Generate Events	Generate Events
	Drop and Generate Events	Drop and Generate Events
No	Any	Disabled

Note the following in the table:

- If a rule is disabled in the base policy, or set to Generate Events, the recommended state is always Generate Events.

  For example, if the base policy is No Rules Active, in which all rules are disabled, there will be no recommendations to Drop and Generate Events.

- Recommendations to Drop and Generate Events are made only for rules already set to Drop and Generate Events in the base policy.

  If you want a rule to be set to Drop and Generate events and the rule was disabled or set to Generate Events in the base policy, you must manually reset the rule state.

When you generate recommendations without changing the advanced settings for Firepower recommended rules, the system recommends rule state changes for all hosts in your entire discovered network.

By default, the system generates recommendations only for rules with low or medium overhead, and generates recommendations to disable rules.

The system does not recommend a rule state for an intrusion rule that is based on a vulnerability that you disable using the Impact Qualification feature.

The system always recommends that you enable a local rule associated with a third-party vulnerability mapped to a host.

The system does not make state recommendations for unmapped local rules.

Related Topics

Deactivating Individual Vulnerabilities, on page 2348
Third-Party Product Mappings, on page 1846

Advanced Settings for Firepower Recommendations

Include all differences between recommendations and rule states in policy reports

By default, an intrusion policy report lists the policy's enabled rules, that is, rules set to either Generate Events or Drop and Generate Events. Enabling the **Include all differences** option also lists the rules whose recommended states differ from their saved states. For information on policy reports, see Policy Reports, on page 356.

Networks to Examine

Specifies the monitored networks or individual hosts to examine for recommendations. You can specify a single IP address or address block, or a comma-separated list comprised of either or both.

Lists of addresses within the hosts that you specify are linked with an OR operation except for negations, which are linked with an AND operation after all OR operations are calculated.

If you want to dynamically adapt active rule processing for specific packets based on host information, you can also enable adaptive profile updates.

Recommendation Threshold (By Rule Overhead)

Prevents the system from recommending or automatically enabling intrusion rules with a higher overhead than the threshold you choose.

Overhead is based on the rule's potential impact on system performance and the likelihood that the rule may generate false positives. Permitting rules with higher overhead usually results in more recommendations, but can affect system performance. You can view the overhead rating for a rule in the rule detail view on the intrusion Rules page.

Note that the system does not factor rule overhead into recommendations to disable rules. Also, local rules are considered to have no overhead, unless they are mapped to a third-party vulnerability.

Generating recommendations for rules with the overhead rating at a particular setting does not preclude you from generating recommendations with different overhead, then generating recommendations again for the original overhead setting. You get the same rule state recommendations for each overhead setting each time you generate recommendations for the same rule set, regardless of the number of times you generate recommendations or how many different overhead settings you generate with. For example, you can generate recommendations with overhead set to medium, then to high, then finally to medium again; if the hosts and applications on your network have not changed, both sets of recommendations with overhead set to medium are then the same for that rule set.

Accept Recommendations to Disable Rules

Specifies whether the system disables intrusion rules based on Firepower recommendations.

Accepting recommendations to disable rules restricts your rule coverage. Omitting recommendations to disable rules augments your rule coverage.

Related Topics

Firepower System IP Address Conventions, on page 15
Adaptive Profile Updates and Firepower Recommended Rules, on page 1813

Generating and Applying Firepower Recommendations

Starting or stopping use of Firepower recommendations may take several minutes, depending on the size of your network and intrusion rule set.

The system builds a separate network map for each leaf domain. In a multidomain deployment, if you enable this feature in an intrusion policy in an ancestor domain, the system generates recommendations using data from all descendant leaf domains. This can enable intrusion rules tailored to assets that may not exist in all leaf domains, which can affect performance.

Before you begin

- Firepower recommendations have the following requirements:
  - FTD License—Threat
  - Classic License—Protection
  - User Roles—Admin or Intrusion Admin
- Configure a network discovery policy before you begin with the steps. Configure the network discovery policy to define internal hosts so that the Firepower recommendations are suitable. See, Network Discovery Customization, on page 1918.

Procedure

Step 1 In the intrusion policy editor's navigation pane, click **Firepower Recommendations**.

Step 2 (Optional) Configure advanced settings; see Advanced Settings for Firepower Recommendations, on page 1519.

Step 3 Generate and apply recommendations.

- **Generate and Use Recommendations**—Generates recommendations and changes rule states to match. Only available if you have never generated recommendations.
- **Generate Recommendations**—Regardless of whether you are using recommendations, generates new recommendations but does not change rule states to match.
- **Update Recommendations**—If you are using recommendations, generates recommendations and changes rule states to match. Otherwise, generates new recommendations without changing rule states.
- **Use Recommendations**—Changes rule states to match any unimplemented recommendations.
- **Do Not Use Recommendations**—Stops use of recommendations. If you manually changed a rule's state before you applied recommendations, the rule state returns to the value you gave it. Otherwise, the rule state returns to its default value.

When you generate recommendations, the system displays a summary of the recommended changes. To view a list of rules where the system recommends a state change, click **View** next to the newly proposed rule state.

Step 4 Evaluate and adjust the recommendations you implemented.

Even if you accept most Firepower recommendations, you can override individual recommendations by setting rule states manually; see Setting Intrusion Rule States, on page 1504.

Step 5 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

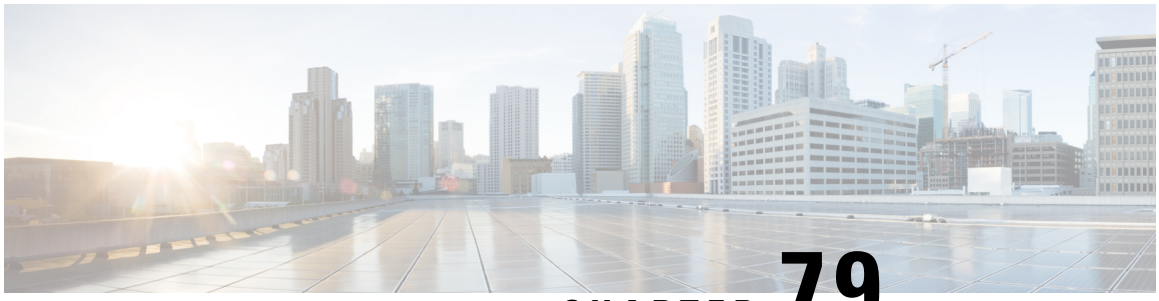
If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Automating Firepower Recommendations, on page 187



CHAPTER 79

Sensitive Data Detection

The following topics explain sensitive data detection and how to configure it:

- Sensitive Data Detection Basics, on page 1523
- Global Sensitive Data Detection Options, on page 1524
- Individual Sensitive Data Type Options, on page 1525
- System-Provided Sensitive Data Types, on page 1526
- License Requirements for Sensitive Data Detection, on page 1527
- Requirements and Prerequisites for Sensitive Data Detection, on page 1527
- Configuring Sensitive Data Detection, on page 1527
- Monitored Application Protocols and Sensitive Data, on page 1529
- Selecting Application Protocols to Monitor, on page 1529
- Special Case: Sensitive Data Detection in FTP Traffic, on page 1530
- Custom Sensitive Data Types, on page 1531

Sensitive Data Detection Basics

Sensitive data such as Social Security numbers, credit card numbers, driver's license numbers, and so on may be leaked onto the Internet, intentionally or accidentally. The system provides a sensitive data preprocessor that can detect and generate events on sensitive data in ASCII text, which can be particularly useful in detecting accidental data leaks.

Global sensitive data preprocessor options control how the preprocessor functions. You can modify global options that specify the following:

- whether the preprocessor replaces all but the last four credit card or Social Security numbers in triggering packets
- which destination hosts on your network to monitor for sensitive data
- how many total occurrences of all data types in a single session result in an event

Individual data types identify the sensitive data you can detect and generate events on in your specified destination network traffic. You can modify default settings for data type options that specify the following:

- a threshold that must be met for a detected data type to generate a single per-session event
- the destination ports to monitor for each data type

- the application protocols to monitor for each data type

You can create and modify custom data types to detect data patterns that you specify. For example, a hospital might create a data type to protect patient numbers, or a university might create a data type to detect student numbers that have a unique numbering pattern.

The system detects sensitive data per TCP session by matching individual data types against traffic. You can modify the default settings for each data type and for global options that apply to all data types in your intrusion policy. The Firepower System provides predefined, commonly used data types. You can also create custom data types.

A sensitive data preprocessor rule is associated with each data type. You enable sensitive data detection and event generation for each data type by enabling the corresponding preprocessor rule for the data type. A link on the configuration page takes you to a filtered view of sensitive data rules on the Rules page, where you can enable and disable rules and configure other rule attributes.

When you save changes to your intrusion policy, you are given the option to automatically enable the sensitive data preprocessor if the rule associated with a data type is enabled and sensitive data detection is disabled.



Tip The sensitive data preprocessor can detect sensitive data in unencrypted Microsoft Word files that are uploaded and downloaded using FTP or HTTP; this is possible because of the way Word files group ASCII text and formatting commands separately.

The system does not detect encrypted or obfuscated sensitive data, or sensitive data in a compressed or encoded format such as a Base64-encoded email attachment. For example, the system would detect the phone number (555)123-4567, but not an obfuscated version where each number is separated by spaces, as in (5 5 5) 1 2 3 - 4 5 6 7, or by intervening HTML code, such as (555)-123-4567. However, the system would detect, for example, the HTML coded number (555)-123-4567 where no intervening codes interrupt the numbering pattern.

Global Sensitive Data Detection Options

Global sensitive data options are policy-specific and apply to all data types.

Mask

Replaces with Xs all but the last four digits of credit card numbers and Social Security numbers in the triggering packet. The masked numbers appear in the intrusion event packet view in the web interface and in downloaded packets.

Networks

Specifies the destination host or hosts to monitor for sensitive data. You can specify a single IP address, address block, or a comma-separated list of either or both. The system interprets a blank field as `any`, meaning any destination IP address.

The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Global Threshold

Specifies the total number of all occurrences of all data types during a single session that the preprocessor must detect in any combination before generating a global threshold event. You can specify 1 through 65535.

Cisco recommends that you set the value for this option higher than the highest threshold value for any individual data type that you enable in your policy.

Note the following points regarding global thresholds:

- You must enable preprocessor rule 139:1 to detect and generate events and, in an inline deployment, drop offending packets on combined data type occurrences.
- The preprocessor generates up to one global threshold event per session.
- Global threshold events are independent of individual data type events; that is, the preprocessor generates an event when the global threshold is reached, regardless of whether the event threshold for any individual data type has been reached, and vice versa.

Related Topics

Firepower System IP Address Conventions, on page 15

Individual Sensitive Data Type Options

At a minimum, each custom data type must specify an event threshold and at least one port or application protocol to monitor.

Each system-provided data type uses an otherwise inaccessible `sd_pattern` keyword to define a built-in data pattern to detect in traffic. You can also create custom data types for which you use simple regular expressions to specify your own data patterns.

Sensitive data types display in all intrusion policies where Sensitive Data Detection is enabled. System-provided data types display as read-only. For custom data types, the name and pattern fields display as read-only, but you can set the other options to policy-specific values.

In a multidomain deployment, the system displays sensitive data types created in the current domain, which you can edit. It also displays data types created in ancestor domains, which you can edit in a limited way. For ancestor data types, the name and pattern fields display as read-only, but you can set the other options to policy-specific values.

Table 113: Individual Data Type Options

Option	Description
Data Type	Specifies the unique name for the data type.
Threshold	Specifies the number of occurrences of the data type when the system generates an event. You can specify 1 through 255. Note that the preprocessor generates one event for a detected data type per session. Note also that global threshold events are independent of individual data type events; that is, the preprocessor generates an event when the data type event threshold is reached, regardless of whether the global event threshold has been reached, and vice versa.

Option	Description
Destination Ports	Specifies destination ports to monitor for the data type. You can specify a single port, a comma-separated list of ports, or `any`, meaning any destination port.
Application Protocols	Specifies up to eight application protocols to monitor for the data type. You must activate application detectors to identify application protocols to monitor. Note that, for Classic devices, this feature requires a Control license.
Pattern	Specifies the pattern to detect. This field is only present for custom data types.

Related Topics

Activating and Deactivating Detectors, on page 1889

System-Provided Sensitive Data Types

Each intrusion policy includes system-provided data types for detecting commonly used data patterns such as credit card numbers, email addresses, U.S. phone numbers, and U.S. Social Security numbers with and without dashes.

Each system-provided data type is associated with a single sensitive data preprocessor rule that has a generator ID (GID) of 138. You must enable the associated sensitive data rule in the intrusion policy to generate events and, in an inline deployment, drop offending packets for each data type that you want to use in your policy.

The following table describes each data type and lists the corresponding preprocessor rule.

Table 114: System-Provided Sensitive Data Types

Data Type	Description	Preprocessor GID:SID
Credit Card Numbers	Matches Visa®, MasterCard®, Discover® and American Express® fifteen- and sixteen-digit credit card numbers, with or without their normal separating dashes or spaces; also uses the Luhn algorithm to verify credit card check digits.	138:2
Email Addresses	Matches email addresses.	138:5
U.S. Phone Numbers	Matches U.S. phone numbers adhering to the pattern `(\d{3}) ?\d{3}-\d{4}`.	138:6
U.S. Social Security Numbers Without Dashes	Matches 9-digit U.S. Social Security numbers that have valid 3-digit area numbers, valid 2-digit group numbers, and do not have dashes.	138:4
U.S. Social Security Numbers With Dashes	Matches 9-digit U.S. Social Security numbers that have valid 3-digit area numbers, valid 2-digit group numbers, and dashes.	138:3

To reduce false positives from 9-digit numbers other than Social Security numbers, the preprocessor uses an algorithm to validate the 3-digit area number and 2-digit group number that precede the 4-digit serial number in each Social Security number. The preprocessor validates Social Security group numbers through November 2009.

License Requirements for Sensitive Data Detection

FTD License

Threat

Classic License

Protection, or as indicated in a procedure.

Requirements and Prerequisites for Sensitive Data Detection

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Intrusion Admin

Configuring Sensitive Data Detection

Because sensitive data detection can have a high impact on the performance of your Firepower System, Cisco recommends that you adhere to the following guidelines:

- Choose the No Rules Active default policy as your base intrusion policy.
- Ensure that the following settings are enabled in the corresponding network analysis policy:
  - **FTP and Telnet Configuration** under **Application Layer Preprocessors**
  - **IP Defragmentation** and **TCP Stream Configuration** under **Transport/Network Layer Preprocessors**.

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Before you begin

For classic devices, this procedure requires the Protection or Control license.

Procedure

Step 1 Choose **Policies** > **Access Control** > **Intrusion**

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Advanced Settings** in the navigation panel.

Step 4 If **Sensitive Data Detection** under **Specific Threat Detection** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **Sensitive Data Detection**.

Step 6 You have the following choices:

- Modify the global settings as described in Global Sensitive Data Detection Options, on page 1524.
- Choose a data type in the **Targets** section, and modify the data type configuration as described in Individual Sensitive Data Type Options, on page 1525.
- If you want to inspect custom sensitive data, create a custom data type; see Custom Sensitive Data Types, on page 1531.

Step 7 Add or remove application protocols to monitor for a data type; see Monitored Application Protocols and Sensitive Data, on page 1529.

Note To detect sensitive data in FTP traffic, you must add the `Ftp data` application protocol.

Step 8 Optionally, to display sensitive data preprocessor rules, click **Configure Rules for Sensitive Data Detection**.

You can enable or disable any of the listed rules. You can also configure sensitive data rules for any of the other actions available on the Rules page, such as rule suppression, rate-based attack prevention, and so on; see Intrusion Rule Types, on page 1489 for more information.

Step 9 To save changes you made in this policy since the last policy commit, click **Policy Information** in the navigation panel, then click **Commit Changes**.

If you enable sensitive data preprocessor rules in your policy without enabling sensitive data detection, you are prompted to enable sensitive data detection when you save changes to your policy.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want to generate intrusion events, enable Sensitive Data Detection rules 138:2, 138:3, 138:4, 138:5, 138:6, 138:>999999, or 139:1. For more information, see Intrusion Rule States, on page 1503, Global Sensitive Data Detection Options, on page 1524, System-Provided Sensitive Data Types, on page 1526, and Custom Sensitive Data Types, on page 1531.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Special Case: Sensitive Data Detection in FTP Traffic, on page 1530

Monitored Application Protocols and Sensitive Data

You can specify up to eight application protocols to monitor for each data type. At least one detector must be enabled for each application protocol you select. By default, all system-provided detectors are activated. If no detector is enabled for an application protocol, the system automatically enables all system-provided detectors for the application; if none exist, the system enables the most recently modified user-defined detector for the application.

You must specify at least one application protocol or port to monitor for each data type. However, except in the case where you want to detect sensitive data in FTP traffic, Cisco recommends for the most complete coverage that you specify corresponding ports when you specify application protocols. For example, if you specify HTTP, you might also configure the well-known HTTP port 80. If a new host on your network implements HTTP, the system monitors port 80 during the interval when it is discovering the new HTTP application protocol.

In the case where you want to detect sensitive data in FTP traffic, you must specify the `FTP data` application protocol; there is no advantage in specifying a port number.

Related Topics

Activating and Deactivating Detectors, on page 1889
Special Case: Sensitive Data Detection in FTP Traffic, on page 1530

Selecting Application Protocols to Monitor

You can specify application protocols to monitor in both system-provided and custom sensitive data types. The application protocols you select are policy-specific.

Before you begin

For classic devices, this procedure requires the Control license.

Procedure

Step 1 Choose **Policies** > **Access Control** > **Intrusion**.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Advanced Settings** in the navigation panel.

Step 4 If **Sensitive Data Detection** under **Specific Threat Detection** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **Sensitive Data Detection**.

Step 6 Click the name of a data type under **Data Types**.

Step 7 Click **Edit** () next to the **Application Protocols** field.

Step 8 You have the following choices:

- To add application protocols for monitoring, choose one or more application protocols from the **Available** list, then click right arrow (>). You can add up to eight application protocols for monitoring.
- To remove an application protocol from monitoring, choose it from the **Enabled** list, then click left arrow (<).

Step 9 Click **OK**.

Step 10 To save changes you made in this policy since the last policy commit, click **Policy Information** in the navigation pane, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Special Case: Sensitive Data Detection in FTP Traffic, on page 1530

Special Case: Sensitive Data Detection in FTP Traffic

You usually determine which traffic to monitor for sensitive data by specifying the ports to monitor or specifying application protocols in deployments.

However, specifying ports or application protocols is not sufficient for detecting sensitive data in FTP traffic. Sensitive data in FTP traffic is found in traffic for the FTP application protocol, which occurs intermittently and uses a transient port number, making it difficult to detect. To detect sensitive data in FTP traffic, you **must** include the following in your configuration:

- Specify the `FTP data` application protocol to enable detection of sensitive data in FTP traffic.

  In the special case of detecting sensitive data in FTP traffic, specifying the `FTP data` application protocol does not invoke detection; instead, it invokes the rapid processing of the FTP/Telnet processor to detect sensitive data in FTP traffic.

- Ensure that the FTP Data detector, which is enabled by default, is enabled.
- Ensure that your configuration includes at least one port to monitor for sensitive data.
- Ensure that the file policy is enabled for the Access Control Policy.

Note that it is not necessary to specify an FTP port except in the unlikely case where you only want to detect sensitive data in FTP traffic. Most sensitive data configurations will include other ports such as HTTP or email ports. In the case where you do want to specify only one FTP port and no other ports to monitor, Cisco recommends that you specify the FTP command port `23`.

Related Topics

The FTP/Telnet Decoder, on page 1700
Activating and Deactivating Detectors, on page 1889
Configuring Sensitive Data Detection, on page 1527

Custom Sensitive Data Types

Each custom data type you create also creates a single sensitive data preprocessor rule that has a Generator ID (GID) of 138 and a Snort ID (SID) of 1000000 or greater, that is, a SID for a local rule. In a multidomain deployment, the system prepends a domain number to the SID of any custom rule created in or imported into a descendant domain. For example, a rule added in the Global domain would have a SID of 1000000 or greater, and rules added in descendant domains would have SIDs of [domain number]000000 or greater.

You must enable the associated sensitive data rule to enable detection, generate events and, in an inline deployment, drop offending packets for each custom data type that you want to use in your policy.

To help you enable sensitive data rules, a link on the configuration page takes you to a filtered view of the intrusion policy Rules page that displays all system-provided and custom sensitive data rules. You can also display custom sensitive data rules along with any custom local rules by choosing the local filtering category on the intrusion policy Rules page. Note that custom sensitive data rules are not listed on the intrusion rules editor page (**Objects** > **Intrusion Rules**).

Once you create a custom data type, you can enable it in any intrusion policy in the system or, for multidomain deployments, in the current domain. To enable a custom data type, you must enable the associated sensitive data rule in any policy that you want to use to detect that custom data type.

Data Patterns in Custom Sensitive Data Types

You define the data pattern for a custom data type using a simple set of regular expressions comprised of the following:

- three metacharacters
- escaped characters that allow you to use the metacharacters as literal characters
- six character classes

Metacharacters are literal characters that have special meaning within regular expressions.

Table 115: Sensitive Data Pattern Metacharacters

Metacharacter	Description	Example
`?`	Matches zero or one occurrence of the preceding character or escape sequence; that is, the preceding character or escape sequence is optional.	`colou?r` matches `color` or `colour`
`{n}`	Matches the preceding character or escape sequence n times.	For example, `\d{2}` matches `55`, `12`, and so on; `\l{3}` matches `AbC`, `www`, and so on; `\w{3}` matches `a1B`, 25C, and so on; `x{5}` matches `xxxxx`

Metacharacter	Description	Example
\	Allows you to use metacharacters as actual characters and is also used to specify a predefined character class.	`\?` matches a question mark, `\\` matches a backslash, `\d` matches numeric characters, and so on

You must use a backslash to escape certain characters for the sensitive data preprocessor to interpret them correctly as literal characters.

Table 116: Escaped Sensitive Data Pattern Characters

Use this escaped character...	To represent this literal character...
\?	?
\{	{
\}	}
\\	\

When defining a custom sensitive data pattern, you can use character classes.

Table 117: Sensitive Data Pattern Character Classes

Character Class	Description	Character Class Definition
\d	Matches any numeric ASCII character 0-9	0-9
\D	Matches any byte that is not a numeric ASCII character	not 0-9
\l (lowercase “ell”)	Matches any ASCII letter	a-zA-Z
\L	Matches any byte that is not an ASCII letter	not a-zA-Z
\w	Matches any ASCII alphanumeric character Note that, unlike PCRE regular expressions, this does not include an underscore (_).	a-zA-Z0-9
\W	Matches any byte that is not an ASCII alphanumeric character	not a-zA-Z0-9

The preprocessor treats characters entered directly, instead of as part of a regular expression, as literal characters. For example, the data pattern 1234 matches `1234`.

The following data pattern example, which is used in system-provided sensitive data rule 138:4, uses the escaped digits character class, the multiplier and option-specifier metacharacters, and the literal dash (-) and left and right parentheses () characters to detect U.S. phone numbers:

```
(\d{3}) ?\d{3}-\d{4}
```

Exercise caution when creating custom data patterns. Consider the following alternative data pattern for detecting phone numbers which, although using valid syntax, could cause many false positives:

```
(?\d{3})? ?\d{3}-?\d{4}
```

Because the second example combines optional parentheses, optional spaces, and optional dashes, it would detect, among others, phone numbers in the following desirable patterns:

- `(555)123-4567`
- `555123-4567`
- `5551234567`

However, the second example pattern would also detect, among others, the following potentially invalid patterns, resulting in false positives:

- `(555 1234567`
- `555)123-4567`
- `555) 123-4567`

Consider finally, for illustration purposes only, an extreme example in which you create a data pattern that detects the lowercase letter `a` using a low event threshold in all destination traffic on a small company network. Such a data pattern could overwhelm your system with literally millions of events in only a few minutes.

Configuring Custom Sensitive Data Types

In a multidomain deployment, the system displays sensitive data types created in the current domain, which you can edit. It also displays data types created in ancestor domains, which you can edit in a limited way. For ancestor data types, the name and pattern fields display as read-only, but you can set the other options to policy-specific values.

You cannot delete a data type if the sensitive data rule for that data type is enabled in any intrusion policy.

Procedure

Step 1 Choose **Policies** > **Access Control** > **Intrusion**

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Advanced Settings** in the navigation panel.

Step 4 If **Sensitive Data Detection** under **Specific Threat Detection** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **Sensitive Data Detection**.

Step 6 Click **Add** ( ) next to **Data Types**.

Step 7 Enter a name for the data type.

Step 8 Enter the pattern you want to detect with this data type; see Data Patterns in Custom Sensitive Data Types, on page 1531.

Step 9 Click **OK**.

Step 10 Optionally, click the data type name, and modify the options described in Individual Sensitive Data Type Options, on page 1525.

Step 11 Optionally, delete a custom data type by clicking **Delete** (), then **OK** to confirm.

Note If the sensitive data rule for that data type is enabled in any intrusion policy, the system warns that you cannot delete the data type. You must disable the sensitive data rule in affected policies before attempting the deletion again; see Setting Intrusion Rule States, on page 1504.

Step 12 To save changes you made in this policy since the last policy commit, click **Policy Information** in the navigation panel, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Enable the associated custom sensitive data preprocessing rule in each policy where you want to use that data type; see Setting Intrusion Rule States, on page 1504.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Editing Custom Sensitive Data Types, on page 1534

Editing Custom Sensitive Data Types

You can edit all fields in custom sensitive data types. Note, however, that when you modify the name or pattern field, these settings change in all intrusion policies on the system. You can set the other options to policy-specific values.

In a multidomain deployment, the system displays sensitive data types created in the current domain, which you can edit. It also displays data types created in ancestor domains, which you can edit in a limited way. For ancestor data types, the name and pattern fields display as read-only, but you can set the other options to policy-specific values.

Procedure

Step 1 Choose **Policies** > **Access Control** > **Intrusion**

Step 2 Click **Edit** () next to the policy you want to edit.

If **View** () appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Advanced Settings** in the navigation panel.

Step 4 If **Sensitive Data Detection** under **Specific Threat Detection** is disabled, click **Enabled**.

Step 5 Click **Edit** next to **Sensitive Data Detection**.

Step 6 In the **Targets** section, click the name of the custom data type.

Step 7 Click **Edit Data Type Name and Pattern**.

Step 8 Modify the data type name and pattern; see Data Patterns in Custom Sensitive Data Types, on page 1531.

Step 9 Click **OK**.

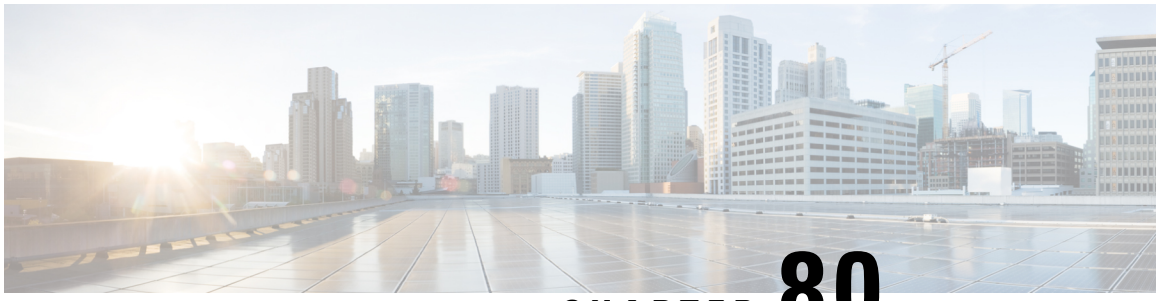
Step 10 Set the remaining options to policy-specific values; see Individual Sensitive Data Type Options, on page 1525.

Step 11 To save changes you made in this policy since the last policy commit, click **Policy Information** in the navigation panel, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.



CHAPTER 80

Globally Limiting Intrusion Event Logging

The following topics describe how to globally limit intrusion event logging:

- Global Rule Thresholding Basics, on page 1537
- Global Rule Thresholding Options, on page 1538
- License Requirements for Global Thresholds, on page 1539
- Requirements and Prerequisites for Global Thresholds, on page 1540
- Configuring Global Thresholds, on page 1540
- Disabling the Global Threshold, on page 1541

Global Rule Thresholding Basics

The global rule threshold sets limits for event logging by an intrusion policy. You can set a global rule threshold across all traffic to limit how often the policy logs events from a specific source or destination and displays those events per specified time period. You can also set thresholds per shared object rule, standard text rule, or preprocessor rule in the policy. When you set a global threshold, that threshold applies for each rule in the policy that does not have an overriding specific threshold. Thresholds can prevent you from being overwhelmed with a large number of events.

Every intrusion policy contains a default global rule threshold that applies by default to all intrusion rules and preprocessor rules. This default threshold limits the number of events on traffic going to a destination to one event per 60 seconds.

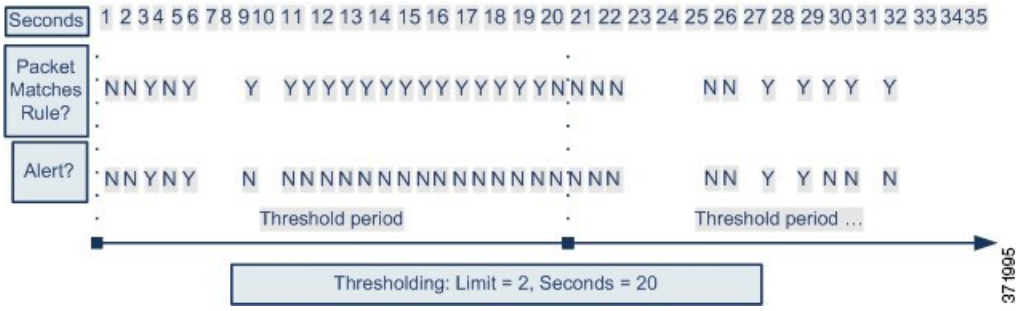
You can:

- Change the global threshold.
- Disable the global threshold.
- Override the global threshold by setting individual thresholds for specific rules.

  For example, you might set a global limit threshold of five events every 60 seconds, but then set a specific threshold of ten events for every 60 seconds for SID 1315. All other rules generate no more than five events in each 60-second period, but the system generates up to ten events for each 60-second period for SID 1315.

Tip A global or individual threshold on a managed device with multiple CPUs may result in a higher number of events than expected.

The following diagram demonstrates how the global rule thresholding works. In this example, an attack is in progress for a specific rule. The global limit threshold is set to limit event generation for each rule to two events every 20 seconds. Note that the period starts at one second and ends at 21 seconds. After the period ends, the cycle starts again and the next two rule matches generate events, then the system does not generate any more events during that period.



Global Rule Thresholding Options

The default threshold limits event generation for each rule to one event every 60 seconds on traffic going to the same destination. The default values for the global rule thresholding options are:

- **Type** — Limit
- **Track By** — Destination
- **Count** — 1
- **Seconds** — 60

You can modify these default values as follows:

Table 118: Thresholding Types

Option	Description
Limit	Logs and displays events for the specified number of packets (specified by the count argument) that trigger the rule during the specified time period. For example, if you set the type to **Limit**, the **Count** to `10`, and the **Seconds** to `60`, and 14 packets trigger the rule, the system stops logging events for the rule after displaying the first 10 that occur within the same minute.
Threshold	Logs and displays a single event when the specified number of packets (specified by the count argument) trigger the rule during the specified time period. Note that the counter for the time restarts after you hit the threshold count of events and the system logs that event. For example, you set the type to **Threshold**, **Count** to `10`, and **Seconds** to `60`, and the rule triggers 10 times by second 33. The system generates one event, then resets the Seconds and Count counters to `0`. The rule then triggers another 10 times in the next 25 seconds. Because the counters reset to `0` at second 33, the system logs another event.

Option	Description
Both	Logs and displays an event once per specified time period, after the specified number (count) of packets trigger the rule. For example, if you set the type to **Both**, **Count** to `2`, and **Seconds** to `10`, the following event counts result: • If the rule is triggered once in 10 seconds, the system does not generate any events (the threshold is not met) • If the rule is triggered twice in 10 seconds, the system generates one event (the threshold is met when the rule triggers the second time) • If the rule is triggered four times in 10 seconds, the system generates one event (the threshold is met when the rule triggered the second time and following events are ignored)

The **Track By** option determines whether the event instance count is calculated per source or destination IP address.

You can also specify the number of instances and time period that define the threshold, as follows:

Table 119: Thresholding Instance/Time Options

Option	Description
Count	For a **Limit** threshold, the number of event instances per specified time period per tracking IP address or address range required to meet the threshold. For a **Threshold** threshold, the number of rule matches you want to use as your threshold.
Seconds	For a **Limit** threshold, the number of seconds that make up the time period when attacks are tracked. For a **Threshold** threshold, the number of seconds that elapse before the count resets. If you set the threshold type to **Limit**, the tracking to **Source**, **Count** to 10, and **Seconds** to 10, the system logs and displays the first 10 events that occur in 10 seconds from a given source port. If only seven events occur in the first 10 seconds, the system logs and displays those, if 40 events occur in the first 10 seconds, the system logs and displays 10, then begins counting again when the 10-second time period elapses.

Related Topics

Configuring Global Thresholds, on page 1540
Intrusion Event Thresholds, on page 1505

License Requirements for Global Thresholds

FTD License

Threat

Classic License

Protection

Requirements and Prerequisites for Global Thresholds

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Intrusion Admin

Configuring Global Thresholds

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Policies** > **Access Control** > **Intrusion**.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Advanced Settings** in the navigation panel.

Step 4 If **Global Rule Thresholding** under **Intrusion Rule Thresholds** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **Global Rule Thresholding**.

Step 6 Using **Type**, specify the type of threshold that will apply over the time you specify in the **Seconds** field.

Step 7 Using **Track By**, specify the tracking method.

Step 8 Enter a value in the **Count** field.

Step 9 Enter a value in the **Seconds** field.

Step 10 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Global Rule Thresholding Options, on page 1538
Configuring Intrusion Rules in Layers, on page 1473
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

Disabling the Global Threshold

You can disable global thresholding in the highest policy layer if you want to threshold events for specific rules rather than applying thresholding to every rule by default.

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Policies** > **Access Control** > **Intrusion**

Step 2 Click **Edit** () next to the policy you want to edit.

If **View** () appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Advanced Settings** in the navigation panel.

Step 4 Next to **Global Rule Thresholding** under **Intrusion Rule Thresholds**, click **Disabled**.

Step 5 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

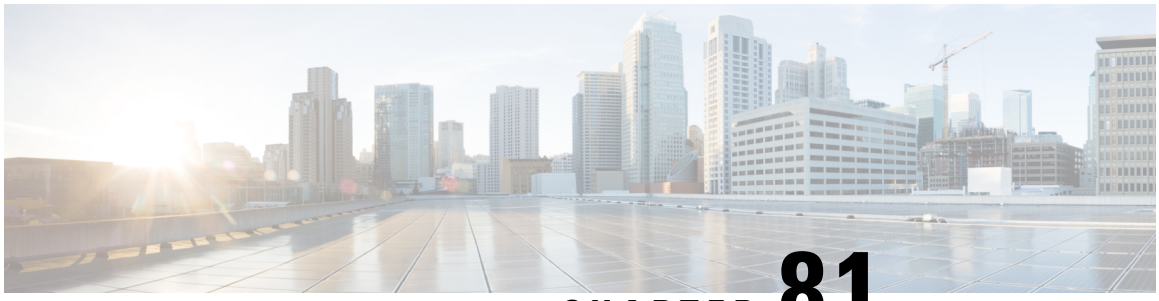
If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461
Configuring Intrusion Rules in Layers, on page 1473



CHAPTER 81

The Intrusion Rules Editor

The following topics describe how to use the intrusion rules editor:

- An Introduction to Intrusion Rule Editing, on page 1543
- License Requirements for the Intrusion Rule Editor, on page 1544
- Requirements and Prerequisites for the Intrusion Rule Editor, on page 1544
- Rule Anatomy, on page 1544
- Custom Rule Creation, on page 1556
- Searching for Rules, on page 1561
- Rule Filtering on the Intrusion Rules Editor Page, on page 1562
- Keywords and Arguments in Intrusion Rules, on page 1565

An Introduction to Intrusion Rule Editing

An *intrusion rule* is a set of keywords and arguments that the system uses to detect attempts to exploit vulnerabilities on your network. As the system analyzes network traffic, it compares packets against the conditions specified in each rule. If the packet data matches all the conditions specified in a rule, the rule triggers. If a rule is an *alert rule*, it generates an intrusion event. If it is a *pass rule*, it ignores the traffic. For a *drop* rule in an inline deployment, the system drops the packet and generates an event. You can view and evaluate intrusion events from the Firepower Management Center web interface.

The Firepower System provides two types of intrusion rules: shared object rules and standard text rules. The Cisco Talos Intelligence Group (Talos) can use shared object rules to detect attacks against vulnerabilities in ways that traditional standard text rules cannot. You cannot create shared object rules. When you write your own intrusion rule, you create a standard text rule.

You can write custom standard text rules to tune the types of events you are likely to see. Note that while this documentation sometimes discusses rules targeted to detect specific exploits, the most successful rules target traffic that may attempt to exploit known vulnerabilities rather than specific known exploits. By writing rules and specifying the rule's event message, you can more easily identify traffic that indicates attacks and policy evasions.

When you enable a custom standard text rule in a custom intrusion policy, keep in mind that some rule keywords and arguments require that traffic first be decoded or preprocessed in a certain way. This chapter explains the options you must configure in your network analysis policy, which governs preprocessing. Note that if you disable a required preprocessor, the system automatically uses it with its current settings, although the preprocessor remains disabled in the network analysis policy web interface.

Caution Make sure you use a controlled network environment to test any intrusion rules that you write before you use the rules in a production environment. Poorly written intrusion rules may seriously affect the performance of the system.

In a multidomain deployment, the system displays rules created in the current domain, which you can edit. It also displays rules created in ancestor domains, which you cannot edit. To view and edit rules created in a lower domain, switch to that domain. The system-provided intrusion rules belong to the Global domain. Administrators in descendant domains can make local editable copies of these system rules.

License Requirements for the Intrusion Rule Editor

FTD License

Threat

Classic License

Protection

Requirements and Prerequisites for the Intrusion Rule Editor

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Intrusion Admin

Rule Anatomy

All standard text rules contain two logical sections: the rule header and the rule options. The rule header contains:

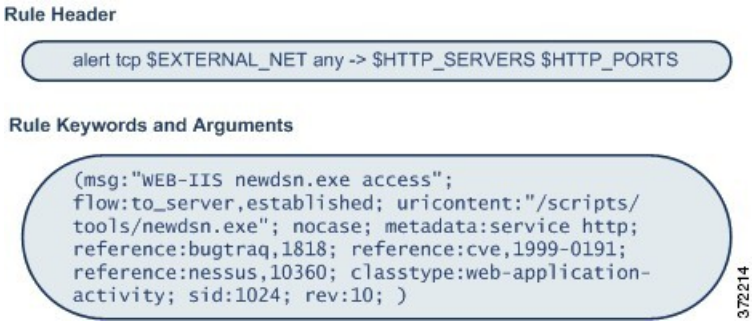
- the rule's action or type
- the protocol
- the source and destination IP addresses and netmasks
- direction indicators showing the flow of traffic from source to destination

- the source and destination ports

The rule options section contains:

- event messages
- keywords and their parameters and arguments
- patterns that a packet's payload must match to trigger the rule
- specifications of which parts of the packet the rules engine should inspect

The following diagram illustrates the parts of a rule:



Note that the options section of a rule is the section enclosed in parentheses. The intrusion rules editor provides an easy-to-use interface to help you build standard text rules.

The Intrusion Rule Header

Every standard text rule and shared object rule has a rule header containing parameters and arguments. The following illustrates parts of a rule header:



The following table describes each part of the rule header shown above.

Table 120: Rule Header Values

Rule Header Component	Example Value	This Value...
Action	`alert`	Generates an intrusion event when triggered.
Protocol	`tcp`	Tests TCP traffic only.
Source IP Address	`$EXTERNAL_NET`	Tests traffic coming from any host that is not on your internal network.

Rule Header Component	Example Value	This Value...
Source Ports	`any`	Tests traffic coming from any port on the originating host.
Operator	`->`	Tests external traffic (destined for the web servers on your network).
Destination IP Address	`$HTTP_SERVERS`	Tests traffic to be delivered to any host specified as a web server on your internal network.
Destination Ports	`$HTTP_PORTS`	Tests traffic delivered to an HTTP port on your internal network.

Note The previous example uses default variables, as do most intrusion rules.

Related Topics

Variable Sets, on page 412

Intrusion Rule Header Action

Each rule header includes a parameter that specifies the action the system takes when a packet triggers a rule. Rules with the action set to *alert* generate an intrusion event against the packet that triggered the rule and log the details of that packet. Rules with the action set to *pass* do not generate an event against, or log the details of, the packet that triggered the rule.

Note In an inline deployment, rules with the rule state set to *Drop and Generate Events* generate an intrusion event against the packet that triggered the rule. Also, if you apply a drop rule in a passive deployment, the rule acts as an alert rule.

By default, pass rules override alert rules. You can create pass rules to prevent packets that meet criteria defined in the pass rule from triggering the alert rule in specific situations, rather than disabling the alert rule. For example, you might want a rule that looks for attempts to log into an FTP server as the user “anonymous” to remain active. However, if your network has one or more legitimate anonymous FTP servers, you could write and activate a pass rule that specifies that, for those specific servers, anonymous users do not trigger the original rule.

Within the intrusion rules editor, you select the rule type from the **Action** list.

Intrusion Rule Header Protocol

In each rule header, you must specify the protocol of the traffic the rule inspects. You can specify the following network protocols for analysis:

- ICMP (Internet Control Message Protocol)
- IP (Internet Protocol)

Note The system ignores port definitions in an intrusion rule header when the protocol is set to `ip`.

- TCP (Transmission Control Protocol)
- UDP (User Datagram Protocol)

Use **IP** as the protocol type to examine all protocols assigned by IANA, including TCP, UDP, ICMP, IGMP, and many more.

Note You cannot currently write rules that match patterns in the next header (for example, the TCP header) in an IP payload. Instead, content matches begin with the last decoded protocol. As a workaround, you can match patterns in TCP headers by using rule options.

Within the Intrusion Rules editor, you select the protocol type from the **Protocol** list.

Related Topics

Intrusion Rule Header Protocol, on page 1546

Intrusion Rule Header Direction

Within the rule header, you can specify the direction that the packet must travel for the rule to inspect it. The following table describes these options.

Table 121: Directional Options in Rule Headers

Use...	To Test...
Directional	only traffic from the specified source IP address to the specified destination IP address
Bidirectional	all traffic traveling between the specified source and destination IP addresses

Intrusion Rule Header Source and Destination IP Addresses

Restricting packet inspection to the packets originating from specific IP addresses or destined to a specific IP address reduces the amount of packet inspection the system must perform. This also reduces false positives by making the rule more specific and removing the possibility of the rule triggering against packets whose source and destination IP addresses do not indicate suspicious behavior.

Tip The system recognizes only IP addresses and does not accept host names for source or destination IP addresses.

Within the intrusion rules editor, you specify source and destination IP addresses in the **Source IPs** and **Destination IPs** fields.

When writing standard text rules, you can specify IPv4 and IPv6 addresses in a variety of ways, depending on your needs. You can specify a single IP address, `any`, IP address lists, CIDR notation, prefix lengths, or a

network variable. Additionally, you can indicate that you want to exclude a specific IP address or set of IP addresses. When specifying IPv6 addresses, you can use any addressing convention defined in RFC 4291.

IP Address Syntax in Intrusion Rules

The following table summarizes the various ways you can specify source and destination IP addresses.

Table 122: Source/Destination IP Address Syntax

To Specify...	Use...	Example
any IP address	`any`	`any`
a specific IP address	the IP address Note that you would not mix IPv4 and IPv6 source and destination addresses in the same rule.	`192.168.1.1` `2001:db8::abcd`
a list of IP addresses	brackets (`[]`) to enclose the IP addresses and commas to separate them	`[192.168.1.1,192.168.1.15]` `[2001:db8::b3ff, 2001:db8::0202]`
a block of IP addresses	IPv4 CIDR block or IPv6 address prefix notation	`192.168.1.0/24` `2001:db8::/32`
anything except a specific IP address or set of addresses	the `!` character before the IP address or addresses you want to negate	`!192.168.1.15` `!2001:db8::0202:b3ff:fe1e`
anything in a block of IP addresses except one or more specific IP addresses	a block of addresses followed by a list of negated addresses or blocks	`[10.0.0/8, !10.2.3.4, !10.1.0.0/16]` `[2001:db8::/32, !2001:db8::8329, !2001:db8::0202]`
IP addresses defined by a network variable	the variable name, in uppercase letters, preceded by `$` Note that preprocessor rules can trigger events regardless of the hosts defined by network variables used in intrusion rules.	`$HOME_NET`
all IP addresses except addresses defined by an IP address variable	the variable name, in uppercase letters, preceded by `!$`	`!$HOME_NET`

The following descritpions provide additional information on some of the IP address entry methods.

Any IP Address

You can specify the word `any` as a rule source or destination IP address to indicate any IPv4 or IPv6 address.

For example, the following rule uses the argument **`any`** in the **Source IPs** and **Destination IPs** fields and evaluates packets with any IPv4 or IPv6 source or destination address:

```
alert tcp any any -> any any
```

You can also specify `::` to indicate any IPv6 address.

Multiple IP Addresses

You can list individual IP addresses by separating the IP addresses with commas and, optionally, by surrounding non-negated lists with brackets, as shown in the following example:

```
[192.168.1.100,192.168.1.103,192.168.1.105]
```

You can list IPv4 and IPv6 addresses alone or in any combination, as shown in the following example:

```
[192.168.1.100,2001:db8::1234,192.168.1.105]
```

Note that surrounding an IP address list with brackets, which was required in earlier software releases, is not required. Note also that, optionally, you can enter lists with a space before or after each comma.



Note You must surround negated lists with brackets.

You can also use IPv4 Classless Inter-Domain Routing (CIDR) notation or IPv6 prefix lengths to specify address blocks. For example:

- 192.168.1.0/24 specifies the IPv4 addresses in the 192.168.1.0 network with a subnet mask of 255.255.255.0, that is, 192.168.1.0 through 192.168.1.255.
- 2001:db8::/32 specifies the IPv6 addresses in the 2001:db8:: network with a prefix length of 32 bits, that is, 2001:db8:: through 2001:db8:ffff:ffff:ffff:ffff:ffff:ffff.



Tip If you need to specify a block of IP addresses but cannot express it using CIDR or prefix length notation alone, you can use CIDR blocks and prefix lengths in an IP address list.

IP Addresses Negation

You can use an exclamation point (`!`) to negate a specified IP address. That is, you can match any IP address with the exception of the specified IP address or addresses. For example, `!192.168.1.1` specifies any IP address other than 192.168.1.1, and `!2001:db8:ca2e::fa4c` specifies any IP address other than 2001:db8:ca2e::fa4c.

To negate a list of IP addresses, place `!` before a bracketed list of IP addresses. For example, `![192.168.1.1,192.168.1.5]` would define any IP address other than 192.168.1.1 or 192.168.1.5.



Note You must use brackets to negate a list of IP addresses.

Be careful when using the negation character with IP address lists. For example, if you use `[!192.168.1.1,!192.168.1.5]` to match any address that is not 192.168.1.1 or 192.168.1.5, the system interprets this syntax as "anything that is not 192.168.1.1, **or** anything that is not 192.168.1.5."

Because 192.168.1.5 is not 192.168.1.1, and 192.168.1.1 is not 192.168.1.5, both IP addresses match the IP address value of `[!192.168.1.1,!192.168.1.5]`, and it is essentially the same as using "`any`."

Instead, use `![192.168.1.1,192.168.1.5]`. The system interprets this as "**not** 192.168.1.1 **and not** 192.168.1.5," which matches any IP address other than those listed between brackets.

Note that you cannot logically use negation with `any` which, if negated, would indicate no address.

Related Topics

Variable Sets, on page 412

Intrusion Rule Header Source and Destination Ports

Within the intrusion rules editor, you specify source and destination ports in the **Source Port** and **Destination Port** fields.

Port Syntax in Intrusion Rules

The Firepower System uses a specific type of syntax to define the port numbers used in rule headers.

Note The system ignores port definitions in an intrusion rule header when the protocol is set to `ip`.

You can list ports by separating the ports with commas, as shown in the following example:

```
80, 8080, 8138, 8600-9000, !8650-8675
```

Optionally, the following example shows how you can surround a port list with brackets, which was required in previous software versions but is no longer required:

```
[80, 8080, 8138, 8600-9000, !8650-8675]
```

Note that you **must** surround negated port lists in brackets, as shown in the following example:

```
![20, 22, 23]
```

The following table summarizes the syntax you can use:

Table 123: Source/Destination Port Syntax

To Specify...	Use	Example
any port	`any`	`any`
a specific port	the port number	`80`
a range of ports	a dash between the first and last port number in the range	`80-443`
all ports less than or equal to a specific port	a dash before the port number	`-21`
all ports greater than or equal to a specific port	a dash after the port number	`80-`
all ports except a specific port or range of ports	the `!` character before the port, port list, or range of ports you want to negate Note that you can logically use negation with all port designations except any, which if negated would indicate *no port*.	`!20`
all ports defined by a port variable	the variable name, in uppercase letter, preceded by `$`	`$HTTP_PORTS`

To Specify...	Use	Example
all ports except ports defined by a port variable	the variable name, in uppercase letter, preceded by `!$`	`!$HTTP_PORTS`

Intrusion Event Details

As you construct a standard text rule, you can include contextual information that describes the vulnerability that the rule detects in exploit attempts. You can also include external references to vulnerability databases and define the priority that the event holds in your organization. When analysts see the event, they then have information about the priority, exploit, and known mitigation readily available.

Message

You can specify meaningful text that appears as a message when the rule triggers. The message gives immediate insight into the nature of the vulnerability that the rule detects attempts to exploit. You can use any printable standard ASCII characters except curly braces (`{}`). The system strips quotes that completely surround the message.



Tip You must specify a rule message. Also, the message cannot consist of white space only, one or more quotation marks only, one or more apostrophes only, or any combination of just white space, quotation marks, or apostrophes.

To define the event message in the intrusion rules editor, you enter the event message in the **Message** field.

Classification

For each rule, you can specify an attack classification that appears in the packet display of the event. The following table lists the name and number for each classification.

Table 124: Rule Classifications

Number	Classification Name	Description
1	not-suspicious	Not Suspicious Traffic
2	unknown	Unknown Traffic
3	bad-unknown	Potentially Bad Traffic
4	attempted-recon	Attempted Information Leak
5	successful-recon-limited	Information Leak
6	successful-recon-largescale	Large Scale Information Leak
7	attempted-dos	Attempted Denial of Service
8	successful-dos	Denial of Service
9	attempted-user	Attempted User Privilege Gain

Number	Classification Name	Description
10	unsuccessful-user	Unsuccessful User Privilege Gain
11	successful-user	Successful User Privilege Gain
12	attempted-admin	Attempted Administrator Privilege Gain
13	successful-admin	Successful Administrator Privilege Gain
14	rpc-portmap-decode	Decode of an RPC Query
15	shellcode-detect	Executable Code was Detected
16	string-detect	A Suspicious String was Detected
17	suspicious-filename-detect	A Suspicious Filename was Detected
18	suspicious-login	An Attempted Login Using a Suspicious Username was Detected
19	system-call-detect	A System Call was Detected
20	tcp-connection	A TCP Connection was Detected
21	trojan-activity	A Network Trojan was Detected
22	unusual-client-port-connection	A Client was Using an Unusual Port
23	network-scan	Detection of a Network Scan
24	denial-of-service	Detection of a Denial of Service Attack
25	non-standard-protocol	Detection of a Non-Standard Protocol or Event
26	protocol-command-decode	Generic Protocol Command Decode
27	web-application-activity	Access to a Potentially Vulnerable Web Application
28	web-application-attack	Web Application Attack
29	misc-activity	Misc Activity
30	misc-attack	Misc Attack
31	icmp-event	Generic ICMP Event
32	inappropriate-content	Inappropriate Content was Detected
33	policy-violation	Potential Corporate Privacy Violation
34	default-login-attempt	Attempt to Login By a Default Username and Password
35	sdf	Sensitive Data
36	malware-cnc	Known malware command and control traffic

Number	Classification Name	Description
37	client-side-exploit	Known client side exploit attempt
38	file-format	Known malicious file or file based exploit

Custom Classification

If you want more customized content for the packet display description of the events generated by a rule you define, you can create a custom classification.

Argument	Description
Classification Name	The name of the classification. The page is difficult to read if you use more than 40 characters. The following characters are not supported: `<>()\'"&$;` and the space character.
Classification Description	A description of the classification. You can use alphanumeric characters and spaces. The following characters are not supported: `<>()\'"&$;`
Priority	High, medium, or low.

Custom Priority

By default, the priority of a rule derives from the event classification for the rule. However, you can override the classification priority for a rule by adding the `priority` keyword to the rule and selecting a high, medium, or low priority. For example, to assign a high priority for a rule that detects web application attacks, add the `priority` keyword to the rule and select **high** as the priority.

Custom Reference

You can use the `reference` keyword to add references to external web sites and additional information about the event. Adding a reference provides analysts with an immediately available resource to help them identify why the packet triggered a rule. The following table lists some of the external systems that can provide data on known exploits and attacks.

Table 125: External Attack Identification Systems

System ID	Description	Example ID
`bugtraq`	Bugtraq page	`8550`
`cve`	Common Vulnerabilities and Exposure ID	`2020-9607`
`mcafee`	McAfee page	`98574`
`url`	Website reference	`www.example.com?exploit=14`
`msb`	Microsoft security bulletin	`MS11-082`

System ID	Description	Example ID
`nessus`	Nessus page	`10039`
`secure-url`	Secure Website Reference (https://...)	`intranet/exploits/exploit=14` Note that you can use `secure-url` with any secure website.

You specify a reference by entering a reference value, as follows:

```
id_system,id
```

where id_system is the system being used as a prefix, and id is the CVE ID number, Arachnids ID, or URL (without `http://`).

For example, to specify the Adobe Acrobat and Reader issue documented in CVE-2020-9607, enter the value:

```
cve,2020-9607
```

Note the following when adding references to a rule:

- Do not use a space after the comma.
- Do not use uppercase letters in the system ID.

Related Topics

Adding a Custom Classification, on page 1554
Defining an Event Priority, on page 1555
Defining an Event Reference, on page 1555

Adding a Custom Classification

In a multidomain deployment, the system displays custom classifications created in the current domain, and you can set the priorities for these classifications. It also displays custom classifications created in ancestor domains, but you cannot set the priorities for these classifications. To view and edit custom classifications created in a lower domain, switch to that domain.

Procedure

Step 1 While creating or editing a rule, choose **Edit Classifications** from the **Classification** drop-down list.

If **View Classifications** displays instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 2 Enter a **Classification Name** and **Classification Description** as described in Intrusion Event Details, on page 1551.

Step 3 Choose a priority for the classification from the **Priority** drop-down list.

Step 4 Click **Add**.

Step 5 Click **Done**.

What to do next

- Continue with creating or editing the rule. See Writing New Rules, on page 1557 or Modifying Existing Rules, on page 1558 for more information.

Related Topics

Custom Rule Creation, on page 1556

Defining an Event Priority

Procedure

Step 1 While creating or editing a rule, choose `priority` from the **Detection Options** drop-down list.

Step 2 Click **Add Option**.

Step 3 Choose a value from the **priority** drop-down list.

Step 4 Click **Save**.

What to do next

- Continue with creating or editing the rule. See Writing New Rules, on page 1557 or Modifying Existing Rules, on page 1558 for more information.

Related Topics

Custom Rule Creation, on page 1556

Defining an Event Reference

Procedure

Step 1 While creating or editing a rule, choose `reference` from the **Detection Options** drop-down list.

Step 2 Click **Add Option**.

Step 3 Enter a value in the **reference** field as described in Intrusion Event Details, on page 1551.

Step 4 Click **Save**.

What to do next

- Continue with creating or editing the rule. See Writing New Rules, on page 1557 or Modifying Existing Rules, on page 1558 for more information.

Related Topics

Custom Rule Creation, on page 1556

Custom Rule Creation

You can create a custom intrusion rule by:

- creating your own standard text rules
- saving existing standard text rules as new
- saving system-provided shared object rules as new
- in a multidomain deployment, saving ancestor rules as new in a descendant domain
- importing a local rule file

The system saves the custom rule in the local rule category, regardless of the method you used to create it.

When you create a custom intrusion rule, the system assigns it a unique rule number, which has the format `GID:SID:Rev`. The elements of this number are:

GID

Generator ID. For all standard text rules, this value is 1. For all shared object rules you save as new, this value is 3.

SID

Snort ID. Indicates whether the rule is a local rule of a system rule. When you create a new rule, the system assigns the next available SID for a local rule.

SID numbers for local rules start at 1000000, and the SID for each new local rule is incremented by one. In a multidomain deployment, the system prepends a domain number to the SID of any custom rule created in or imported into a descendant domain. For example, a rule added in the Global domain would have a SID of 1000000 or greater, and rules added in descendant domains would have SIDs of [domain number]000000 or greater.

Rev

The revision number. For a new rule, the revision number is one. Each time you modify a custom rule the revision number increments by one.

In a custom standard text rule, you set the rule header settings and the rule keywords and arguments. You can use the rule header settings to focus the rule to only match traffic using a specific protocol and traveling to or from specific IP addresses or ports.

In a custom system-provided standard text rule or shared object rule, you are limited to modifying rule header information such as the source and destination ports and IP addresses. You cannot modify the rule keywords or arguments.

Modifying header information for a shared object rule and saving your changes creates a new instance of the rule with a generator ID (GID) of 3 and the next available SID for a custom rule. The system links the new instance of the shared object rule to the reserved `soid` keyword, which maps the rule you create to the rule created by the Cisco Talos Intelligence Group (Talos). You can delete instances of a shared object rule that you create, but you cannot delete shared object rules created by Talos.

Writing New Rules

Procedure

Step 1 Access the intrusion rules using either of the following methods:

- Choose **Policies** > **Access Control** > **Intrusion**, and click **Intrusion Rules**.
- Choose **Objects** > **Intrusion Rules**.

Step 2 Click **Create Rule**.

Step 3 Enter a value in the **Message** field.

Step 4 Choose a value from each of the following drop-down lists:

- **Classification**
- **Action**
- **Protocol**
- **Direction**

Step 5 Enter values in the following fields:

- **Source IPs**
- **Destination IPs**
- **Source Port**
- **Destination Port**

The system uses the value `any` if you do not specify a value for these fields.

Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.

Step 6 Choose a value from the **Detection Options** drop-down list.

Step 7 Click **Add Option**.

Step 8 Enter any arguments for the keyword you added.

Step 9 Optionally, repeat steps 6 to 8.

Step 10 If you added multiple keywords, you can:

- Reorder keywords — Click the up or down arrow next to the keyword you want to move.
- Delete a keyword — Click the **X** next to that keyword.

Step 11 Click **Save As New**.

What to do next

- Enable your new or changed rules within the appropriate intrusion policy; see Viewing Intrusion Rules in an Intrusion Policy, on page 1491.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Modifying Existing Rules

You can modify custom intrusion rules. In a multidomain deployment, you can modify custom intrusion rules that belong to the current domain only.

You can save system-provided rules and rules belonging to ancestor domains as new custom rules in the local rule category, which you can then modify.

Procedure

Step 1 Access the intrusion rules using either of the following methods:

- Choose **Policies** > **Access Control** > **Intrusion**, and click **Intrusion Rules**.
- Choose **Objects** > **Intrusion Rules**.

Step 2 Locate the rule you want to modify. You have the following choices:

- Navigate through the folders to the rule.
- Search for the rule; see Searching for Rules, on page 1561.
- Filter for the group to which the rule belongs; see Filtering Rules, on page 1565.

Step 3 Click **Edit** ( ) next to the rule or, in the case of search results, click the rule message.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Modify the rule as appropriate for the rule type.

Note Do not modify the protocol for a shared object rule; doing so would render the rule ineffective.

Step 5 You have the following choices:

- Click **Save** if you are editing a custom rule and want to overwrite the current version of that rule.
- Click **Save As New** if you are editing a system-provided rule or any rule belonging to an ancestor domain, or if you are editing a custom rule and want to save the changes as a new rule.

What to do next

- If you want to use the local modification of the rule instead of the system-provided rule, deactivate the system-provided rule by using the procedures at Intrusion Rule States, on page 1503 and activate the local rule.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Searching for Rules, on page 1561
Rule Filtering on the Intrusion Rules Editor Page, on page 1562

Viewing Rule Documentation

From the Rule Edit page, you can view rule documentation supplied by the Cisco Talos Intelligence Group (Talos). While viewing, you can click external references to view additional informaton provided by Talos. You can also click **Context Explorer** to view contextual information for events generated by the rule.

Procedure

Step 1 Access an intrusion rule using either of the following methods:

- Choose **Policies** > **Access Control** > **Intrusion**, and click **Intrusion Rules**.
- Choose **Objects** > **Intrusion Rules**.

Step 2 Locate the rule you want to view. You have the following choices:

- Navigate through the folders to the rule.
- Search for the rule; see Searching for Rules, on page 1561.
- Filter for the group to which the rule belongs; see Filtering Rules, on page 1565.

Step 3 Click **Edit** () next to the rule or, in the case of search results, click the rule message.

If **View** () appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Click **View Documentation**.

Step 5 Optionally, click either of the following links:

- References—see Keyword Filtering, on page 1563 and *Custom Reference* in Intrusion Event Details, on page 1551 for information on available external references.
- **Context Explorer**—see for information on viewing event data for the rule in the context explorer.

Tip Selecting an external link closes the documentation pop-up window; to exit the rule edit page without modifying the rule, select any menu path.

Adding Comments to Intrusion Rules

You can add comments to any intrusion rule. Such comments can be helpful to provide context and additional information about the rule and the exploit or policy violation it identifies.

In a multidomain deployment, the system displays comments created in the current domain, which you can delete. It also displays comments created in ancestor domains, which you cannot delete. To view comments created in a lower domain, switch to that domain.

Procedure

Step 1 Access the intrusion rules using either of the following methods:

- Choose **Policies** > **Access Control** > **Intrusion**, and click **Intrusion Rules**.

- Choose **Objects** > **Intrusion Rules**.

Step 2 Locate the rule you want to annotate. You have the following choices:

- Navigate through the folders to the rule.
- Search for the rule; see Searching for Rules, on page 1561.
- Filter for the group where the rule belongs; see Filtering Rules, on page 1565.

Step 3 Click **Edit** ( ) next to the rule or, in the case of search results, click the rule message.

If **View** ( ) appears next to a rule instead, the rule belongs to an ancestor policy, or you do not have permission to modify the rule.

Step 4 Click **Rule Comment**.

Step 5 Enter your comment in the text box.

Step 6 Click **Add Comment**.

Tip You can also add and view rule comments in an intrusion event's packet view.

What to do next

- Continue with creating or editing the rule. See Writing New Rules, on page 1557 or Modifying Existing Rules, on page 1558 for more information.

Related Topics

Searching for Rules, on page 1561
Event Information Fields, on page 2273

Deleting Custom Rules

You can delete custom rules if the rules are not currently enabled in an intrusion policy. You cannot delete either standard text rules or shared object rules provided by the system. In a multidomain deployment, you can delete local rules created in the current domain only.

The system stores deleted rules in the deleted category, and you can use a deleted rule as the basis for a new rule. The Rules page in an intrusion policy does not display the deleted category, so you cannot enable deleted custom rules.

Tip Custom rules include shared object rules that you save with modified header information. The system also saves these in the local rule category and lists them with a GID of 3. You can delete your modified version of a shared object rule, but you cannot delete the original shared object rule.

Procedure

Step 1 Access the intrusion rules using either of the following methods:

- Choose **Policies** > **Access Control** > **Intrusion**, and click **Intrusion Rules**.
- Choose **Objects** > **Intrusion Rules**.

Step 2 You have two choices:

- Delete all local rules — Click **Delete Local Rules**, then click **OK**.
- Delete a single rule — Choose `Local Rules` from the **Group Rules By** drop-down, click **Delete** () next to a rule you want to delete, and click **OK** to confirm the deletion.

Related Topics

Intrusion Rule States, on page 1503

Searching for Rules

The Firepower System provides thousands of standard text rules, and the Cisco Talos Intelligence Group (Talos) continues to add rules as new vulnerabilities and exploits are discovered. You can easily search for specific rules so that you can activate, deactivate, or edit them.

Procedure

Step 1 Access the intrusion rules using either of the following methods:

- Choose **Policies** > **Access Control** > **Intrusion**, and click **Intrusion Rules**.
- Choose **Objects** > **Intrusion Rules**.

Step 2 Click **Search** on the toolbar.

Step 3 Add search criteria.

Step 4 Click **Search**.

What to do next

- If you want to view or edit a located rule (or a copy of the rule, if it is a system rule), click the hyperlinked rule message. See Writing New Rules, on page 1557 or Modifying Existing Rules, on page 1558 for more information.

Search Criteria for Intrusion Rules

The following table describes the available search options:

Table 126: Rule Search Criteria

Option	Description
Signature ID	To search for a single rule based on Snort ID (SID), enter an SID number. To search for multiple rules, enter a comma-separated list of SID numbers. This field has an 80-character limit.

Option	Description
Generator ID	To search for standard text rules, select **1**. To search for shared object rules, select **3**.
Message	To search for a rule with a specific message, enter a single word from the rule message in the **Message** field. For example, to search for DNS exploits, you would enter `DNS`, or to search for buffer overflow exploits, enter `overflow`.
Protocol	To search rules that evaluate traffic of a specific protocol, select the protocol. If you do not select a protocol, search results contain rules for all protocols.
Source Port	To search for rules that inspect packets originating from a specified port, enter a source port number or a port-related variable.
Destination Port	To search for rules that inspect packets destined for a specific port, enter a destination port number or a port-related variable.
Source IP	To search for rules that inspect packets originating from a specified IP address, enter a source IP address or an IP address-related variable.
Destination IP	To search for rules that inspect packets destined for a specified IP address, enter a destination IP address or an IP address-related variable.
Keyword	To search for specific keywords, you can use the keyword search options. You select a keyword and enter a keyword value for which to search. You can also precede the keyword value with an exclamation point (`!`) to match any value other than the specified value.
Category	To search for rules in a specific category, select the category from the **Category** list.
Classification	To search for rules that have a specific classification, select the classification name from the **Classification** list.
Rule State	To search for rules within a specific policy and a specific rule state, select the policy from the first **Rule State** list, and choose a state from the second list to search for rules set to **Generate Events**, **Drop and Generate Events**, or **Disabled**.

Rule Filtering on the Intrusion Rules Editor Page

You can filter the rules on the intrusion rules editor page to display a subset of rules. This can be useful, for example, when you want to modify a rule or change its state but have difficulty finding it among the thousands of rules available.

When you enter a filter, the page displays any folder that includes at least one matching rule, or a message when no rule matches.

Filtering Guidelines

Your filter can include special keywords and their arguments, character strings, and literal character strings in quotes, with spaces separating multiple filter conditions. A filter cannot include regular expressions, wild card characters, or any special operator such as a negation character (!), a greater than symbol (>), less than symbol (<), and so on.

All keywords, keyword arguments, and character strings are case-insensitive. Except for the `gid` and `sid` keywords, all arguments and strings are treated as partial strings. Arguments for `gid` and `sid` return only exact matches.

You can expand a folder on the original, unfiltered page and the folder remains expanded when the subsequent filter returns matches in that folder. This can be useful when the rule you want to find is in a folder that contains a large number of rules.

You cannot constrain a filter with a subsequent filter. Any filter you enter searches the entire rules database and returns all matching rules. When you enter a filter while the page still displays the result of a previous filter, the page clears and returns the result of the new filter instead.

You can use the same features with rules in a filtered or unfiltered list. For example, you can edit rules in a filtered or unfiltered list on the intrusion rules editor page. You can also use any of the options in the context menu for the page.



Tip Filtering may take significantly longer when the combined total of rules in all sub-groups is large because rules appear in multiple categories, even when the total number of unique rules is much smaller.

Keyword Filtering

Each rule filter can include one or more keywords in the format:

```
keyword:argument
```

where keyword is one of the keywords in the following table and argument is a single, case-insensitive, alphanumeric string to search for in the specific field or fields relevant to the keyword.

Arguments for all keywords except `gid` and `sid` are treated as partial strings. For example, the argument `123` returns `"12345"`, `"41235"`, `"45123"`, and so on. The arguments for `gid` and `sid` return only exact matches; for example, `sid:3080` returns only SID 3080.



Tip You can search for a partial SID by filtering with one or more character strings.

The following table describes the specific filtering keywords and arguments you can use to filter rules.

Table 127: Rule Filter Keywords

Keyword	Description	Example
`arachnids`	Returns one or more rules based on all or part of the Arachnids ID in a rule reference.	`arachnids:181`
`bugtraq`	Returns one or more rules based on all or part of the Bugtraq ID in a rule reference.	`bugtraq:2120`
`cve`	Returns one or more rules based on all or part of the CVE number in a rule reference.	`cve:2003-0109`

Keyword	Description	Example
`gid`	The argument `1` returns standard text rules. The argument `3` returns shared object rules.	`gid:3`
`mcafee`	Returns one or more rules based on all or part of the McAfee ID in a rule reference.	`mcafee:10566`
`msg`	Returns one or more rules based on all or part of the rule Message field, also known as the event message.	`msg:chat`
`nessus`	Returns one or more rules based on all or part of the Nessus ID in a rule reference.	`nessus:10737`
`ref`	Returns one or more rules based on all or part of a single alphanumeric string in a rule reference or in the rule Message field.	`ref:MS03-039`
`sid`	Returns the rule with the exact Snort ID.	`sid:235`
`url`	Returns one or more rules based on all or part of the URL in a rule reference.	`url:faqs.org`

Related Topics

Defining an Event Reference, on page 1555
Intrusion Event Details, on page 1551
Preprocessor Generator IDs, on page 2267

Character String Filtering

Each rule filter can include one or more alphanumeric character strings. Character strings search the rule **Message** field, Snort ID (SID), and Generator ID (GID). For example, the string `123` returns the strings `"Lotus123"`, `"123mania"`, and so on in the rule message, and also returns SID 6123, SID 12375, and so on.

All character strings are case-insensitive and are treated as partial strings. For example, any of the strings `ADMIN`, `admin`, or `Admin` return `"admin"`, `"CFADMIN"`, `"Administrator"` and so on.

You can enclose character strings in quotes to return exact matches. For example, the literal string `"overflow attempt"` in quotes returns only that exact string, whereas a filter comprised of the two strings `overflow` and `attempt` without quotes returns `"overflow attempt"`, `"overflow multipacket attempt"`, `"overflow with evasion attempt"`, and so on.

Related Topics

Intrusion Event Details, on page 1551
Preprocessor Generator IDs, on page 2267

Combination Keyword and Character String Filtering

You can narrow filter results by entering any combination of keywords, character strings, or both, separated by spaces. The result includes any rule that matches all the filter conditions.

You can enter multiple filter conditions in any order. For example, each of the following filters returns the same rules:

- `url:at login attempt cve:200`
- `login attempt cve:200 url:at`
- `login cve:200 attempt url:at`

Filtering Rules

On the Intrusion Rules page, you can filter rules into subsets so you can more easily find specific rules. You can then use any of the page features, including choosing any of the features available in the context menu.

Rule filtering can be particularly useful to locate a specific rule to edit.

Procedure

Step 1 Access the intrusion rules using either of the following methods:

- Choose **Policies** > **Access Control** > **Intrusion**, and click **Intrusion Rules**.
- Choose **Objects** > **Intrusion Rules**.

Step 2 Prior to filtering, you have the following choices:

- Expand any rule group you want to expand. Some rule groups also have sub-groups that you can expand.

  Expanding a group on the original, unfiltered page can be useful when you expect that a rule might be in that group. The group remains expanded when the subsequent filter results in a match in that folder, and when you return to the original, unfiltered page by clicking filter **Clear** ( ).

- Choose a different grouping method from the **Group Rules By** drop-down list.

Step 3 Enter filter constraints in the text box next to **Filter** ( ) under the **Group Rules By** list.

Step 4 Press Enter.

Note Clear the current filtered list by clicking filter **Clear** ( ).

Keywords and Arguments in Intrusion Rules

Using the rules language, you can specify the behavior of a rule by combining keywords. Keywords and their associated values (called *arguments*) dictate how the system evaluates packets and packet-related values that the rules engine tests. The Firepower System currently supports keywords that allow you to perform inspection functions, such as content matching, protocol-specific pattern matching, and state-specific matching. You can define up to 100 arguments per keyword, and combine any number of compatible keywords to create highly specific rules. This helps decrease the chance of false positives and false negatives and focus the intrusion information you receive.

Note that you can also use adaptive profile updates in passive deployments to dynamically adapt active rule processing for specific packets based on rule metadata and host information.

Keywords described in this section are listed under Detection Options in the rules editor.

Related Topics

About Adaptive Profiles, on page 1811

The content and protected_content Keywords

Use the `content` keyword or the `protected_content` keyword to specify content that you want to detect in a packet.

You should almost always follow a `content` or `protected_content` keyword by modifiers that indicate where the content should be searched for, whether the search is case sensitive, and other options.

Note that all content matches must be true for the rule to trigger an event, that is, each content match has an AND relationship with the others.

Note also that, in an inline deployment, you can set up rules that match malicious content and then replace it with your own text string of equal length.

content

When you use the `content` keyword, the rules engine searches the packet payload or stream for that string. For example, if you enter `/bin/sh` as the value for one of the `content` keywords, the rules engine searches the packet payload for the string `/bin/sh`.

Match content using either an ASCII string, hexadecimal content (binary byte code), or a combination of both. Surround hexadecimal content with pipe characters (|) in the keyword value. For example, you can mix hexadecimal content and ASCII content using something that looks like `|90C8 C0FF FFFF|/bin/sh`.

You can specify multiple content matches in a single rule. To do this, use additional instances of the `content` keyword. For each content match, you can indicate that content matches must be found in the packet payload or stream for the rule to trigger.



Caution You may invalidate your intrusion policy if you create a rule that includes only one `content` keyword and that keyword has the **Not** option selected.

protected_content

The `protected_content` keyword allows you to encode your search content string before configuring the rule argument. The original rule author uses a hash function (SHA-512, SHA-256, or MD5) to encode the string before configuring the keyword.

When you use the `protected_content` keyword instead of the `content` keyword, there is no change to how the rules engine searches the packet payload or stream for that string and most of the keyword options function as expected. The following table summarizes the exceptions, where the `protected_content` keyword options differ from the `content` keyword options.

Table 128: protected_content Option Exceptions

Option	Description
Hash Type	New option for the `protected_content` rule keyword.
Case Insensitive	Not supported

Option	Description
Within	Not supported
Depth	Not supported
Length	New option for the `protected_content` rule keyword.
Use Fast Pattern Matcher	Not supported
Fast Pattern Matcher Only	Not supported
Fast Pattern Matcher Offset and Length	Not supported

Cisco recommends that you include at least one `content` keyword in rules that include a `protected_content` keyword to ensure that the rules engine uses the fast pattern matcher, which increases processing speed and improves performance. Position the content `keyword` before the `protected_content` keyword in the rule. Note that the rules engine uses the fast pattern matcher when a rule includes at least one `content` keyword, regardless of whether you enable the `content` keyword Use Fast Pattern Matcher argument.

Caution You may invalidate your intrusion policy if you create a rule that includes only one `protected_content` keyword and that keyword has the **Not** option selected.

Related Topics

Custom Rule Creation, on page 1556
Basic content and protected_content Keyword Arguments, on page 1567
The replace Keyword, on page 1577

Basic content and protected_content Keyword Arguments

You can constrain the location and case-sensitivity of content searches with parameters that modify the `content` or `protected_content` keyword. Configure options that modify the `content` or `protected_content` keyword to specify the content for which you want to search.

Case Insensitive

Note This option is **not** supported when configuring the `protected_content` keyword.

You can instruct the rules engine to ignore case when searching for content matches in ASCII strings. To make your search case-insensitive, check **Case Insensitive** when specifying a content search.

Hash Type

Note This option is **only** configurable with the `protected_content` keyword.

Use the **Hash Type** drop-down to identify the hash function you used to encode your search string. The system supports SHA-512, SHA-256, and MD5 hashing for `protected_content` search strings. If the length of your hashed content does not match the selected hash type, the system does **not** save the rule.

The system automatically selects the Cisco-set default value. When **Default** is selected, no specific hash function is written into the rule and the system assumes SHA-512 for the hash function.

Raw Data

The **Raw Data** option instructs the rules engine to analyze the original packet payload before analyzing the normalized payload data (decoded by a network analysis policy) and does not use an argument value. You can use this keyword when analyzing telnet traffic to check the telnet negotiation options in the payload before normalization.

You cannot use the **Raw Data** option together in the same `content` or `protected_content` keyword with any HTTP content option.

Tip You can configure the HTTP Inspect preprocessor **Client Flow Depth** and **Server Flow Depth** options to determine whether raw data is inspected in HTTP traffic, and how much raw data is inspected.

Not

Select the **Not** option to search for content that does not match the specified content. If you create a rule that includes a `content` or `protected_content` keyword with the **Not** option selected, you must also include in the rule at least one other `content` or `protected_content` keyword without the **Not** option selected.

Caution Do not create a rule that includes only one `content` or `protected_content` keyword if that keyword has the **Not** option selected. You may invalidate your intrusion policy.

For example, SMTP rule 1:2541:9 includes three `content` keywords, one of which has the **Not** option selected. A custom rule based on this rule would be invalid if you removed all of the `content` keywords except the one with the **Not** option selected. Adding such a rule to your intrusion policy could invalidate the policy.

Tip You cannot select the **Not** check box and the **Use Fast Pattern Matcher** check box with the same `content` keyword.

content and protected_content Keyword Search Locations

You can use search location options to specify where to begin searching for the specified content and how far to continue searching.

Permitted Combinations: content Search Location Arguments

You can use either of two `content` location pairs to specify where to begin searching for the specified content and how far to continue searching, as follows:

- Use **Offset** and **Depth** together to search relative to the beginning of the packet payload.
- Use **Distance** and **Within** together to search relative to the current search location.

When you specify only one of a pair, the default for the other option in the pair is assumed.

You cannot mix the **Offset** and **Depth** options with the **Distance** and **Within** options. For example, you cannot pair **Offset** and **Within**. You can use any number of location options in a rule.

When no location is specified, the defaults for **Offset** and **Depth** are assumed; that is, the content search starts at the beginning of the packet payload and continues to the end of the packet.

You can also use an existing `byte_extract` variable to specify the value for a location option.



Tip You can use any number of location options in a rule.

Related Topics

The byte_extract Keyword, on page 1582

Permitted Combinations: protected_content Search Location Arguments

Use the required **Length** `protected_content` location option in combination with either the **Offset** or **Distance** location option to specify where to begin searching for the specified content and how far to continue searching, as follows:

- Use **Length** and **Offset** together to search for the protected string relative to the beginning of the packet payload.
- Use **Length** and **Distance** together to search for the protected string relative to the current search location.



Tip You cannot mix the **Offset** and **Distance** options within a single keyword configuration, but you can use any number of location options in a rule.

When no location is specified, the defaults are assumed; that is, the content search starts at the beginning of the packet payload and continues to the end of the packet.

You can also use an existing `byte_extract` variable to specify the value for a location option.

Related Topics

The byte_extract Keyword, on page 1582

content and protected_content Search Location Arguments

Depth



Note This option is **only** supported when configuring the `content` keyword.

Specifies the maximum content search depth, in bytes, from the beginning of the offset value, or if no offset is configured, from the beginning of the packet payload.

For example, in a rule with a content value of `cgi-bin/phf`, and `offset` value of `3`, and a `depth` value of `22`, the rule starts searching for a match to the `cgi-bin/phf` string at byte 3, and stops after processing 22 bytes (byte 25) in packets that meet the parameters specified by the rule header.

You must specify a value that is greater than or equal to the length of the specified content, up to a maximum of 65535 bytes. You cannot specify a value of 0.

The default depth is to search to the end of the packet.

Distance

Instructs the rules engine to identify subsequent content matches that occur a specified number of bytes after the previous successful content match.

Because the distance counter starts at byte 0, specify one less than the number of bytes you want to move forward from the last successful content match. For example, if you specify 4, the search begins at the fifth byte.

You can specify a value of -65535 to 65535 bytes. If you specify a negative `Distance` value, the byte you start searching on may fall outside the beginning of a packet. Any calculations will take into account the bytes outside the packet, even though the search actually starts on the first byte in the packet. For example, if the current location in the packet is the fifth byte, and the next content rule option specifies a `Distance` value of -10 and a `Within` value of 20, the search starts at the beginning of the payload and the `Within` option is adjusted to 15.

The default distance is 0, meaning the current location in the packet subsequent to the last content match.

Length



Note This option is **only** supported when configuring the `protected_content` keyword.

The **Length** `protected_content` keyword option indicates the length, in bytes, of the unhashed search string.

For example, if you used the content `Sample1` to generate a secure hash, use `7` for the **Length** value. You **must** enter a value in this field.

Offset

Specifies in bytes where in the packet payload to start searching for content relative to the beginning of the packet payload. You can specify a value of 65535 to 65535 bytes.

Because the offset counter starts at byte 0, specify one less than the number of bytes you want to move forward from the beginning of the packet payload. For example, if you specify 7, the search begins at the eighth byte.

The default offset is 0, meaning the beginning of the packet.

Within



Note This option is **only** supported when configuring the `content` keyword.

The **Within** option indicates that, to trigger the rule, the next content match must occur within the specified number of bytes after the end of the last successful content match. For example, if you specify a **Within** value of `8`, the next content match must occur within the next eight bytes of the packet payload or it does not meet the criteria that triggers the rule.

You can specify a value that is greater than or equal to the length of the specified content, up to a maximum of 65535 bytes.

The default for **Within** is to search to the end of the packet.

Overview: HTTP content and protected_content Keyword Arguments

HTTP `content` or `protected_content` keyword options let you specify where to search for content matches within an HTTP message decoded by the HTTP Inspect preprocessor.

Two options search status fields in HTTP responses:

- **HTTP Status Code**
- **HTTP Status Message**

Note that although the rules engine searches the raw, unnormalized status fields, these options are listed here separately to simplify explanation below of the restrictions to consider when combining other raw HTTP fields and normalized HTTP fields.

Five options search normalized fields in HTTP requests, responses, or both, as appropriate :

- **HTTP URI**
- **HTTP Method**
- **HTTP Header**
- **HTTP Cookie**
- **HTTP Client Body**

Three options search raw (unnormalized) non-status fields in HTTP requests, responses, or both, as appropriate:

- **HTTP Raw URI**
- **HTTP Raw Header**
- **HTTP Raw Cookie**

Use the following guidelines when selecting HTTP `content` options:

- HTTP `content` options apply only to TCP traffic.
- To avoid a negative impact on performance, select only those parts of the message where the specified content might appear.

  For example, when traffic is likely to include large cookies such as those in shopping cart messages, you might search for the specified content in the HTTP header but not in HTTP cookies.

- To take advantage of HTTP Inspect preprocessor normalization, and to improve performance, any HTTP-related rule you create should at a minimum include at least one `content` or `protected_content` keyword with an **HTTP URI**, **HTTP Method**, **HTTP Header**, or **HTTP Client Body** option selected.

- You cannot use the `replace` keyword in conjunction with HTTP `content` or `protected_content` keyword options.

You can specify a single normalized HTTP option or status field, or use normalized HTTP options and status fields in any combination to target a content area to match. However, note the following restrictions when using HTTP field options:

- You cannot use the **Raw Data** option together in the same `content` or `protected_content` keyword with any HTTP option.
- You cannot use a raw HTTP field option (**HTTP Raw URI**, **HTTP Raw Header**, or **HTTP Raw Cookie**) together in the same `content` or `protected_content` keyword with its normalized counterpart (**HTTP URI**, **HTTP Header**, or **HTTP Cookie**, respectively).
- You cannot select **Use Fast Pattern Matcher** in combination with one or more of the following HTTP field options:

  HTTP Raw URI, **HTTP Raw Header**, **HTTP Raw Cookie**, **HTTP Cookie**, **HTTP Method**, **HTTP Status Message**, or **HTTP Status Code**

  However, you can include the options above in a `content` or `protected_content` keyword that also uses the fast pattern matcher to search one of the following normalized fields:

  HTTP URI, **HTTP Header**, or **HTTP Client Body**

  For example, if you select **HTTP Cookie**, **HTTP Header**, and **Use Fast Pattern Matcher**, the rules engine searches for content in both the HTTP cookie and the HTTP header, but the fast pattern matcher is applied only to the HTTP header, not to the HTTP cookie.

- When you combine restricted and unrestricted options, the fast pattern matcher searches only the unrestricted fields you specify to test whether to pass the rule to the intrusion rules editor for complete evaluation, including evaluation of the restricted fields.

Related Topics

content Keyword Fast Pattern Matcher Arguments, on page 1575

HTTP content and protected_content Keyword Arguments

HTTP URI

Select this option to search for content matches in the normalized request URI field.

Note that you cannot use this option in combination with the `pcre` keyword HTTP URI (U) option to search the same content.



Note A pipelined HTTP request packet contains multiple URIs. When **HTTP URI** is selected and the rules engine detects a pipelined HTTP request packet, the rules engine searches all URIs in the packet for a content match.

HTTP Raw URI

Select this option to search for content matches in the normalized request URI field.

Note that you cannot use this option in combination with the `pcre` keyword HTTP URI (U) option to search the same content.



Note A pipelined HTTP request packet contains multiple URIs. When **HTTP URI** is selected and the rules engine detects a pipelined HTTP request packet, the rules engine searches all URIs in the packet for a content match.

HTTP Method

Select this option to search for content matches in the request method field, which identifies the action such as GET and POST to take on the resource identified in the URI.

HTTP Header

Select this option to search for content matches in the normalized header field, except for cookies, in HTTP requests; also in responses when the HTTP Inspect preprocessor **Inspect HTTP Responses** option is enabled.

Note that you cannot use this option in combination with the `pcre` keyword HTTP header (H) option to search the same content.

HTTP Raw Header

Select this option to search for content matches in the raw header field, except for cookies, in HTTP requests; also in responses when the HTTP Inspect preprocessor I**nspect HTTP Responses** option is enabled.

Note that you cannot use this option in combination with the `pcre` keyword HTTP raw header (D) option to search the same content.

HTTP Cookie

Select this option to search for content matches in any cookie identified in a normalized HTTP client request header; also in response set-cookie data when the HTTP Inspect preprocessor **Inspect HTTP Responses** option is enabled. Note that the system treats cookies included in the message body as body content.

You must enable the HTTP Inspect preprocessor **Inspect HTTP Cookies** option to search only the cookie for a match; otherwise, the rules engine searches the entire header, including the cookie.

Note the following:

- You cannot use this option in combination with the `pcre` keyword HTTP cookie (C) option to search the same content.
- The `Cookie:` and `Set-Cookie:` header names, leading spaces on the header line, and the `CRLF` that terminates the header line are inspected as part of the header and not as part of the cookie.

HTTP Raw Cookie

Select this option to search for content matches in any cookie identified in a raw HTTP client request header; also in response set-cookie data when the HTTP Inspect preprocessor **Inspect HTTP Responses** option is enabled; note that the system treats cookies included in the message body as body content.

You must enable the HTTP Inspect preprocessor **Inspect HTTP Cookies** option to search only the cookie for a match; otherwise, the rules engine searches the entire header, including the cookie.

Note the following:

- You cannot use this option in combination with the `pcre` keyword HTTP raw cookie (K) option to search the same content.

- The `Cookie:` and `Set-Cookie:` header names, leading spaces on the header line, and the `CRLF` that terminates the header line are inspected as part of the header and not as part of the cookie.

HTTP Client Body

Select this option to search for content matches in the message body in an HTTP client request.

Note that for this option to function, you must specify a value of 0 to 65535 for the HTTP Inspect preprocessor **HTTP Client Body Extraction Depth** option.

HTTP Status Code

Select this option to search for content matches in the 3-digit status code in an HTTP response.

You must enable the HTTP Inspect preprocessor **Inspect HTTP Responses** option for this option to return a match.

HTTP Status Message

Select this option to search for content matches in the textual description that accompanies the status code in an HTTP response.

You must enable the HTTP Inspect preprocessor **Inspect HTTP Responses** option for this option to return a match.

Related Topics

pcre Modifier Options, on page 1590
Server-Level HTTP Normalization Options, on page 1709

Overview: content Keyword Fast Pattern Matcher



Note These options are **not** supported when configuring the `protected_content` keyword.

The fast pattern matcher quickly determines which rules to evaluate before passing a packet to the rules engine. This initial determination improves performance by significantly reducing the number of rules used in packet evaluation.

By default, the fast pattern matcher searches packets for the longest content specified in a rule; this is to eliminate as much as possible needless evaluation of a rule. Consider the following example rule fragment:

```
alert tcp any any -> any 80 (msg:"Exploit"; content:"GET";
http_method; nocase; content:"/exploit.cgi"; http_uri;
nocase;)
```

Almost all HTTP client requests contain the content `GET`, but few will contain the content `/exploit.cgi`. Using `GET` as the fast pattern content would cause the rules engine to evaluate this rule in most cases and would rarely result in a match. However, most client `GET` requests would not be evaluated using `/exploit.cgi`, thus increasing performance.

The rules engine evaluates the packet against the rule only when the fast pattern matcher detects the specified content. For example, if one `content` keyword in a rule specifies the content `short`, another specifies `longer`, and a third specifies `longest`, the fast pattern matcher will use the content `longest` and the rule will be evaluated only if the rules engine finds `longest` in the payload.

content Keyword Fast Pattern Matcher Arguments

Use Fast Pattern Matcher

Use this option to specify a shorter search pattern for the fast pattern matcher to use. Ideally, the pattern you specify is less likely to be found in the packet than the longest pattern and, therefore, more specifically identifies the targeted exploit.

Note the following restrictions when selecting **Use Fast Pattern Matcher** and other options in the same `content` keyword:

- You can specify **Use Fast Pattern Matcher** only one time per rule.
- You cannot use **Distance**, **Within**, **Offset**, or **Depth** when you select **Use Fast Pattern Matcher** in combination with **Not**.
- You cannot select Use Fast Pattern Matcher in combination with any of the following HTTP field options:

  HTTP Raw URI, **HTTP Raw Header**, **HTTP Raw Cookie**, **HTTP Cookie**, **HTTP Method**, **HTTP Status Message**, or **HTTP Status Code**

  However, you can include the options above in a `content` keyword that also uses the fast pattern matcher to search one of the following normalized fields:

  HTTP URI, **HTTP Header**, or **HTTP Client Body**

  For example, if you select **HTTP Cookie**, **HTTP Header**, and **Use Fast Pattern Matcher**, the rules engine searches for content in both the HTTP cookie and the HTTP header, but the fast pattern matcher is applied only to the HTTP header, not to the HTTP cookie.

  Note that you cannot use a raw HTTP field option (**HTTP Raw URI**, **HTTP Raw Header**, or **HTTP Raw Cookie**) together in the same `content` keyword with its normalized counterpart (**HTTP URI**, **HTTP Header**, or **HTTP Cookie**, respectively).

  When you combine restricted and unrestricted options, the fast pattern matcher searches only the unrestricted fields you specify to test whether to pass the packet to the rules engine for complete evaluation, including evaluation of the restricted fields.

- Optionally, when you select **Use Fast Pattern Matcher** you can also select **Fast Pattern Matcher Only** or **Fast Pattern Matcher Offset and Length**, but not both.
- You cannot use the fast pattern matcher when inspecting Base64 data.

Fast Pattern Matcher Only

This option allows you to use the `content` keyword only as a fast pattern matcher option and not as a rule option. You can use this option to conserve resources when rules engine evaluation of the specified content is not necessary. For example, consider a case where a rule requires only that the content `12345` be anywhere in the payload. When the fast pattern matcher detects the pattern, the packet can be evaluated against additional keywords in the rule. There is no need for the rules engine to reevaluate the packet to determine if it includes the pattern `12345`.

You would not use this option when the rule contains other conditions relative to the specified content. For example, you would not use this option to search for the content `1234` if another rule condition sought to determine if `abcd` occurs before `1234`. In this case, the rules engine could not determine the relative location because specifying **Fast Pattern Matcher Only** instructs the rules engine not to search for the specified content.

Note the following conditions when using this option:

- The specified content is location-independent; that is, it may occur anywhere in the payload; thus, you cannot use positional options (**Distance**, **Within**, **Offset**, **Depth**, or **Fast Pattern Matcher Offset and Length**).
- You cannot use this option in combination with **Not**.
- You cannot use this option in combination with **Fast Pattern Matcher Offset and Length**.
- The specified content will be treated as case-insensitive, because all patterns are inserted into the fast pattern matcher in a case-insensitive manner; this is handled automatically, so it is not necessary to select **Case Insensitive** when you select this option.
- You should not immediately follow a `content` keyword that uses the **Fast Pattern Matcher Only** option with the following keywords, which set the search location relative to the current search location:
  - `isdataat`
  - `pcre`
  - `content` when **Distance** or **Within** is selected
  - `content` when **HTTP URI** is selected
  - `asn1`
  - `byte_jump`
  - `byte_test`
  - `byte_math`
  - `byte_extract`
  - `base64_decode`

Fast Pattern Matcher Offset and Length

The **Fast Pattern Matcher Offset and Length** option allows you to specify a portion of the content to search. This can reduce memory consumption in cases where the pattern is very long and only a portion of the pattern is sufficient to identify the rule as a likely match. When a rule is selected by the fast pattern matcher, the entire pattern is evaluated against the rule.

You determine the portion for the fast pattern matcher to use by specifying in bytes where to begin the search (offset) and how far into the content (length) to search, using the syntax:

```
offset,length
```

For example, for the content:

```
1234567
```

if you specify the number of offset and length bytes as:

```
1,5
```

the fast pattern matcher searches only for the content `23456`.

Note that you cannot use this option together with **Fast Pattern Matcher Only**.

Related Topics

Overview: HTTP content and protected_content Keyword Arguments, on page 1571
The base64_decode and base64_data Keywords, on page 1652

The replace Keyword

You can use the `replace` keyword in an inline deployment to replace specified content or to replace content in SSL traffic detected by the Cisco SSL Appliance.

To use the `replace` keyword, construct a custom standard text rule that uses the `content` keyword to look for a specific string. Then use the `replace` keyword to specify a string to replace the content. The replace value and content value must be the same length.



Note You **cannot** use the `replace` keyword to replace hashed content in a `protected_content` keyword.

Optionally, you can enclose the replacement string in quotation marks for backward compatibility with previous Firepower System software versions. If you do not include quotation marks, they are added to the rule automatically so the rule is syntactically correct. To include a leading or trailing quotation mark as part of the replacement text, you must use a backslash to escape it, as shown in the following example:

```
"replacement text plus \"quotation\" marks""
```

A rule can contain multiple `replace` keywords, but only one per `content` keyword. Only the first instance of the content found by the rule is replaced.

The following are example uses of the `replace` keyword:

- If the system detects an incoming packet that contains an exploit, you can replace the malicious string with a harmless one. Sometimes this technique is more successful than simply dropping the offending packet. In some attack scenarios, the attacker simply resends the dropped packet until it bypasses your network defenses or floods your network. By substituting one string for another rather than dropping the packet, you may trick the attacker into believing that the attack was launched against a target that was not vulnerable.
- If you are concerned about reconnaissance attacks that try to learn whether you are running a vulnerable version of, for example, a web server, then you can detect the outgoing packet and replace the banner with your own text.



Note Make sure that you set the rule state to Generate Events in the inline intrusion policy where you want to use the replace rule; setting the rule to Drop and Generate events would cause the packet to drop, which would prevent replacing the content.

As part of the string replacement process, the system automatically updates the packet checksums so that the destination host can receive the packet without error.

Note that you cannot use the `replace` keyword in combination with HTTP request message `content` keyword options.

Related Topics

The content and protected_content Keywords, on page 1566
Overview: HTTP content and protected_content Keyword Arguments, on page 1571

The byte_jump Keyword

The `byte_jump` keyword calculates the number of bytes defined in a specified byte segment, and then skips that number of bytes within the packet, either forward from the end of the specified byte segment, or from the beginning or end of the packet payload, or from a point relative to the last content match, depending on the options you specify. This is useful in packets where a specific segment of bytes describe the number of bytes included in variable data within the packet.

The following table describes the arguments required by the `byte_jump` keyword.

Table 129: Required byte_jump Arguments

Argument	Description
Bytes	The number of bytes to pick up from the packet. If used without DCE/RPC, the allowed values are 0 to 10, with the following restrictions: • If used with the `From End` argument, bytes can be 0. If Bytes is 0, the extracted value is 0. • If you specify a number of bytes other than 1, 2, or 4, you must specify a Number Type (hexadecimal, octal, or decimal.) If used with DCE/RPC, allowed values are 1, 2, and 4.
Offset	The number of bytes into the payload to start processing. The `offset` counter starts at byte 0, so calculate the `offset` value by subtracting 1 from the number of bytes you want to jump forward from the beginning of the packet payload or the last successful content match. You can specify -65535 to 65535 bytes. You can also use an existing `byte_extract` variable or `byte_math` result to specify the value for this argument.

The following table describes options you can use to define how the system interprets the values you specified for the required arguments.

Table 130: Additional Optional byte_jump Arguments

Argument	Description
Relative	Makes the offset relative to the last pattern found in the last successful content match.
Align	Rounds the number of converted bytes up to the next 32-bit boundary.

Argument	Description
Multiplier	Indicates the value by which the rules engine should multiply the `byte_jump` value obtained from the packet to get the final `byte_jump` value. That is, instead of skipping the number of bytes defined in a specified byte segment, the rules engine skips that number of bytes multiplied by an integer you specify with the Multiplier argument.
Post Jump Offset	The number of bytes -65535 through 65535 to skip forward or backward after applying other `byte_jump` arguments. A positive value skips forward and a negative value skips backward. Leave the field blank or enter `0` to disable. Note that some `byte_jump` arguments do not apply when you select the **DCE/RPC** argument.
From Beginning	Indicates that the rules engine should skip the specified number of bytes in the payload starting from the beginning of the packet payload, instead of from the current position in the packet.
From End	The jump will originate from the byte that follows the last byte of the buffer.
Bitmask	Applies the specified hexadecimal bitmask using the AND operator to the bytes extracted from the Bytes argument. A bitmask can be 1 to 4 bytes. The result will be right-shifted by the number of bits equal to the number of trailing zeros in the mask.

You can specify only one of **DCE/RPC**, **Endian**, or **Number Type**.

If you want to define how the `byte_jump` keyword calculates the bytes, you can choose from the arguments described in the following table. If you do not select a byte-ordering argument, the rules engine uses big endian byte order.

Table 131: Byte-Ordering byte_jump Arguments

Argument	Description
Big Endian	Processes data in big endian byte order, which is the default network byte order.
Little Endian	Processes data in little endian byte order.
DCE/RPC	Specifies a `byte_jump` keyword for traffic processed by the DCE/RPC preprocessor. The DCE/RPC preprocessor determines big endian or little endian byte order, and the **Number Type** and **Endian** arguments do not apply. When you enable this argument, you can also use `byte_jump` in conjunction with other specific DCE/RPC keywords.

Define how the system views string data in a packet by using one of the arguments in the following table.

Table 132: Number Type Arguments

Argument	Description
Hexadecimal String	Represents converted string data in hexadecimal format.
Decimal String	Represents converted string data in decimal format.
Octal String	Represents converted string data in octal format.

For example, if the values you set for `byte_jump` are as follows:

- Bytes = 4
- Offset = 12
- Relative enabled
- Align enabled

the rules engine calculates the number described in the four bytes that appear 13 bytes after the last successful content match, and skips ahead that number of bytes in the packet. For instance, if the four calculated bytes in a specific packet were `00 00 00 1F`, the rules engine would convert this to 31. Because `align` is specified (which instructs the engine to move to the next 32-bit boundary), the rules engine skips ahead 32 bytes in the packet.

Alternately, if the values you set for `byte_jump` are as follows:

- Bytes = 4
- Offset = 12
- From Beginning enabled
- Multiplier = 2

the rules engine calculates the number described in the four bytes that appear 13 bytes after the beginning of the packet. Then, the engine multiplies that number by two to obtain the total number of bytes to skip. For instance, if the four calculated bytes in a specific packet were `00 00 00 1F`, the rules engine would convert this to 31, then multiply it by two to get 62. Because From Beginning is enabled, the rules engine skips the first 63 bytes in the packet.

Related Topics

The byte_extract Keyword, on page 1582
DCE/RPC Keywords, on page 1614

The byte_test Keyword

The `byte_test` keyword tests the specified byte segment against the Value argument and its operator.

The following table describes the required arguments for the `byte_test` keyword.

Table 133: Required byte_test Arguments

Argument	Description
Bytes	The number of bytes to calculate from the packet. If used without DCE/RPC, the allowed values are 1 to 10. However, if you specify a number of bytes other than 1, 2, or 4, you must specify a Number Type (hexadecimal, octal, or decimal.). If used with DCE/RPC, allowed values are 1, 2, and 4.
Value	Value to test, including its operator. Supported operators: <, >, =, !, &, ^, !>, !<, !=, !&, or !^. For example, if you specify `!1024`, `byte_test` would convert the specified number, and if it did not equal 1024, it would generate an event (if all other keyword parameters matched). Note that ! and != are equivalent. You can also use an existing `byte_extract` variable or `byte_math` result to specify the value for this argument.
Offset	The number of bytes into the payload to start processing. The `offset` counter starts at byte 0, so calculate the `offset` value by subtracting 1 from the number of bytes you want to count forward from the beginning of the packet payload or the last successful content match. You can use an existing `byte_extract` variable or `byte_math` result to specify the value for this argument.

You can further define how the system uses `byte_test` arguments with the arguments described in the following table.

Table 134: Additional Optional byte_test Arguments

Argument	Description
Bitmask	Applies the specified hexadecimal bitmask using the AND operator to the bytes extracted from the Bytes argument. A bitmask can be 1 to 4 bytes. The result will be right-shifted by the number of bits equal to the number of trailing zeros in the mask.
Relative	Makes the offset relative to the last successful pattern match.

You can specify only one of **DCE/RPC**, **Endian**, or **Number Type**.

To define how the `byte_test` keyword calculates the bytes it tests, choose from the arguments in the following table. If you do not select a byte-ordering argument, the rules engine uses big endian byte order.

Table 135: Byte-Ordering byte_test Arguments

Argument	Description
Big Endian	Processes data in big endian byte order, which is the default network byte order.
Little Endian	Processes data in little endian byte order.
DCE/RPC	Specifies a `byte_test` keyword for traffic processed by the DCE/RPC preprocessor. The DCE/RPC preprocessor determines big endian or little endian byte order, and the **Number Type** and **Endian** arguments do not apply. When you enable this argument, you can also use `byte_test` in conjunction with other specific DCE/RPC keywords.

You can define how the system views string data in a packet by using one of the arguments in the following table.

Table 136: Number Type byte-test Arguments

Argument	Description
Hexadecimal String	Represents converted string data in hexadecimal format.
Decimal String	Represents converted string data in decimal format.
Octal String	Represents converted string data in octal format.

For example, if the value for `byte_test` is specified as the following:

- Bytes = 4
- Operator and Value > 128
- Offset = 8
- Relative enabled

The rules engine calculates the number described in the four bytes that appear 9 bytes away from (relative to) the last successful content match, and, if the calculated number is larger than 128 bytes, the rule is triggered.

Related Topics

The byte_extract Keyword, on page 1582
DCE/RPC Keywords, on page 1614

The byte_extract Keyword

You can use the `byte_extract` keyword to read a specified number of bytes from a packet into a variable. You can then use the variable later in the same rule as the value for specific arguments in certain other detection keywords.

This is useful, for example, for extracting data size from packets where a specific segment of bytes describes the number of bytes included in data within the packet. For example, a specific segment of bytes might say that subsequent data is comprised of four bytes; you can extract the data size of four bytes to use as your variable value.

You can use `byte_extract` to create up to two separate variables in a rule concurrently. You can redefine a `byte_extract` variable any number of times; entering a new `byte_extract` keyword with the same variable name and a different variable definition overwrites the previous definition of that variable.

The following table describes the arguments required by the `byte_extract` keyword.

Table 137: Required byte_extract Arguments

Argument	Description
Bytes to Extract	The number of bytes to pick up from the packet. If you specify a number of bytes other than 1, 2, or 4, you must specify a Number Type (hexadecimal, octal, or decimal.)
Offset	The number of bytes into the payload to begin extracting data. You can specify -65535 to 65535 bytes. The offset counter starts at byte 0, so calculate the offset value by subtracting 1 from the number of bytes you want to count forward. For example, specify `7` to count forward 8 bytes. The rules engine counts forward from the beginning of the packet payload or, if you also specify **Relative**, after the last successful content match. Note that you can specify negative numbers only when you also specify **Relative**. You can use an existing `byte_math` result to specify the value for this argument.
Variable Name	The variable name to use in arguments for other detection keywords. You can specify an alphanumeric string that must begin with a letter.

To further define how the system locates the data to extract, you can use the arguments described in the following table.

Table 138: Additional Optional byte_extract Arguments

Argument	Description
Multiplier	A multiplier for the value extracted from the packet. You can specify 0 to 65535. If you do not specify a multiplier, the default value is 1.
Align	Rounds the extracted value to the nearest 2-byte or 4-byte boundary. When you also select **Multiplier**, the system applies the multiplier before the alignment.
Relative	Makes **Offset** relative to the end of the last successful content match instead of the beginning of the payload.
Bitmask	Applies the specified hexadecimal bitmask using the AND operator to the bytes extracted from the Bytes to Extract argument. A bitmask can be 1 to 4 bytes. The result will be right-shifted by the number of bits equal to the number of trailing zeros in the mask.

You can specify only one of **DCE/RPC**, **Endian**, or **Number Type**.

To define how the `byte_extract` keyword calculates the bytes it tests, you can choose from the arguments in the following table. If you do not select a byte-ordering argument, the rules engine uses big endian byte order.

Table 139: Byte-Ordering byte_extract Arguments

Argument	Description
Big Endian	Processes data in big endian byte order, which is the default network byte order.
Little Endian	Processes data in little endian byte order.
DCE/RPC	Specifies a `byte_extract` keyword for traffic processed by the DCE/RPC preprocessor. The DCE/RPC preprocessor determines big endian or little endian byte order, and the **Number Type** and **Endian** arguments do not apply. When you enable this argument, you can also use `byte_extract` in conjunction with other specific DCE/RPC keywords.

You can specify a number type to read data as an ASCII string. To define how the system views string data in a packet, you can select one of the arguments in the following table.

Table 140: Number Type byte_extract arguments

Argument	Description
Hexadecimal String	Reads extracted string data in hexadecimal format.
Decimal String	Reads extracted string data in decimal format.
Octal String	Reads extracted string data in octal format.

For example, if the value for `byte_extract` is specified as the following:

- Bytes to Extract = 4
- Variable Name = var
- Offset = 8
- Relative = enabled

the rules engine reads the number described in the four bytes that appear 9 bytes away from (relative to) the last successful content match into a variable named `var`, which you can specify later in the rule as the value for certain keyword arguments.

The following table lists the keyword arguments where you can specify a variable defined in the `byte_extract` keyword.

Table 141: Arguments Accepting a byte_extract Variable

Keyword	Argument
content	Depth, Offset, Distance, Within
byte_jump	Offset
byte_test	Offset, Value
byte_math	RValue, Offset

Keyword	Argument
isdataat	Offset

Related Topics

The DCE/RPC Preprocessor, on page 1686
DCE/RPC Keywords, on page 1614
Basic content and protected_content Keyword Arguments, on page 1567
The byte_jump Keyword, on page 1578
The byte_test Keyword, on page 1580
Packet Characteristics, on page 1635

The byte_math Keyword

The `byte_math` keyword performs a mathematical operation on an extracted value and a specified value or existing variable, and stores the outcome in a new resulting variable. You can then use the resulting variable as an argument in other keywords.

You can use multiple `byte_math` keywords in a rule to perform multiple byte_math operations.

The following table describes the arguments required by the `byte_math` keyword.

Table 142: Required byte_math Arguments

Argument	Description
Bytes	The number of bytes to calculate from the packet. If used without DCE/RPC, the allowed values are 1 to 10: • Bytes can be 1 to 10 when the operator is +, -. *, or /. • Bytes can be 1 to 4 when the operator is << or >>. • If you specify a number of bytes other than 1, 2, or 4, you must specify a Number Type (hexadecimal, octal, or decimal.) If used with DCE/RPC, allowed values are 1, 2, and 4.
Offset	The number of bytes into the payload to start processing. The `offset` counter starts at byte 0, so calculate the `offset` value by subtracting 1 from the number of bytes you want to jump forward from the beginning of the packet payload or (if you specified Relative) from the last successful content match. You can specify -65535 to 65535 bytes. You can also specify the `byte_extract` variable here.
Operator	+, -, *, /, <<, or >>
RValue	The value following the operator. This can be an unsigned integer or a variable passed from `byte_extract`.

Argument	Description
Result Variable	The name of the variable into which the result of the `byte_math` calculation will be stored. You can use this variable as an argument in other keywords. This value is stored as an unsigned integer. This variable name: • Must use alphanumeric characters • Must not begin with a number • May include special characters supported by the Microsoft filename/variable name convention • Cannot consist entirely of special characters

The following table describes options you can use to define how the system interprets the values you specified for the required arguments.

Table 143: Additional Optional byte_math Arguments

Argument	Description
Relative	Makes the offset relative to the last pattern found in the last successful content match instead of the beginning of the payload.
Bitmask	Applies the specified hexadecimal bitmask using the AND operator to the bytes extracted from the Bytes argument. A bitmask can be 1 to 4 bytes. The result will be right-shifted by the number of bits equal to the number of trailing zeros in the mask.

You can specify only one of **DCE/RPC**, **Endian**, or **Number Type**.

If you want to define how the `byte_math` keyword calculates the bytes, you can choose from the arguments described in the following table. If you do not select a byte-ordering argument, the rules engine uses big endian byte order.

Table 144: Byte-Ordering byte_math Arguments

Argument	Description
Big Endian	Processes data in big endian byte order, which is the default network byte order.
Little Endian	Processes data in little endian byte order.
DCE/RPC	Specifies a `byte_math` keyword for traffic processed by the DCE/RPC preprocessor. The DCE/RPC preprocessor determines big endian or little endian byte order, and the **Number Type** and **Endian** arguments do not apply. When you enable this argument, you can also use `byte_math` in conjunction with other specific DCE/RPC keywords.

Define how the system views string data in a packet by using one of the arguments in the following table.

Table 145: Number Type Arguments

Argument	Description
Hexadecimal String	Represents string data in hexadecimal format.
Decimal String	Represents string data in decimal format.
Octal String	Represents string data in octal format.

For example, if the values you set for `byte_math` are as follows:

- Bytes = 2
- Offset = 0
- Operator = *
- RValue = height
- Result Variable = area

the rules engine extracts the number described in the first two bytes in the packet and multiplies it by the RValue (which uses the existing variable, `height`) to create the new variable, `area`.

Table 146: Arguments Accepting a byte_math Variable

Keyword	Argument
byte_jump	Offset
byte_test	Offset, Value
byte_extract	Offset
isdataat	Offset

Overview: The pcre Keyword

The `pcre` keyword allows you to use Perl-compatible regular expressions (PCRE) to inspect packet payloads for specified content. You can use PCRE to avoid writing multiple rules to match slight variations of the same content.

Regular expressions are useful when searching for content that could be displayed in a variety of ways. The content may have different attributes that you want to account for in your attempt to locate it within a packet's payload.

Note that the regular expression syntax used in intrusion rules is a subset of the full regular expression library and varies in some ways from the syntax used in commands in the full library. When adding a `pcre` keyword using the intrusion rules editor, enter the full value in the following format:

```
!/pcre/ ismxAEGRBUIPHDMCKSY
```

where:

- ! is an optional negation (use this if you want to match patterns that **do not** match the regular expression).
- /pcre/ is a Perl-compatible regular expression.
- ismxAEGRBUIPHDMCKSY is any combination of modifier options.

Also note that you must escape the characters listed in the following table for the rules engine to interpret them correctly when you use them in a PCRE to search for specific content in a packet payload.

Table 147: Escaped PCRE Characters

You must escape...	with a backslash...	or Hex code...
# (hash mark)	\#	\x23
; (semicolon)	\;	\x3B
| (vertical bar)	|	\x7C
: (colon)	\:	\x3A

You can also use m?regex?, where ? is a delimiter other than /. You may want to use this in situations where you need to match a forward slash within a regular expression and do not want to escape it with a backslash. For example, you might use m?regex? ismxAEGRBUIPHDMCKSY where regex is your Perl-compatible regular expression and ismxAEGRBUIPHDMCKSY is any combination of modifier options.

Tip Optionally, you can surround your Perl-compatible regular expression with quote characters, for example, pcre_expression or "pcre_expression". The option of using quotes accommodates experienced users accustomed to previous versions when quotes were required instead of optional. The intrusion rules editor does not display quotation marks when you display a rule after saving it.

pcre Syntax

The pcre keyword accepts standard Perl-compatible regular expression (PCRE) syntax. The following sections describe that syntax.

Tip While this section describes the basic syntax you may use for PCRE, you may want to consult an online reference or book dedicated to Perl and PCRE for more advanced information.

Metacharacters

Metacharacters are literal characters that have special meaning within regular expressions. When you use them within a regular expression, you must "escape" them by preceding them with a backslash.

The following table describes the metacharacters you can use with PCRE and gives examples of each.

Table 148: PCRE Metacharacters

Metacharacter	Description	Example
`.`	Matches any character except newlines. If `s` is used as a modifying option, it also includes newline characters.	`abc.` matches `abcd`, `abc1`, `abc#`, and so on.
`*`	Matches zero or more occurrences of a character or expression.	`abc*` matches `abc`, `abcc`, `abccc`, `abccccc`, and so on.
`?`	Matches zero or one occurrence of a character or expression.	`abc?` matches `abc`.
`+`	Matches one or more occurrences of a character or expression.	`abc+` matches `abc`, `abcc`, `abccc`, `abccccc`, and so on.
`()`	Groups expressions.	`(abc)+` matches `abc`, `abcabc`, `abcabcabc` and so on.
`{}`	Specifies a limit for the number of matches for a character or expression. If you want to set a lower and upper limit, separate the lower limit and upper limit with a comma.	`a{4,6}` matches `aaaa`, `aaaaa`, or `aaaaaa`. `(ab){2}` matches `abab`.
`[]`	Allows you to define character classes, and matches any character or combination of characters described in the set.	`[abc123]` matches `a` or `b` or `c`, and so on.
`^`	Matches content at the beginning of a string. Also used for negation, if used within a character class.	`^in` matches the "in" in `info`, but not in `bin`. `[^a]` matches anything that does not contain `a`.
`$`	Matches content at the end of a string.	`ce$` matches the "`ce`" in `announce`, but not `cent`.
`|`	Indicates an OR expression.	`(MAILTO|HELP)` matches `MAILTO` or `HELP`.
`\`	Allows you to use metacharacters as actual characters and is also used to specify a predefined character class.	`\.` matches a period, `\*` matches an asterisk, `\\` matches a backslash and so on. `\d` matches the numeric characters, `\w` matches alphanumeric characters, and so on.

Character Classes

Character classes include alphabetic characters, numeric characters, alphanumeric characters, and white space characters. While you can create your own character classes within brackets, you can use the predefined classes as shortcuts for different types of character types. When used without additional qualifiers, a character class matches a single digit or character.

The following table describes and provides examples of the predefined character classes accepted by PCRE.

Table 149: PCRE Character Classes

Character Class	Description	Character Class Definition
\d	Matches a numeric character ("digit").	[0-9]
\D	Matches anything that is not an numeric character.	[^0-9]

Character Class	Description	Character Class Definition
\w	Matches an alphanumeric character (“word”).	[a-zA-Z0-9_]
\W	Matches anything that is not an alphanumeric character.	[^a-zA-Z0-9_]
\s	Matches white space characters, including spaces, carriage returns, tabs, newlines, and form feeds.	[\r\t\n\f]
\S	Matches anything that is not a white space character.	[^ \r\t\n\f]

pcre Modifier Options

You can use modifying options after you specify regular expression syntax in the `pcre` keyword’s value. These modifiers perform Perl, PCRE, and Snort-specific processing functions. Modifiers always appear at the end of the PCRE value, and appear in the following format:

```
/pcre/ismxAEGRBUIPHDMCKSY
```

where `ismxAEGRBUPHMC` can include any of the modifying options that appear in the following tables.

Tip Optionally, you can surround the regular expression and any modifying options with quotes, for example, `“/pcre/ismxAEGRBUIPHDMCKSY”`. The option of using quotes accommodates experienced users accustomed to previous versions when quotes were required instead of optional. The intrusion rules editor does not display quotation marks when you display a rule after saving it.

The following table describes options you can use to perform Perl processing functions.

Table 150: Perl-Related Post Regular Expression Options

Option	Description
i	Makes the regular expression case-insensitive.
s	The dot character (.) describes all characters except the newline or `\n` character. You can use `"s"` as an option to override this and have the dot character match all characters, including the newline character.
m	By default, a string is treated as a single line of characters, and `^` and `$` match the beginning and ending of a specific string. When you use `"m"` as an option, `^` and `$` match content immediately before or after any newline character in the buffer, as well as at the beginning or end of the buffer.
x	Ignores white space data characters that may appear within the pattern, except when escaped (preceded by a backslash) or included inside a character class.

The following table describes the PCRE modifiers you can use after the regular expression.

Table 151: PCRE-Related Post Regular Expression Options

Option	Description
A	The pattern must match at the beginning of the string (same as using `^` in a regular expression).
E	Sets `$` to match only at the end of the subject string. (Without `E`, `$` also matches immediately before the final character if it is a newline, but not before any other newline characters).
G	By default, `*` `+` and `?` are "greedy," which means that if two or more matches are found, they will choose the longest match. Use the G character to change this so that these characters always choose the first match unless followed by a question mark character (`?`). For example, `*?` `+?` and `??` would be greedy in a construct using the G modifier, and any incidences of `*`, `+`, or `?` without the additional question mark will not be greedy.

The following table describes the Snort-specific modifiers that you can use after the regular expression.

.

Table 152: Snort-Specific Post Regular Expression Modifiers

Option	Description
R	Searches for matching content relative to the end of the last match found by the rules engine.
B	Searches for the content within data before it is decoded by a preprocessor (this option is similar to using the `Raw Data` argument with the `content` or `protected_content` keyword).
U	Searches for the content within the URI of a normalized HTTP request message decoded by the HTTP Inspect preprocessor. Note that you cannot use this option in combination with the `content` or `protected_content` keyword **HTTP URI** option to search the same content. Note that a pipelined HTTP request packet contains multiple URIs. A PCRE expression that includes the U option causes the rules engine to search for a content match only in the first URI in a pipelined HTTP request packet. To search all URIs in the packet, use the `content` or `protected_content` keyword with **HTTP URI** selected, either with or without an accompanying PCRE expression that uses the U option.
I	Searches for the content within the URI of a raw HTTP request message decoded by the HTTP Inspect preprocessor. Note that you cannot use this option in combination with the `content` or `protected_content` keyword **HTTP Raw URI** option to search the same content
P	Searches for the content within the body of a normalized HTTP request message decoded by the HTTP Inspect preprocessor.

Option	Description
H	Searches for the content within the header, excluding cookies, of an HTTP request or response message decoded by the HTTP Inspect preprocessor. Note that you cannot use this option in combination with the `content` or `protected_content` keyword **HTTP Header** option to search the same content.
D	Searches for the content within the header, excluding cookies, of a raw HTTP request or response message decoded by the HTTP Inspect preprocessor. Note that you cannot use this option in combination with the `content` or `protected_content` keyword **HTTP Raw Header** option to search the same content.
M	Searches for the content within the method field of a normalized HTTP request message decoded by the HTTP Inspect preprocessor; the method field identifies the action such as GET, PUT, CONNECT, and so on to take on the resource identified in the URI.
C	When the HTTP Inspect preprocessor **Inspect HTTP Cookies** option is enabled, searches for the normalized content within any cookie in an HTTP request header, and also within any set-cookie in an HTTP response header when the preprocessor **Inspect HTTP Responses** option is enabled. When **Inspect HTTP Cookies** is not enabled, searches the entire header, including the cookie or set-cookie data. Note the following: • Cookies included in the message body are treated as body content. • You cannot use this option in combination with the `content` or `protected_content` keyword **HTTP Cookie** option to search the same content. • The `Cookie:` and `Set-Cookie:` header names, leading spaces on the header line, and the `CRLF` that terminates the header line are inspected as part of the header and not as part of the cookie.
K	When the HTTP Inspect preprocessor **Inspect HTTP Cookies** option is enabled, searches for the raw content within any cookie in an HTTP request header, and also within any set-cookie in an HTTP response header when the preprocessor **Inspect HTTP Responses** option is enabled. When **Inspect HTTP Cookies** is not enabled, searches the entire header, including the cookie or set-cookie data. Note the following: • Cookies included in the message body are treated as body content. • You cannot use this option in combination with the `content` or `protected_content` keyword **HTTP Raw Cookie** option to search the same content. • The `Cookie:` and `Set-Cookie:` header names, leading spaces on the header line, and the `CRLF` that terminates the header line are inspected as part of the header and not as part of the cookie.
S	Searches the 3-digit status code in an HTTP response.
Y	Searches the textual description that accompanies the status code in an HTTP response.



Note Do not use the U option in combination with the R option. This could cause performance problems. Also, do not use the U option in combination with any other HTTP content option (I, P, H, D, M, C, K, S, or Y).

Related Topics

Overview: HTTP content and protected_content Keyword Arguments, on page 1571

pcre Example Keyword Values

The following examples show values that you could enter for `pcre`, with descriptions of what each example would match.

- **`/feedback[(\d{0,1})]?\.cgi/U`**

This example searches packet payload for `feedback`, followed by zero or one numeric character, followed by `.cgi`, and located only in URI data.

This example would match:

- `feedback.cgi`
- `feedback1.cgi`
- `feedback2.cgi`
- `feedback3.cgi`

This example would **not** match:

- `feedbacka.cgi`
- `feedback11.cgi`
- `feedback21.cgi`
- `feedbackzb.cgi`

- **`/^ez(\w{3,5})\.cgi/iU`**

This example searches packet payload for `ez` at the beginning of a string, followed by a word of 3 to 5 letters, followed by `.cgi`. The search is case-insensitive and only searches URI data.

This example would match:

- `EZBoard.cgi`
- `ezman.cgi`
- `ezadmin.cgi`
- `EZAdmin.cgi`

This example would **not** match:

- `ezez.cgi`
- `fez.cgi`

- `abcezboard.cgi`
- `ezboardman.cgi`
- **`/mail(file|seek)\.cgi/U`**

This example searches packet payload for `mail`, followed by either `file` or `seek`, in URI data.

This example would match:

- `mailfile.cgi`
- `mailseek.cgi`

This example would **not** match:

- `MailFile.cgi`
- `mailfilefile.cgi`
- **`m?http\\x3a\x2f\x2f.*(\n|\t)+?U`**

This example searches packet payload for URI content for a tab or newline character in an HTTP request, after any number of characters. This example uses `m?`regex`?` to avoid using `http\:\/\/` in the expression. Note that the colon is preceded by a backslash.

This example would match:

- `http://www.example.com?scriptvar=x&othervar=\n\..\..`
- `http://www.example.com?scriptvar=\t`

This example would **not** match:

- `ftp://ftp.example.com?scriptvar=&othervar=\n\..\..`
- `http://www.example.com?scriptvar=|/bin/sh -i|`
- **`m?http\\x3a\x2f\x2f.*=\|.*\|+?sU`**

This example searches packet payload for a URL with any number of characters, including newlines, followed by an equal sign, and pipe characters that contain any number of characters or white space. This example uses `m?`regex`?` to avoid using `http\:\/\/` in the expression.

This example would match:

- `http://www.example.com?value=|/bin/sh/ -i|`
- `http://www.example.com?input=|cat /etc/passwd|`

This example would **not** match:

- `ftp://ftp.example.com?value=|/bin/sh/ -i|`
- `http://www.example.com?value=x&input?|cat /etc/passwd|`
- **`/[0-9a-f]{2}\:[0-9a-f]{2}\:[0-9a-f]{2}\:[0-9a-f]{2}\:[0-9a-f]{2}\:[0-9a-f]{2}/i`**

This example searches packet payload for any MAC address. Note that it escapes the colon characters with backslashes.

The metadata Keyword

You can use the `metadata` keyword to add your own descriptive information to a rule. You can also use the `metadata` keyword with `service` arguments to identify applications and ports in network traffic. You can use the information you add to organize or identify rules in ways that suit your needs, and you can search rules for information you add and for `service` arguments.

The system validates metadata based on the argument format:

key value

where *key* and *value* provide a combined description separated by a space. This is the format used by the Cisco Talos Intelligence Group (Talos) for adding metadata to rules provided by Cisco.

Alternatively, you can also use the format:

key = *value*

For example, you could use the *key value* format to identify rules by author and date, using a category and sub-category as follows:

```
author SnortGuru_20050406
```

You can use multiple `metadata` keywords in a rule. You can also use commas to separate multiple *key value* arguments in a single `metadata` keyword, as seen in the following example:

```
author SnortGuru_20050406, revised_by SnortUser1_20050707,
revised_by SnortUser2_20061003,
revised_by SnortUser1_20070123
```

You are not limited to using a *key value* or *key=value* format; however, you should be aware of limitations resulting from validation based on these formats.

Restricted Characters to Avoid

Note the following character restrictions:

- Do not use a semicolon (;) or colon (:).
- The system interprets a comma as a separator for multiple *key value* or *key=value* arguments. For example:

  key value,*key value*,*key value*
- The system interprets the equal to (=) character or space character as separators between *key* and *value*. For example:

  key value

  key=value

All other characters are permitted.

Reserved Metadata to Avoid

Avoid using the following words in a `metadata` keyword, either as a single argument or as the *key* in a *key value* argument; these are reserved for use by Talos:

```
application
engine
impact_flag
os
policy
rule-type
rule-flushing
soid
```



Note Contact Support for assistance in adding restricted metadata to local rules that might not otherwise function as expected.

Impact Level 1

You can use the following reserved *key value* argument in a `metadata` keyword:

```
impact_flag red
```

This *key value* argument sets the impact flag to red (level 1) for a local rule you import or a custom rule you create using the intrusion rules editor.

Note that when Talos includes the `impact_flag red` argument in a rule provided by Cisco, Talos has determined that a packet triggering the rule indicates that the source or destination host is potentially compromised by a virus, trojan, or other piece of malicious software.

Related Topics

Best Practices for Importing Local Intrusion Rules, on page 147
The Intrusion Events Clipboard, on page 2286

Service Metadata

The system detects applications running on the hosts in your network and inserts application protocol information into your network traffic; it does this regardless of the configuration of your discovery policy. You can use `metadata` keyword `service` arguments in a TCP or UDP rule to match application protocols and ports in your network traffic. You can combine one or more `service` application arguments in a rule with a single port argument.

Service Applications

You can use the `metadata` keyword with `service` as the *key* and an application as the *value* to match packets with the identified application protocol. For example, the following *key value* argument in a `metadata` keyword associates the rule with HTTP traffic:

```
service http
```

You can identify multiple applications separated by commas. For example:

```
service http, service smtp, service ftp
```

Caution Adaptive profiling **must** be enabled (its default state) as described in Configuring Adaptive Profiles, on page 1814 for intrusion rules to use service metadata.

The following table describes the most common application values used with the `service` keyword.

Note Contact Support for assistance if you have difficulty identifying applications not in the table.

Table 153: service Values

Value	Description
`cvs`	Concurrent Versions System
`dcerpc`	Distributed Computing Environment/Remote Procedure Calls System
`dns`	Domain Name System
`finger`	Finger user information protocol
`ftp`	File Transfer Protocol
`ftp-data`	File Transfer Protocol (Data Channel)
`http`	Hypertext Transfer Protocol
`imap`	Internet Message Access Protocol
`isakmp`	Internet Security Association and Key Management Protocol
`mysql`	My Structured Query Language
`netbios-dgm`	NETBIOS Datagram Service
`netbios-ns`	NETBIOS Name Service
`netbios-ssn`	NETBIOS Session Service
`nntp`	Network News Transfer Protocol
`oracle`	Oracle Net Services
`shell`	OS Shell
`pop2`	Post Office Protocol, version 2
`pop3`	Post Office Protocol, version 3
`smtp`	Simple Mail Transfer Protocol
`snmp`	Simple Network Management Protocol

Value	Description
`ssh`	Secure Shell network protocol
`sunrpc`	Sun Remote Procedure Call Protocol
`telnet`	Telnet network protocol
`tftp`	Trivial File Transfer Protocol
`x11`	X Window System

Service Ports

You can use the `metadata` keyword with `service` as the *key* and a specified port argument as the *value* to define how the rule matches ports in combination with applications.

You can specify any of the port values in the table below, one value per rule.

Table 154: service Port Values

Value	Description
`else-ports` or `unknown`	The system applies the rule if either of the following conditions is met: • The packet application is known and matches the rule application. • The packet application is unknown and packet ports match the rule ports. The `else-ports` and `unknown` values produce the default behavior that the system uses when `service` specifies an application protocol with no port modifier.
`and-ports`	The system applies the rule if the packet application is known and matches the rule application, and the packet port matches the ports in the rule header. You cannot use `and-ports` in a rule that does not specify an application.
`or-ports`	The system applies the rule if any of the following conditions are met: • The packet application is known and matches the rule application. • The packet application is unknown and packet port matches the rule ports. • The packet application does not match the rule application and packet ports match the rule ports. • The rule does not specify an application and packet ports match the rule ports.

Note the following:

- You must include a `service` application argument with the `service and-ports` argument.
- If a rule specifies more than one of the values in the table above, the system applies the last one that appears in the rule.
- Port and application arguments can be in any order.

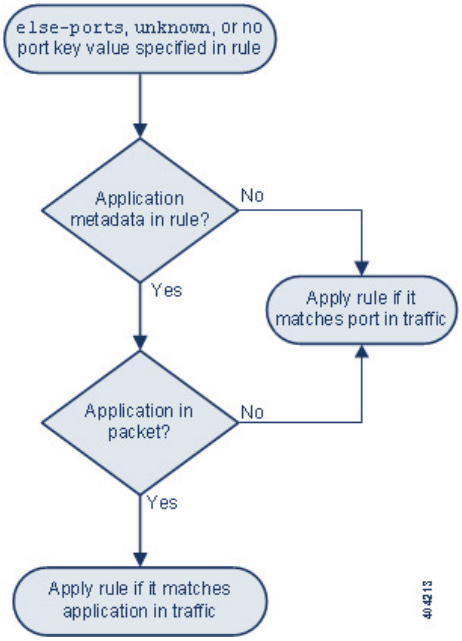
Except for the `and-ports` value, you can include a `service` port argument with or without one or more `service` application arguments. For example:

```
service or-ports, service http, service smtp
```

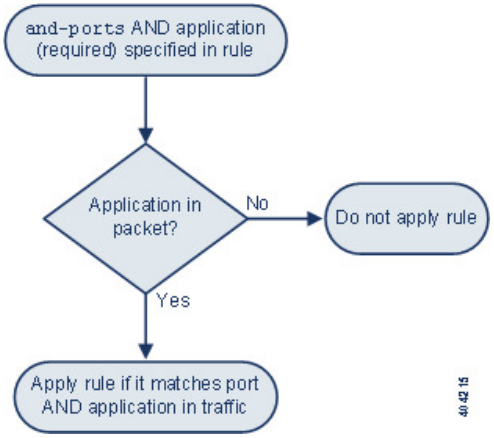
Applications and Ports in Traffic

The diagrams below illustrate the application and port combinations that intrusion rules support, and the results of applying these rule constraints to packet data.

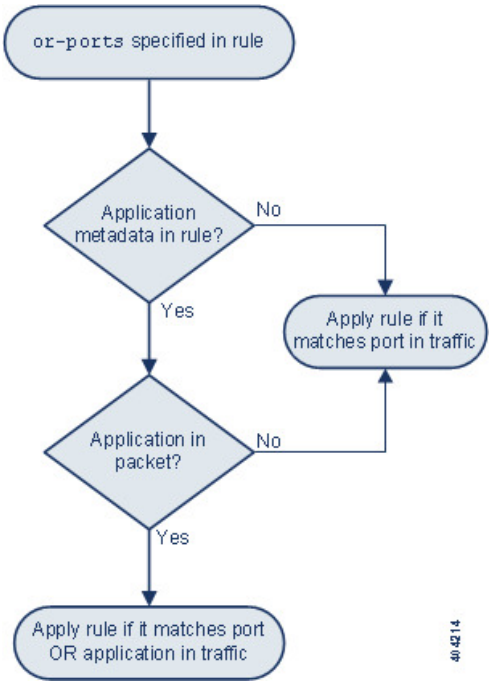
Host application protocol else source/destination ports:



Host application protocol and source/destination ports:



Host application protocol or source/destination ports:



Example Matches

The following sample rules using the `metadata` keyword with `service` arguments are shown with examples of data they match and do not match:

- `alert tcp any any -> any [80,8080] (metadata:service and-ports, service http, service smtp;)`

Example Matches	Example Non-Matches
• HTTP traffic over TCP port 80 • HTTP traffic over TCP port 8080 • SMTP traffic over TCP port 80 • SMTP traffic over TCP port 8080	• POP3 traffic on ports 80 or 8080 • Traffic of unknown application on ports 80 or 8080 • HTTP traffic on port 9999

- `alert tcp any any -> any [80,8080] (metadata:service or-ports, service http;)`

Example Matches	Example Non-Matches
• HTTP traffic on any port • SMTP traffic on port 80 • SMTP traffic on port 8080 • Traffic of unknown application on port 80 and 8080	• Non-HTTP and non-SMTP traffic on ports other than 80 or 8080

- Any of the following rules:
  - `alert tcp any any -> any [80,8080] metadata:service else-ports, service http;)`
  - `alert tcp any any -> any [80,8080] metadata:service unknown, service http;)`
  - `alert tcp any any -> any [80,8080] metadata:service http;)`

Example Matches	Example Non-Matches
• HTTP traffic on any port • port 80 if packet application is unknown • port 8080 if packet application is unknown	• SMTP traffic on ports 80 or 8080 • POP3 traffic on ports 80 or 8080

Metadata Search Guidelines

To search for rules that use the `metadata` keyword, select the `metadata` keyword on the rules Search page and, optionally, type any portion of the metadata. For example, you can type:

- `search` to display all rules where you have used `search` for *key*.
- `search http` to display all rules where you have used `search` for *key* and `http` for *value*.
- `author snortguru` to display all rules where you have used `author` for *key* and `SnortGuru` for *value*.
- `author s` to display all rules where you have used `author` for *key* and any terms such as `SnortGuru` or `SnortUser1` or `SnortUser2` for *value*.

Tip When you search for both *key* and *value* , use the same connecting operator (equal to [=] or a space character) in searches that is used in the *key value* argument in the rule; searches return different results depending on whether you follow *key* with equal to (=) or a space character.

Note that regardless of the format you use to add metadata, the system interprets your metadata search term as all or part of a *key value* or *key=value* argument. For example, the following would be valid metadata that does not follow a *key value* or *key=value* format:

```
ab cd ef gh
```

However, the system would interpret each space in the example as a separator between a *key* and *value* . Thus, you could successfully locate a rule containing the example metadata using any of the following searches for juxtaposed and single terms:

```
cd ef
ef gh
ef
```

but you would not locate the rule using the following search, which the system would interpret as a single *key value* argument:

```
ab ef
```

Related Topics

Searching for Rules, on page 1561

IP Header Values

You can use keywords to identify possible attacks or security policy violations in the IP headers of packets.

fragbits

The `fragbits` keyword inspects the fragment and reserved bits in the IP header. You can check each packet for the Reserved Bit, the More Fragments bit, and the Don't Fragment bit in any combination.

Table 155: Fragbits Argument Values

Argument	Description
R	Reserved bit
M	More Fragments bit
D	Don’t Fragment bit

To further refine a rule using the `fragbits` keyword, you can specify any operator described in the following table after the argument value in the rule.

Table 156: Fragbit Operators

Operator	Description
plus sign (**+**)	The packet must match against all specified bits.
asterisk (*****)	The packet can match against any of the specified bits.
exclamation point (**!**)	The packet meets the criteria if none of the specified bits are set.

For example, to generate an event against packets that have the Reserved Bit set (and possibly any other bits), use `R+` as the `fragbits` value.

id

The `id` keyword tests the IP header fragment identification field against the value you specify in the keyword’s argument. Some denial-of-service tools and scanners set this field to a specific number that is easy to detect. For example, in SID 630, which detects a Synscan portscan, the `id` value is set to `39426`, the static value used as the ID number in packets transmitted by the scanner.

Note `id` argument values must be numeric.

ipopts

The `IPopts` keyword allows you to search packets for specified IP header options. The following table lists the available argument values.

Table 157: IPoption Arguments

Argument	Description
rr	record route
eol	end of list
nop	no operation
ts	time stamp
sec	IP security option
lsrr	loose source routing
ssrr	strict source routing
satid	stream identifier

Analysts most frequently watch for strict and loose source routing because these options may be an indication of a spoofed source IP address.

ip_proto

The `ip_proto` keyword allows you to identify packets with the IP protocol specified as the keyword's value. You can specify the IP protocols as a number, 0 through 255. You can combine these numbers with the following operators: `<`, `>`, or `!`. For example, to inspect traffic with any protocol that is not ICMP, use `!1` as a value to the `ip_proto` keyword. You can also use the `ip_proto` keyword multiple times in a single rule; note, however, that the rules engine interprets multiple instances of the keyword as having a Boolean AND relationship. For example, if you create a rule containing `ip_proto:!3; ip_proto:!6`, the rule ignores traffic using the GGP protocol AND the TCP protocol.

tos

Some networks use the type of service (ToS) value to set precedence for packets traveling on that network. The `tos` keyword allows you to test the packet's IP header ToS value against the value you specify as the keyword's argument. Rules using the `tos` keyword will trigger on packets whose ToS is set to the specified value and that meet the rest of the criteria set forth in the rule.



Note Argument values for `tos` must be numeric.

The ToS field has been deprecated in the IP header protocol and replaced with the Differentiated Services Code Point (DSCP) field.

ttl

A packet's time-to-live (ttl) value indicates how many hops it can make before it is dropped. You can use the `ttl` keyword to test the packet's IP header ttl value against the value, or range of values, you specify as the keyword's argument. It may be helpful to set the `ttl` keyword parameter to a low value such as 0 or 1, as low time-to-live values are sometimes indicative of a traceroute or intrusion evasion attempt. (Note, though, that

the appropriate value for this keyword depends on your managed device placement and network topology.) Use syntax as follows:

- Use an integer from 0 to 255 to set a specific value for the TTL value. You can also precede the value with an equal (=) sign (for example, you can specify `5` or `=5`).
- Use a hyphen (-) to specify a range of TTL values (for example, `0-2` specifies all values 0 through 2, `-5` specifies all values 0 through 5, and `5-` specifies all values 5 through 255).
- Use the greater than (>) sign to specify TTL values greater than a specific value (for example, `>3` specifies all values greater than 3).
- Use the greater than and equal to signs (>=) to specify TTL values greater than or equal to a specific value (for example, `>=3` specifies all values greater than or equal to 3).
- Use the less than (<) sign to specify TTL values less than a specific value (for example, `<3` specifies all values less than 3).
- Use the less than and equal to signs (<=) to specify TTL values less than or equal to a specific value (for example, `<=3` specifies all values less than or equal to 3).

ICMP Header Values

The Firepower System supports keywords that you can use to identify attacks and security policy violations in the headers of ICMP packets. Note, however, that predefined rules exist that detect most ICMP types and codes. Consider enabling an existing rule or creating a local rule based on an existing rule; you may be able to find a rule that meets your needs more quickly than if you build an ICMP rule from scratch.

icmp_id and icmp_seq

The ICMP identification and sequence numbers help associate ICMP replies with ICMP requests. In normal traffic, these values are dynamically assigned to packets. Some covert channel and Distributed Denial of Server (DDoS) programs use static ICMP ID and sequence values. The following keywords allow you to identify ICMP packets with static values.

Keyword	Definition
icmp_id	Inspects an ICMP echo request or reply packet's ICMP ID number. Use a numeric value that corresponds with the ICMP ID number as the argument for the `icmp_id` keyword.
icmp_seq	The `icmp_seq` keyword inspects an ICMP echo request or reply packet's ICMP sequence. Use a numeric value that corresponds with the ICMP sequence number as the argument for the `icmp_seq` keyword.

itype

Use the `itype` keyword to look for packets with specific ICMP message type values. You can specify either a valid ICMP type value or an invalid ICMP type value to test for different types of traffic. For example, attackers may set ICMP type values out of range to cause denial of service and flooding attacks.

You can specify a range for the `itype` argument value using less than (<) and greater than (>).

For example:

- <35
- >36
- 3<>55

icode

ICMP messages sometimes include a code value that provides details when a destination is unreachable.

You can use the `icode` keyword to identify packets with specific ICMP code values. You can choose to specify either a valid ICMP code value or an invalid ICMP code value to test for different types of traffic.

You can specify a range for the `icode` argument value using less than (<) and greater than (>).

For example:

- to find values less than 35, specify `<35`.
- to find values greater than 36, specify `>36`.
- to find values between 3 and 55, specify `3<>55`.



Tip You can use the `icode` and `itype` keywords together to identify traffic that matches both. For example, to identify ICMP traffic that contains an ICMP Destination Unreachable code type with an ICMP Port Unreachable code type, specify an `itype` keyword with a value of 3 (for Destination Unreachable) and an `icode` keyword with a value of 3 (for Port Unreachable).

TCP Header Values and Stream Size

The Firepower System supports keywords that are designed to identify attacks attempted using TCP headers of packets and TCP stream size.

ack

You can use the `ack` keyword to compare a value against a packet's TCP acknowledgment number. The rule triggers if a packet's TCP acknowledgment number matches the value specified for the `ack` keyword.

Argument values for `ack` must be numeric.

flags

You can use the `flags` keyword to specify any combination of TCP flags that, when set in an inspected packet, cause the rule to trigger.



Note In situations where you would traditionally use `A+` as the value for `flags`, you should instead use the `flow` keyword with a value of `established`. Generally, you should use the `flow` keyword with a value of `stateless` when using flags to ensure that all combinations of flags are detected.

You can either check for or ignore the values described in the following table for the `flag` keyword.

Table 158: flag Arguments

Argument	TCP Flag
Ack	Acknowledges data.
Psh	Data should be sent in this packet.
Syn	A new connection.
Urg	Packet contains urgent data.
Fin	A closed connection.
Rst	An aborted connection.
CWR	An ECN congestion window has been reduced. This was formerly the R1 argument, which is still supported for backward compatibility.
ECE	ECN echo. This was formerly the R2 argument, which is still supported for backward compatibility.

When using the `flags` keyword, you can use an operator to indicate how the system performs matches against multiple flags. The following table describes these operators.

Table 159: Operators Used with flags

Operator	Description	Example
all	The packet must contain all specified flags.	Select `Urg` and `all` to specify that a packet must contain the Urgent flag and may contain any other flags.
any	The packet can contain any of the specified flags.	Select `Ack`, `Psh`, and `any` to specify that either or both the `Ack` and `Psh` flags must be set to trigger the rule, and that other flags may also be set on a packet.
not	The packet must **not** contain the specified flag set.	Select `Urg` and `not` to specify that the Urgent flag is **not** set on packets that trigger this rule.

flow

You can use the `flow` keyword to select packets for inspection by a rule based on session characteristics. The `flow` keyword allows you to specify the direction of the traffic flow to which a rule applies, applying rules to either the client flow or server flow. To specify how the `flow` keyword inspects your packets, you can set the direction of traffic you want analyzed, the state of packets inspected, and whether the packets are part of a rebuilt stream.

Stateful inspection of packets occurs when rules are processed. If you want a TCP rule to ignore stateless traffic (traffic without an established session context), you must add the `flow` keyword to the rule and select the **Established** argument for the keyword. If you want a UDP rule to ignore stateless traffic, you must add the `flow` keyword to the rule and select either the **Established** argument or a directional argument, or both. This causes the TCP or UDP rule to perform stateful inspection of a packet.

When you add a directional argument, the rules engine inspects only those packets that have an established state with a flow that matches the direction specified. For example, if you add the `flow` keyword with the `established` argument and the `From Client` argument to a rule that triggers when a TCP or UDP connection is detected, the rules engine only inspects packets that are sent from the client.



Tip For maximum performance, always include a `flow` keyword in a TCP rule or a UDP session rule.

The following table describes the stream-related arguments you can specify for the `flow` keyword:

Table 160: State-Related flow Arguments

Argument	Description
Established	Triggers on established connections.
Stateless	Triggers regardless of the state of the stream processor.

The following table describes the directional options you can specify for the `flow` keyword:

Table 161: flow Directional Arguments

Argument	Description
To Client	Triggers on server responses.
To Server	Triggers on client responses.
From Client	Triggers on client responses.
From Server	Triggers on server responses.

Notice that `From Server` and `To Client` perform the same function, as do `To Server` and `From Client`. These options exist to add context and readability to the rule. For example, if you create a rule designed to detect an attack from a server to a client, use `From Server`. But, if you create a rule designed to detect an attack from the client to the server, use `From Client`.

The following table describes the stream-related arguments you can specify for the `flow` keyword:

Table 162: Stream-Related flow Arguments

Argument	Description
Ignore Stream Traffic	Does not trigger on rebuilt stream packets.
Only Stream Traffic	Triggers only on rebuilt stream packets.

For example, you can use `To Server, Established, Only Stream Traffic` as the value for the `flow` keyword to detect traffic, traveling from a client to the server in an established session, that has been reassembled by the stream preprocessor.

seq

The `seq` keyword allows you to specify a static sequence number value. Packets whose sequence number matches the specified argument trigger the rule containing the keyword. While this keyword is used rarely, it is helpful in identifying attacks and network scans that use generated packets with static sequence numbers.

window

You can use the `window` keyword to specify the TCP window size you are interested in. A rule containing this keyword triggers whenever it encounters a packet with the specified TCP window size. While this keyword is used rarely, it is helpful in identifying attacks and network scans that use generated packets with static TCP window sizes.

stream_size

You can use the `stream_size` keyword in conjunction with the stream preprocessor to determine the size in bytes of a TCP stream, using the format:

```
direction,operator,bytes
```

where bytes is number of bytes. You must separate each option in the argument with a comma (,).

The following table describes the case-insensitive directional options you can specify for the `stream_size` keyword:

Table 163: stream_size Keyword Directional Arguments

Argument	Description
client	triggers on a stream from the client matching the specified stream size.
server	triggers on a stream from the server matching the specified stream size.
both	triggers on traffic from the client and traffic from the server both matching the specified stream size. For example, the argument `both, >, 200` would trigger when traffic from the client is greater than 200 bytes AND traffic from the server is greater than 200 bytes.
either	triggers on traffic from either the client or the server matching the specified stream size, whichever occurs first. For example, the argument `either, >, 200` would trigger when traffic from the client is greater than 200 bytes OR traffic from the server is greater than 200 bytes.

The following table describes the operators you can use with the `stream_size` keyword:

Table 164: stream_size Keyword Argument Operators

Operator	Description
=	equal to
!=	not equal to
>	greater than

Operator	Description
<	less than
>=	greater than or equal to
<=	less than or equal to

For example, you could use `client, >=, 5001216` as the argument for the `stream_size` keyword to detect a TCP stream traveling from a client to a server and greater than or equal to 5001216 bytes.

The stream_reassembly Keyword

You can use the `stream_reassemble` keyword to enable or disable TCP stream reassembly for a single connection when inspected traffic on the connection matches the conditions of the rule. Optionally, you can use this keyword multiple times in a rule.

Use the following syntax to enable or disable stream reassembly:

```
enable|disable, server|client|both, option, option
```

The following table describes the optional arguments you can use with the `stream_reassemble` keyword.

Table 165: stream_reassemble Optional Arguments

Argument	Description
noalert	Generate no events regardless of any other detection options specified in the rule.
fastpath	Ignore the rest of the connection traffic when there is a match.

For example, the following rule disables TCP client-side stream reassembly without generating an event on the connection where a 200 OK status code is detected in an HTTP response:

```
alert tcp any 80 -> any any (flow:to_client, established; content: “200 OK”;
stream_reassemble:disable, client, noalert
```

SSL Keywords

You can use SSL rule keywords to invoke the Secure Sockets Layer (SSL) preprocessor and extract information about SSL version and session state from packets in an encrypted session.

When a client and server communicate to establish an encrypted session using SSL or Transport Layer Security (TLS), they exchange handshake messages. Although the data transmitted in the session is encrypted, the handshake messages are not.

The SSL preprocessor extracts state and version information from specific handshake fields. Two fields within the handshake indicate the version of SSL or TLS used to encrypt the session and the stage of the handshake.

ssl_state

The `ssl_state` keyword can be used to match against state information for an encrypted session. To check for two or more SSL versions used simultaneously, use multiple `ssl_version` keywords in a rule.

When a rule uses the `ssl_state` keyword, the rules engine invokes the SSL preprocessor to check traffic for SSL state information.

For example, to detect an attacker's attempt to cause a buffer overflow on a server by sending a `ClientHello` message with an overly long challenge length and too much data, you could use the `ssl_state` keyword with `client_hello` as an argument then check for abnormally large packets.

Use a comma-separated list to specify multiple arguments for the SSL state. When you list multiple arguments, the system evaluates them using the OR operator. For example, if you specify `client_hello` and `server_hello` as arguments, the system evaluates the rule against traffic that has a `client_hello` OR a `server_hello`.

You can also negate any argument; for example:

```
!client_hello, !unknown
```

To ensure the connection has reached each of a set of states, multiple rules using the ssl_state rule option should be used. The `ssl_state` keyword takes the following identifiers as arguments:

Table 166: `ssl_state` Arguments

Argument	Purpose
`client_hello`	Matches against a handshake message with `ClientHello` as the message type, where the client requests an encrypted session.
`server_hello`	Matches against a handshake message with `ServerHello` as the message type, where the server responds to the client's request for an encrypted session.
`client_keyx`	Matches against a handshake message with `ClientKeyExchange` as the message type, where the client transmits a key to the server to confirm receipt of a key from the server.
`server_keyx`	Matches against a handshake message with `ServerKeyExchange` as the message type, where the client transmits a key to the server to confirm receipt of a key from the server.
`unknown`	Matches against any handshake message type.

ssl_version

The `ssl_version` keyword can be used to match against version information for an encrypted session. When a rule uses the `ssl_version` keyword, the rules engine invokes the SSL preprocessor to check traffic for SSL version information.

For example, if you know there is a buffer overflow vulnerability in SSL version 2, you could use the `ssl_version` keyword with the `sslv2` argument to identify traffic using that version of SSL.

Use a comma-separated list to specify multiple arguments for the SSL version. When you list multiple arguments, the system evaluates them using the OR operator. For example, if you wanted to identify any encrypted traffic that was not using SSLv2, you could add `ssl_version:ssl_v3,tls1.0,tls1.1,tls1.2` to a rule. The rule would evaluate any traffic using SSL Version 3, TLS Version 1.0, TLS Version 1.1, or TLS Version 1.2.

The `ssl_version` keyword takes the following SSL/TLS version identifiers as arguments:

Table 167: `ssl_version` Arguments

Argument	Purpose
`sslv2`	Matches against traffic encoded using Secure Sockets Layer (SSL) Version 2.
`sslv3`	Matches against traffic encoded using Secure Sockets Layer (SSL) Version 3.
`tls1.0`	Matches against traffic encoded using Transport Layer Security (TLS) Version 1.0.
`tls1.1`	Matches against traffic encoded using Transport Layer Security (TLS) Version 1.1.
`tls1.2`	Matches against traffic encoded using Transport Layer Security (TLS) Version 1.2.

The appid Keyword

You can use the `appid` keyword to identify the application protocol, client application, or web application in a packet. For example, you could target a specific application that you know is susceptible to a specific vulnerability.

Within the `appid` keyword of an intrusion rule, click **Configure AppID** to select one or more applications that you want to detect.

Browsing the Available Applications

When you first start to build the condition, the **Available Applications** list is unconstrained and displays every application the system detects, 100 per page:

- To page through the applications, click the arrows underneath the list.
- To display a pop-up window with summary information about the application's characteristics, as well as Internet search links that you can follow, click **Information** ( ) next to an application.

Using Application Filters

To help you find the applications you want to match, you can constrain the **Available Applications** list in the following ways:

- To search for applications, click the **Search by name** prompt above the list, then type a name. The list updates as you type to display matching applications.
- To constrain the applications by applying a filter, use the **Application Filters** list. The **Available Applications** list updates as you apply filters. For your convenience, the system uses an **Unlock icon** to mark applications that the system can identify only in decrypted traffic—not encrypted or unencrypted.

Note If you select one or more filters in the Application Filters list and also search the **Available Applications** list, your selections and the search-filtered **Available Applications** list are combined using an AND operation.

Selecting Applications

To select a single application, select it and click **Add to Rule**. To select all applications in the current constrained view, right-click and select **Select All**.

Application Layer Protocol Values

Although preprocessors perform most of the normalization and inspection of application layer protocol values, you can continue to inspect application layer values using various preprocessor options.

The RPC Keyword

The `rpc` keyword identifies Open Network Computing Remote Procedure Call (ONC RPC) services in TCP or UDP packets. This allows you to detect attempts to identify the RPC programs on a host. Intruders can use an RPC portmapper to determine if any of the RPC services running on your network can be exploited. They can also attempt to access other ports running RPC without using portmapper. The following table lists the arguments that the `rpc` keyword accepts.

Table 168: rpc Keyword Arguments

Argument	Description
application	The RPC application number
procedure	The RPC procedure invoked
version	The RPC version

To specify the arguments for the `rpc` keyword, use the following syntax:

```
application,procedure,version
```

where application is the RPC application number, procedure is the RPC procedure number, and version is the RPC version number. You must specify all arguments for the `rpc` keyword — if you are not able to specify one of the arguments, replace it with an asterisk (`*`).

For example, to search for RPC portmapper (which is the RPC application indicated by the number 100000), with any procedure or version, use `100000,*,*` as the arguments.

The ASN.1 Keyword

The `asn1` keyword allows you to decode a packet or a portion of a packet, looking for various malicious encodings.

The following table describes the arguments for the `asn1` keyword.

Table 169: asn.1 Keyword Arguments

Argument	Description
Bitstring Overflow	Detects invalid, remotely exploitable bitstring encodings.
Double Overflow	Detects a double ASCII encoding that is larger than a standard buffer. This is known to be an exploitable function in Microsoft Windows, but it is unknown at this time which services may be exploitable.

Argument	Description
Oversize Length	Detects ASN.1 type lengths greater than the supplied argument. For example, if you set the Oversize Length to 500, any ASN.1 type greater than 500 triggers the rule.
Absolute Offset	Sets an absolute offset from the beginning of the packet payload. (Remember that the offset counter starts at byte 0.) For example, if you want to decode SNMP packets, set Absolute Offset to 0 and do not set a Relative Offset. Absolute Offset may be positive or negative.
Relative Offset	This is the relative offset from the last successful content match, `pcre`, or `byte_jump`. To decode an ASN.1 sequence right after the content "foo", set Relative Offset to 0, and do not set an Absolute Offset. Relative Offset may be positive or negative. (Remember that the offset counter starts at 0.)

For example, there is a known vulnerability in the Microsoft ASN.1 Library that creates a buffer overflow, allowing an attacker to exploit the condition with a specially crafted authentication packet. When the system decodes the asn.1 data, exploit code in the packet could execute on the host with system-level privileges or could cause a DoS condition. The following rule uses the `asn1` keyword to detect attempts to exploit this vulnerability:

```
alert tcp $EXTERNAL_NET any -> $HOME_NET 445
(flow:to_server, established; content:"|FF|SMB|73|";
nocase; offset:4; depth:5;
asn1:bitstring_overflow,double_overflow,oversize_length 100,
relative_offset 54;)
```

The above rule generates an event against TCP traffic traveling from any IP address defined in the $EXTERNAL_NET variable, from any port, to any IP address defined in the $HOME_NET variable using port 445. In addition, it only executes the rule on established TCP connections to servers. The rule then tests for specific content in specific locations. Finally, the rule uses the `asn1` keyword to detect bitstring encodings and double ASCII encodings and to identify asn.1 type lengths over 100 bytes in length starting 55 bytes from the end of the last successful content match. (Remember that the `offset` counter starts at byte 0.)

The urilen Keyword

You can use the `urilen` keyword in conjunction with the HTTP Inspect preprocessor to inspect HTTP traffic for URIs of a specific length, less than a maximum length, greater than a minimum length, or within a specified range.

After the HTTP Inspect preprocessor normalizes and inspects the packet, the rules engine evaluates the packet against the rule and determines whether the URI matches the length condition specified by the `urilen` keyword. You can use this keyword to detect exploits that attempt to take advantage of URI length vulnerabilities, for example, by creating a buffer overflow that allows the attacker to cause a DoS condition or execute code on the host with system-level privileges.

Note the following when using the `urilen` keyword in a rule:

- In practice, you always use the `urilen` keyword in combination with the `flow:established` keyword and one or more other keywords.
- The rule protocol is always TCP.
- Target ports are always HTTP ports.

You specify the URI length using a decimal number of bytes, less than (<) and greater than (>).

For example:

- specify `5` to detect a URI 5 bytes long.
- specify `< 5` (separated by one space character) to detect a URI less than 5 bytes long.
- specify `> 5` (separated by one space character) to detect a URI greater than 5 bytes long.
- specify `3 <> 5` (with one space character before and after `<>`) to detect a URI between 3 and 5 bytes long inclusive.

For example, there is a known vulnerability in Novell's server monitoring and diagnostics utility iMonitor version 2.4, which comes with eDirectory version 8.8. A packet containing an excessively long URI creates a buffer overflow, allowing an attacker to exploit the condition with a specially crafted packet that could execute on the host with system-level privileges or could cause a DoS condition. The following rule uses the `urilen` keyword to detect attempts to exploit this vulnerability:

```
alert tcp $EXTERNAL_NET any -> $HOME_NET $HTTP_PORTS
(msg:"EXPLOIT eDirectory 8.8 Long URI iMonitor buffer
overflow attempt"; flow:to_server,established;
urilen:> 8192; uricontent:"/nds/"; nocase;
classtype:attempted-admin; sid:x; rev:1;)
```

The above rule generates an event against TCP traffic traveling from any IP address defined in the $EXTERNAL_NET variable, from any port, to any IP address defined in the $HOME_NET variable using the ports defined in the $HTTP_PORTS variable. In addition, packets are evaluated against the rule only on established TCP connections to servers. The rule uses the `urilen` keyword to detect any URI over 8192 bytes in length. Finally, the rule searches the URI for the specific case-insensitive content `/nds/`.

Related Topics

Intrusion Rule Header Protocol, on page 1546
Intrusion Rule Header Source and Destination Ports, on page 1550
Predefined Default Variables, on page 414

DCE/RPC Keywords

The three DCE/RPC keywords described in the following table allow you to monitor DCE/RPC session traffic for exploits. When the system processes rules with these keywords, it invokes the DCE/RPC preprocessor.

Table 170: DCE/RPC Keywords

Use...	In this way...	To detect...
`dce_iface`	alone	packets identifying a specific DCE/RPC service
`dce_opnum`	preceded by `dce_iface`	packets identifying specific DCE/RPC service operations
`dce_stub_data`	preceded by `dce_iface` + `dce_opnum`	stub data defining a specific operation request or response

Note in the table that you should always precede `dce_opnum` with `dce_iface`, and you should always precede `dce_stub_data` with `dce_iface` + `dce_opnum`.

You can also use these DCE/RPC keywords in combination with other rule keywords. Note that for DCE/RPC rules, you use the `byte_jump`, `byte_test`, `and byte_extract` keywords with their **DCE/RPC** arguments selected.

Cisco recommends that you include at least one `content` keyword in rules that include DCE/RPC keywords to ensure that the rules engine uses the fast pattern matcher, which increases processing speed and improves performance. Note that the rules engine uses the fast pattern matcher when a rule includes at least one `content` keyword, regardless of whether you enable the `content` keyword **Use Fast Pattern Matcher** argument.

You can use the DCE/RPC version and adjoining header information as the matching content in the following cases:

- the rule does not include another `content` keyword
- the rule contains another `content` keyword, but the DCE/RPC version and adjoining information represent a more unique pattern than the other content

  For example, the DCE/RPC version and adjoining information are more likely to be unique than a single byte of content.

You should end qualifying rules with one of the following version and adjoining information content matches:

- For connection-oriented DCE/RPC rules, use the content `|05 00 00|` (for major version 05, minor version 00, and the request PDU (protocol data unit) type 00).
- For connectionless DCE/RPC rules, use the content `|04 00|` (for version 04, and the request PDU type 00).

In either case, position the `content` keyword for version and adjoining information as the last keyword in the rule to invoke the fast pattern matcher without repeating processing already completed by the DCE/RPC preprocessor. Note that placing the `content` keyword at the end of the rule applies to version content used as a device to invoke the fast pattern matcher, and not necessarily to other content matches in the rule.

Related Topics

The DCE/RPC Preprocessor, on page 1686
The content and protected_content Keywords, on page 1566
content Keyword Fast Pattern Matcher Arguments, on page 1575
Overview: The byte_jump and byte_test Keywords
The byte_extract Keyword, on page 1582

dce_iface

You can use the `dce_iface` keyword to identify a specific DCE/RPC service.

Optionally, you can also use `dce_iface` in combination with the `dce_opnum` and `dce_stub_data` keywords to further limit the DCE/RPC traffic to inspect.

A fixed, sixteen-byte Universally Unique Identifier (UUID) identifies the application interface assigned to each DCE/RPC service. For example, the UUID 4b324fc8-670-01d3-1278-5a47bf6ee188 identifies the DCE/RPC lanmanserver service, also known as the srvsvc service, which provides numerous management functions for sharing peer-to-peer printers, files, and SMB named pipes. The DCE/RPC preprocessor uses the UUID and associated header values to track DCE/RPC sessions.

The interface UUID is comprised of five hexadecimal strings separated by hyphens:

```
<4hexbytes>-<2hexbytes>-<2hexbytes>-<2hexbytes>-<6hexbytes>
```

You specify the interface by entering the entire UUID including hyphens, as seen in the following UUID for the netlogon interface:

```
12345678-1234-abcd-ef00-01234567cffb
```

Note that you must specify the first three strings in the UUID in big endian byte order. Although published interface listings and protocol analyzers typically display UUIDs in the correct byte order, you might encounter a need to rearrange the UUID byte order before entering it. Consider the following messenger service UUID shown as it might sometimes be displayed in raw ASCII text with the first three strings in little endian byte order:

```
f8 91 7b 5a 00 ff d0 11 a9 b2 00 c0 4f b6 e6 fc
```

You would specify the same UUID for the `dce_iface` keyword by inserting hyphens and putting the first three strings in big endian byte order as follows:

```
5a7b91f8-ff00-11d0-a9b2-00c04fb6e6fc
```

Although a DCE/RPC session can include requests to multiple interfaces, you should include only one `dce_iface` keyword in a rule. Create additional rules to detect additional interfaces.

DCE/RPC application interfaces also have interface version numbers. You can optionally specify an interface version with an operator indicating that the version equals, does not equal, is less than, or greater than the specified value.

Both connection-oriented and connectionless DCE/RPC can be fragmented in addition to any TCP segmentation or IP fragmentation. Typically, it is not useful to associate any DCE/RPC fragment other than the first with the specified interface, and doing so may result in a large number of false positives. However, for flexibility you can optionally evaluate all fragments against the specified interface.

The following table summarizes the `dce_iface` keyword arguments.

Table 171: `dce_iface` Arguments

Argument	Description
Interface UUID	The UUID, including hyphens, that identifies the application interface of the specific service that you want to detect in DCE/RPC traffic. Any request associated with the specified interface would match the interface UUID.
Version	Optionally, the application interface version number 0 to 65535 and an operator indicating whether to detect a version greater than (>), less than (<), equal to (=), or not equal to (!) the specified value.
All Fragments	Optionally, enable to match against the interface in all associated DCE/RPC fragments and, if specified, on the interface version. This argument is disabled by default, indicating that the keyword matches only if the first fragment or the entire unfragmented packet is associated with the specified interface. Note that enabling this argument may result in false positives.

The dce_opnum Keyword

You can use the `dce_opnum` keyword in conjunction with the DCE/RPC preprocessor to detect packets that identify one or more specific operations that a DCE/RPC service provides.

Client function calls request specific service functions, which are referred to in DCE/RPC specifications as *operations*. An operation number (opnum) identifies a specific operation in the DCE/RPC header. It is likely that an exploit would target a specific operation.

For example, the UUID 12345678-1234-abcd-ef00-01234567cffb identifies the interface for the netlogon service, which provides several dozen different operations. One of these is operation 6, the NetrServerPasswordSet operation.

You should precede a `dce_opnum` keyword with a `dce_iface` keyword to identify the service for the operation.

You can specify a single decimal value 0 to 65535 for a specific operation, a range of operations separated by a hyphen, or a comma-separated list of operations and ranges in any order.

Any of the following examples would specify valid netlogon operation numbers:

```
15
15-18
15, 18-20
15, 20-22, 17
15, 18-20, 22, 24-26
```

The dce_stub_data Keyword

You can use the `dce_stub_data` keyword in conjunction with the DCE/RPC preprocessor to specify that the rules engine should start inspection at the beginning of the stub data, regardless of any other rule options. Packet payload rule options that follow the `dce_stub_data` keyword are applied relative to the stub data buffer.

DCE/RPC stub data provides the interface between a client procedure call and the DCE/RPC run-time system, the mechanism that provides the routines and services central to DCE/RPC. DCE/RPC exploits are identified in the stub data portion of the DCE/RPC packet. Because stub data is associated with a specific operation or function call, you should always precede `dce_stub_data` with `dce_iface` and `dce_opnum` to identify the related service and operation.

The `dce_stub_data` keyword has no arguments.

SIP Keywords

Four SIP keywords allow you to monitor SIP session traffic for exploits.

Note that the SIP protocol is vulnerable to denial of service (DoS) attacks. Rules addressing these attacks can benefit from rate-based attack prevention.

The sip_header Keyword

You can use the `sip_header` keyword to start inspection at the beginning of the extracted SIP request or response header and restrict inspection to header fields.

The `sip_header` keyword has no arguments.

The following example rule fragment points to the SIP header and matches the CSeq header field:

```
alert udp any any -> any 5060 ( sip_header; content:"CSeq"; )
```

Related Topics

Dynamic Intrusion Rule States, on page 1511
Rate-Based Attack Prevention, on page 1800

The sip_body Keyword

You can use the `sip_body` keyword to start inspection at the beginning of the extracted SIP request or response message body and restrict inspection to the message body.

The `sip_body` keyword has no arguments.

The following example rule fragment points to the SIP message body and matches a specific IP address in the c (connection information) field in extracted SDP data:

```
alert udp any any -> any 5060 ( sip_body; content:"c=IN 192.168.12.14"; )
```

Note that rules are not limited to searching for SDP content. The SIP preprocessor extracts the entire message body and makes it available to the rules engine.

The sip_method Keyword

A *method* field in each SIP request identifies the purpose of the request. You can use the `sip_method` keyword to test SIP requests for specific methods. Separate multiple methods with commas.

You can specify any of the following currently defined SIP methods:

```
ack, benotify, bye, cancel, do, info, invite, join, message, notify, options, prack,
publish, quath, refer, register, service, sprack, subscribe, unsubscribe, update
```

Methods are case-insensitive. You can separate multiple methods with commas.

Because new SIP methods might be defined in the future, you can also specify a custom method, that is, a method that is not a currently defined SIP method. Accepted field values are defined in RFC 2616, which allows all characters except control characters and separators such as `=`, `(`, and `)`. See RFC 2616 for the complete list of excluded separators. When the system encounters a specified custom method in traffic, it will inspect the packet header but not the message.

The system supports up to 32 methods, including the 21 currently defined methods and an additional 11 methods. The system ignores any undefined methods that you might configure. Note that the 32 total methods includes methods specified using the **Methods to Check** SIP preprocessor option.

You can specify only one method when you use negation. For example:

```
!invite
```

Note, however, that multiple `sip_method` keywords in a rule are linked with an **AND** operation. For example, to test for all extracted methods except `invite` and `cancel`, you would use two negated `sip_method` keywords:

```
sip_method: !invite
sip_method: !cancel
```

Cisco recommends that you include at least one `content` keyword in rules that include the `sip_method` keyword to ensure that the rules engine uses the fast pattern matcher, which increases processing speed and improves performance. Note that the rules engine uses the fast pattern matcher when a rule includes at least one `content` keyword, regardless of whether you enable the `content` keyword **Use Fast Pattern Matcher** argument.

Related Topics

SIP Preprocessor Options, on page 1724
The content and protected_content Keywords, on page 1566
content Keyword Fast Pattern Matcher Arguments, on page 1575

The sip_stat_code Keyword

A three-digit status code in each SIP response indicates the outcome of the requested action. You can use the `sip_stat_code` keyword to test SIP responses for specific status codes.

You can specify a one-digit response-type number 1-9, a specific three-digit number 100-999, or a comma-separated list of any combination of either. A list matches if any single number in the list matches the code in the SIP response.

The following table describes the SIP status code values you can specify.

Table 172: `sip_stat_code` Values

To detect...	Specify...	For example...	Detects...
a specific status code	the three-digit status code	189	189
any three-digit code that begins with a specified single digit	the single digit	1	1xx; that is, 100, 101, 102, and so on
a list of values	any comma-separated combination of specific codes and single digits	222, 3	222 plus 300, 301, 302, and so on

Note also that the rules engine does not use the fast pattern matcher to search for the value specify using the `sip_stat_code` keyword, regardless of whether your rule includes a `content` keyword.

GTP Keywords

Three GSRP Tunneling Protocol (GTP) keywords allow you to inspect the GTP command channel for GTP version, message type, and information elements. You cannot use GTP keywords in combination with other intrusion rule keywords such as `content` or `byte_jump`. You **must** use the `gtp_version` keyword in each rule that uses the `gtp_info` or `gtp_type` keyword.

The gtp_version Keyword

You can use the `gtp_version` keyword to inspect GTP control messages for GTP version 0, 1, or 2.

Because different GTP versions define different message types and information elements, you must use `gtp_version` when you use the `gtp_type` or `gtp_info` keyword. You can specify the value 0, 1, or 2.

The gtp_type Keyword

Each GTP message is identified by a message type, which is comprised of both a numeric value and a string. You can use the `gtp_type` keyword to inspect traffic for specific GTP message types. Because different GTP versions define different message types and information elements, you must also use `gtp_version` when you use the `gtp_type` or `gtp_info` keyword.

You can specify a defined decimal value for a message type, a defined string, or a comma-separated list of either or both in any combination, as seen in the following example:

```
10, 11, echo_request
```

The system uses an OR operation to match each value or string that you list. The order in which you list values and strings does not matter. Any single value or string in the list matches the keyword. You receive an error if you attempt to save a rule that includes an unrecognized string or an out-of-range value.

Note in the table that different GTP versions sometimes use different values for the same message type. For example, the `sgsn_context_request` message type has a value of 50 in GTPv0 and GTPv1, but a value of 130 in GTPv2.

The `gtp_type` keyword matches different values depending on the version number in the packet. In the example above, the keyword matches the message type value 50 in a GTPv0 or GTPv1 packet and the value 130 in a GTPv2 packet. The keyword does not match a packet when the message type value in the packet is not a known value for the version specified in the packet.

If you specify an integer for the message type, the keyword matches if the message type in the keyword matches the value in the GTP packet, regardless of the version specified in the packet.

The following table lists the defined values and strings recognized by the system for each GTP message type.

Table 173: GTP Message Types

Value	Version 0	Version 1	Version 2
1	echo_request	echo_request	echo_request
2	echo_response	echo_response	echo_response
3	version_not_supported	version_not_supported	version_not_supported
4	node_alive_request	node_alive_request	N/A
5	node_alive_response	node_alive_response	N/A
6	redirection_request	redirection_request	N/A
7	redirection_response	redirection_response	N/A
16	create_pdp_context_request	create_pdp_context_request	N/A
17	create_pdp_context_response	create_pdp_context_response	N/A
18	update_pdp_context_request	update_pdp_context_request	N/A
19	update_pdp_context_response	update_pdp_context_response	N/A
20	delete_pdp_context_request	delete_pdp_context_request	N/A
21	delete_pdp_context_response	delete_pdp_context_response	N/A
22	create_aa_pdp_context_request	init_pdp_context_activation_request	N/A
23	create_aa_pdp_context_response	init_pdp_context_activation_response	N/A
24	delete_aa_pdp_context_request	N/A	N/A
25	delete_aa_pdp_context_response	N/A	N/A
26	error_indication	error_indication	N/A
27	pdu_notification_request	pdu_notification_request	N/A
28	pdu_notification_response	pdu_notification_response	N/A

Value	Version 0	Version 1	Version 2
29	pdu_notification_reject_request	pdu_notification_reject_request	N/A
30	pdu_notification_reject_response	pdu_notification_reject_response	N/A
31	N/A	supported_ext_header_notification	N/A
32	send_routing_info_request	send_routing_info_request	create_session_request
33	send_routing_info_response	send_routing_info_response	create_session_response
34	failure_report_request	failure_report_request	modify_bearer_request
35	failure_report_response	failure_report_response	modify_bearer_response
36	note_ms_present_request	note_ms_present_request	delete_session_request
37	note_ms_present_response	note_ms_present_response	delete_session_response
38	N/A	N/A	change_notification_request
39	N/A	N/A	change_notification_response
48	identification_request	identification_request	N/A
49	identification_response	identification_response	N/A
50	sgsn_context_request	sgsn_context_request	N/A
51	sgsn_context_response	sgsn_context_response	N/A
52	sgsn_context_ack	sgsn_context_ack	N/A
53	N/A	forward_relocation_request	N/A
54	N/A	forward_relocation_response	N/A
55	N/A	forward_relocation_complete	N/A
56	N/A	relocation_cancel_request	N/A
57	N/A	relocation_cancel_response	N/A
58	N/A	forward_srns_contex	N/A
59	N/A	forward_relocation_complete_ack	N/A
60	N/A	forward_srns_contex_ack	N/A
64	N/A	N/A	modify_bearer_command
65	N/A	N/A	modify_bearer_failure_indication
66	N/A	N/A	delete_bearer_command
67	N/A	N/A	delete_bearer_failure_indication

Value	Version 0	Version 1	Version 2
68	N/A	N/A	bearer_resource_command
69	N/A	N/A	bearer_resource_failure_indication
70	N/A	ran_info_relay	downlink_failure_indication
71	N/A	N/A	trace_session_activation
72	N/A	N/A	trace_session_deactivation
73	N/A	N/A	stop_paging_indication
95	N/A	N/A	create_bearer_request
96	N/A	mbms_notification_request	create_bearer_response
97	N/A	mbms_notification_response	update_bearer_request
98	N/A	mbms_notification_reject_request	update_bearer_response
99	N/A	mbms_notification_reject_response	delete_bearer_request
100	N/A	create_mbms_context_request	delete_bearer_response
101	N/A	create_mbms_context_response	delete_pdn_request
102	N/A	update_mbms_context_request	delete_pdn_response
103	N/A	update_mbms_context_response	N/A
104	N/A	delete_mbms_context_request	N/A
105	N/A	delete_mbms_context_response	N/A
112	N/A	mbms_register_request	N/A
113	N/A	mbms_register_response	N/A
114	N/A	mbms_deregister_request	N/A
115	N/A	mbms_deregister_response	N/A
116	N/A	mbms_session_start_request	N/A
117	N/A	mbms_session_start_response	N/A
118	N/A	mbms_session_stop_request	N/A
119	N/A	mbms_session_stop_response	N/A
120	N/A	mbms_session_update_request	N/A
121	N/A	mbms_session_update_response	N/A
128	N/A	ms_info_change_request	identification_request

Value	Version 0	Version 1	Version 2
129	N/A	ms_info_change_response	identification_response
130	N/A	N/A	sgsn_context_request
131	N/A	N/A	sgsn_context_response
132	N/A	N/A	sgsn_context_ack
133	N/A	N/A	forward_relocation_request
134	N/A	N/A	forward_relocation_response
135	N/A	N/A	forward_relocation_complete
136	N/A	N/A	forward_relocation_complete_ack
137	N/A	N/A	forward_access
138	N/A	N/A	forward_access_ack
139	N/A	N/A	relocation_cancel_request
140	N/A	N/A	relocation_cancel_response
141	N/A	N/A	configuration_transfer_tunnel
149	N/A	N/A	detach
150	N/A	N/A	detach_ack
151	N/A	N/A	cs_paging
152	N/A	N/A	ran_info_relay
153	N/A	N/A	alert_mme
154	N/A	N/A	alert_mme_ack
155	N/A	N/A	ue_activity
156	N/A	N/A	ue_activity_ack
160	N/A	N/A	create_forward_tunnel_request
161	N/A	N/A	create_forward_tunnel_response
162	N/A	N/A	suspend
163	N/A	N/A	suspend_ack
164	N/A	N/A	resume
165	N/A	N/A	resume_ack
166	N/A	N/A	create_indirect_forward_tunnel_request

Value	Version 0	Version 1	Version 2
167	N/A	N/A	create_indirect_forward_tunnel_response
168	N/A	N/A	delete_indirect_forward_tunnel_request
169	N/A	N/A	delete_indirect_forward_tunnel_response
170	N/A	N/A	release_access_bearer_request
171	N/A	N/A	release_access_bearer_response
176	N/A	N/A	downlink_data
177	N/A	N/A	downlink_data_ack
179	N/A	N/A	pgw_restart
180	N/A	N/A	pgw_restart_ack
200	N/A	N/A	update_pdn_request
201	N/A	N/A	update_pdn_response
211	N/A	N/A	modify_access_bearer_request
212	N/A	N/A	modify_access_bearer_response
231	N/A	N/A	mbms_session_start_request
232	N/A	N/A	mbms_session_start_response
233	N/A	N/A	mbms_session_update_request
234	N/A	N/A	mbms_session_update_response
235	N/A	N/A	mbms_session_stop_request
236	N/A	N/A	mbms_session_stop_response
240	data_record_transfer_request	data_record_transfer_request	N/A
241	data_record_transfer_response	data_record_transfer_response	N/A
254	N/A	end_marker	N/A
255	pdu	pdu	N/A

The gtp_info Keyword

A GTP message can include multiple information elements, each of which is identified by both a defined numeric value and a defined string. You can use the `gtp_info` keyword to start inspection at the beginning of a specified information element, and restrict inspection to the specified information element. Because different GTP versions define different message types and information elements, you must also use `gtp_version` when you use this keyword.

You can specify either the defined decimal value or the defined string for an information element. You can specify a single value or string, and you can use multiple `gtp_info` keywords in a rule to inspect multiple information elements.

When a message includes multiple information elements of the same type, all are inspected for a match. When information elements occur in an invalid order, only the last instance is inspected.

Note that different GTP versions sometimes use different values for the same information element. For example, the `cause` information element has a value of 1 in GTPv0 and GTPv1, but a value of 2 in GTPv2.

The `gtp_info` keyword matches different values depending on the version number in the packet. In the example above, the keyword matches the information element value 1 in a GTPv0 or GTPv1 packet and the value 2 in a GTPv2 packet. The keyword does not match a packet when the information element value in the packet is not a known value for the version specified in the packet.

If you specify an integer for the information element, the keyword matches if the message type in the keyword matches the value in the GTP packet, regardless of the version specified in the packet.

The following table lists the values and strings recognized by the system for each GTP information element.

Table 174: GTP Information Elements

Value	Version 0	Version 1	Version 2
1	cause	cause	imsi
2	imsi	imsi	cause
3	rai	rai	recovery
4	tlli	tlli	N/A
5	p_tmsi	p_tmsi	N/A
6	qos	N/A	N/A
8	recording_required	recording_required	N/A
9	authentication	authentication	N/A
11	map_cause	map_cause	N/A
12	p_tmsi_sig	p_tmsi_sig	N/A
13	ms_validated	ms_validated	N/A
14	recovery	recovery	N/A
15	selection_mode	selection_mode	N/A
16	flow_label_data_1	teid_1	N/A
17	flow_label_signalling	teid_control	N/A
18	flow_label_data_2	teid_2	N/A
19	ms_unreachable	teardown_ind	N/A

Value	Version 0	Version 1	Version 2
20	N/A	nsapi	N/A
21	N/A	ranap	N/A
22	N/A	rab_context	N/A
23	N/A	radio_priority_sms	N/A
24	N/A	radio_priority	N/A
25	N/A	packet_flow_id	N/A
26	N/A	charging_char	N/A
27	N/A	trace_ref	N/A
28	N/A	trace_type	N/A
29	N/A	ms_unreachable	N/A
71	N/A	N/A	apn
72	N/A	N/A	ambr
73	N/A	N/A	ebi
74	N/A	N/A	ip_addr
75	N/A	N/A	mei
76	N/A	N/A	msisdn
77	N/A	N/A	indication
78	N/A	N/A	pco
79	N/A	N/A	paa
80	N/A	N/A	bearer_qos
80	N/A	N/A	flow_qos
82	N/A	N/A	rat_type
83	N/A	N/A	serving_network
84	N/A	N/A	bearer_tft
85	N/A	N/A	tad
86	N/A	N/A	uli
87	N/A	N/A	f_teid
88	N/A	N/A	tmsi

Value	Version 0	Version 1	Version 2
89	N/A	N/A	cn_id
90	N/A	N/A	s103pdf
91	N/A	N/A	s1udf
92	N/A	N/A	delay_value
93	N/A	N/A	bearer_context
94	N/A	N/A	charging_id
95	N/A	N/A	charging_char
96	N/A	N/A	trace_info
97	N/A	N/A	bearer_flag
99	N/A	N/A	pdn_type
100	N/A	N/A	pti
101	N/A	N/A	drx_parameter
103	N/A	N/A	gsm_key_tri
104	N/A	N/A	umts_key_cipher_quin
105	N/A	N/A	gsm_key_cipher_quin
106	N/A	N/A	umts_key_quin
107	N/A	N/A	eps_quad
108	N/A	N/A	umts_key_quad_quin
109	N/A	N/A	pdn_connection
110	N/A	N/A	pdn_number
111	N/A	N/A	p_tmsi
112	N/A	N/A	p_tmsi_sig
113	N/A	N/A	hop_counter
114	N/A	N/A	ue_time_zone
115	N/A	N/A	trace_ref
116	N/A	N/A	complete_request_msg
117	N/A	N/A	guti
118	N/A	N/A	f_container

Value	Version 0	Version 1	Version 2
119	N/A	N/A	f_cause
120	N/A	N/A	plmn_id
121	N/A	N/A	target_id
123	N/A	N/A	packet_flow_id
124	N/A	N/A	rab_contex
125	N/A	N/A	src_rnc_pdcp
126	N/A	N/A	udp_src_port
127	charge_id	charge_id	apn_restriction
128	end_user_address	end_user_address	selection_mode
129	mm_context	mm_context	src_id
130	pdp_context	pdp_context	N/A
131	apn	apn	change_report_action
132	protocol_config	protocol_config	fq_csid
133	gsn	gsn	channel
134	msisdn	msisdn	emlpp_pri
135	N/A	qos	node_type
136	N/A	authentication_qu	fqdn
137	N/A	tft	ti
138	N/A	target_id	mbms_session_duration
139	N/A	utran_trans	mbms_service_area
140	N/A	rab_setup	mbms_session_id
141	N/A	ext_header	mbms_flow_id
142	N/A	trigger_id	mbms_ip_multicast
143	N/A	omc_id	mbms_distribution_ack
144	N/A	ran_trans	rfsp_index
145	N/A	pdp_context_pri	uci
146	N/A	addi_rab_setup	csg_info
147	N/A	sgsn_number	csg_id

Value	Version 0	Version 1	Version 2
148	N/A	common_flag	cmi
149	N/A	apn_restriction	service_indicator
150	N/A	radio_priority_lcs	detach_type
151	N/A	rat_type	ldn
152	N/A	user_loc_info	node_feature
153	N/A	ms_time_zone	mbms_time_to_transfer
154	N/A	imei_sv	throttling
155	N/A	camel	arp
156	N/A	mbms_ue_context	epc_timer
157	N/A	tmp_mobile_group_id	signalling_priority_indication
158	N/A	rim_routing_addr	tmgi
159	N/A	mbms_config	mm_srvcc
160	N/A	mbms_service_area	flags_srvcc
161	N/A	src_rnc_pdcp	nmbr
162	N/A	addi_trace_info	N/A
163	N/A	hop_counter	N/A
164	N/A	plmn_id	N/A
165	N/A	mbms_session_id	N/A
166	N/A	mbms_2g3g_indicator	N/A
167	N/A	enhanced_nsapi	N/A
168	N/A	mbms_session_duration	N/A
169	N/A	addi_mbms_trace_info	N/A
170	N/A	mbms_session_repetition_num	N/A
171	N/A	mbms_time_to_data	N/A
173	N/A	bss	N/A
174	N/A	cell_id	N/A
175	N/A	pdu_num	N/A
177	N/A	mbms_bearer_capab	N/A

Value	Version 0	Version 1	Version 2
178	N/A	rim_routing_disc	N/A
179	N/A	list_pfc	N/A
180	N/A	ps_xid	N/A
181	N/A	ms_info_change_report	N/A
182	N/A	direct_tunnel_flags	N/A
183	N/A	correlation_id	N/A
184	N/A	bearer_control_mode	N/A
185	N/A	mbms_flow_id	N/A
186	N/A	mbms_ip_multicast	N/A
187	N/A	mbms_distribution_ack	N/A
188	N/A	reliable_inter_rat_handover	N/A
189	N/A	rfsp_index	N/A
190	N/A	fqdn	N/A
191	N/A	evolved_allocation1	N/A
192	N/A	evolved_allocation2	N/A
193	N/A	extended_flags	N/A
194	N/A	uci	N/A
195	N/A	csg_info	N/A
196	N/A	csg_id	N/A
197	N/A	cmi	N/A
198	N/A	apn_ambr	N/A
199	N/A	ue_network	N/A
200	N/A	ue_ambr	N/A
201	N/A	apn_ambr_nsapi	N/A
202	N/A	ggsn_backoff_timer	N/A
203	N/A	signalling_priority_indication	N/A
204	N/A	signalling_priority_indication_nsapi	N/A
205	N/A	high_bitrate	N/A

Value	Version 0	Version 1	Version 2
206	N/A	max_mbr	N/A
251	charging_gateway_addr	charging_gateway_addr	N/A
255	private_extension	private_extension	private_extension

SCADA Keywords

The rules engine uses Modbus and DNP3 rules to access certain protocol fields.

Modbus Keywords

You can use Modbus keywords alone or in combination with other keywords such as `content` and `byte_jump`.

modbus_data

You can use the `modbus_data` keyword to point to the beginning of the Data field in a Modbus request or response.

modbus_func

You can use the `modbus_func` keyword to match against the Function Code field in a Modbus application layer request or response header. You can specify either a single defined decimal value or a single defined string for a Modbus function code.

The following table lists the defined values and strings recognized by the system for Modbus function codes.

Table 175: Modbus Function Codes

Value	String
1	read_coils
2	read_discrete_inputs
3	read_holding_registers
4	read_input_registers
5	write_single_coil
6	write_single_register
7	read_exception_status
8	diagnostics
11	get_comm_event_counter
12	get_comm_event_log
15	write_multiple_coils

Value	String
16	write_multiple_registers
17	report_slave_id
20	read_file_record
21	write_file_record
22	mask_write_register
23	read_write_multiple_registers
24	read_fifo_queue
43	encapsulated_interface_transport

modbus_unit

You can use the `modbus_unit` keyword to match a single decimal value against the Unit ID field in a Modbus request or response header.

DNP3 Keywords

You can use DNP3 keywords alone or in combination with other keywords such as `content` and `byte_jump`.

dnp3_data

You can use the `dnp3_data` keyword to point to the beginning of reassembled DNP3 application layer fragments.

The DNP3 preprocessor reassembles link layer frames into application layer fragments. The `dnp3_data` keyword points to the beginning of each application layer fragment; other rule options can match against the reassembled data within fragments without separating the data and adding checksums every 16 bytes.

dnp3_func

You can use the `dnp3_func` keyword to match against the Function Code field in a DNP3 application layer request or response header. You can specify either a single defined decimal value or a single defined string for a DNP3 function code.

The following table lists the defined values and strings recognized by the system for DNP3 function codes.

Table 176: DNP3 Function Codes

Value	String
0	confirm
1	read
2	write
3	select

Value	String
4	operate
5	direct_operate
6	direct_operate_nr
7	immed_freeze
8	immed_freeze_nr
9	freeze_clear
10	freeze_clear_nr
11	freeze_at_time
12	freeze_at_time_nr
13	cold_restart
14	warm_restart
15	initialize_data
16	initialize_appl
17	start_appl
18	stop_appl
19	save_config
20	enable_unsolicited
21	disable_unsolicited
22	assign_class
23	delay_measure
24	record_current_time
25	open_file
26	close_file
27	delete_file
28	get_file_info
29	authenticate_file
30	abort_file
31	activate_config

Value	String
32	authenticate_req
33	authenticate_err
129	response
130	unsolicited_response
131	authenticate_resp

dnp3_ind

You can use the `dnp3_ind` keyword to match against flags in the Internal Indications field in a DNP3 application layer response header.

You can specify the string for a single known flag or a comma-separated list of flags, as seen in the following example:

```
class_1_events, class_2_events
```

When you specify multiple flags, the keyword matches against any flag in the list. To detect a combination of flags, use the `dnp3_ind` keyword multiple times in a rule.

The following list provides the string syntax recognized by the system for defined DNP3 internal indications flags.

```
class_1_events
class_2_events
class_3_events
need_time
local_control
device_trouble
device_restart
no_func_code_support
object_unknown
parameter_error
event_buffer_overflow
already_executing
config_corrupt
reserved_2
reserved_1
```

dnp3_obj

You can use the `dnp3_obj` keyword to match against DNP3 object headers in a request or response.

DNP3 data is comprised of a series of DNP3 objects of different types such as analog input, binary input, and so on. Each type is identified with a *group* such as analog input group, binary input group, and so on, each of which can be identified by a decimal value. The objects in each group are further identified by an *object variation* such as 16-bit integers, 32-bit integers, short floating point, and so on, each of which specifies the data format of the object. Each type of object variation can also be identified by a decimal value.

You identify object headers by specifying the decimal number for the type of object header group and the decimal number for the type of object variation. The combination of the two defines a specific type of DNP3 object.

Packet Characteristics

You can write rules that only generate events against packets with specific packet characteristics.

dsize

The `dsize` keyword tests the packet payload size. With it, you can use the greater than and less than operators (< and >) to specify a range of values. You can use the following syntax to specify ranges:

```
>number_of_bytes
<number_of_bytes
number_of_bytes<>number_of_bytes
```

For example, to indicate a packet size greater than 400 bytes, use `>400` as the `dtype` value. To indicate a packet size of less than 500 bytes, use `<500`. To specify that the rule trigger against any packet between 400 and 500 bytes inclusive, use `400<>500`.

Caution The `dsize` keyword tests packets before they are decoded by any preprocessors.

isdataat

The `isdataat` keyword instructs the rules engine to verify that data resides at a specific location in the payload.

The following table lists the arguments you can use with the `isdataat` keyword.

Table 177: isdataat Arguments

Argument	Type	Description
Offset	Required	The specific location in the payload. For example, to test that data appears at byte 50 in the packet payload, you would specify `50` as the offset value. A `!` modifier negates the results of the `isdataat` test; it alerts if a certain amount of data is not present within the payload. You can also use an existing `byte_extract` variable or `byte_math` result to specify the value for this argument.
Relative	Optional	Makes the location relative to the last successful content match. If you specify a relative location, note that the counter starts at byte 0, so calculate the location by subtracting 1 from the number of bytes you want to move forward from the last successful content match. For example, to specify that the data must appear at the ninth byte after the last successful content match, you would specify a relative offset of `8`.
Raw Data	Optional	Specifies that the data is located in the original packet payload before decoding or application layer normalization by any Firepower System preprocessor. You can use this argument with **Relative** if the previous content match was in the raw packet data.

For example, in a rule searching for the content `foo`, if the value for `isdataat` is specified as the following:

- Offset = !10
- Relative = enabled

The system alerts if the rules engine does not detect 10 bytes after `foo` before the payload ends.

sameip

The `sameip` keyword tests that a packet's source and destination IP addresses are the same. It does not take an argument.

fragoffset

The `fragoffset` keyword tests the offset of a fragmented packet. This is useful because some exploits (such as WinNuke denial-of-service attacks) use hand-generated packet fragments that have specific offsets.

For example, to test whether the offset of a fragmented packet is 31337 bytes, specify `31337` as the `fragoffset` value.

You can use the following operators when specifying arguments for the `fragoffset` keyword.

Table 178: fragoffset Keyword Argument Operators

Operator	Description
!	not
>	greater than
<	less than

Note that you cannot use the not (`!`) operator in combination with `<` or `>`.

cvs

The `cvs` keyword tests Concurrent Versions System (CVS) traffic for malformed CVS entries. An attacker can use a malformed entry to force a heap overflow and execute malicious code on the CVS server. This keyword can be used to identify attacks against two known CVS vulnerabilities: CVE-2004-0396 (CVS 1.11.x up to 1.11.15, and 1.12.x up to 1.12.7) and CVS-2004-0414 (CVS 1.12.x through 1.12.8, and 1.11.x through 1.11.16). The `cvs` keyword checks for a well-formed entry, and generates alerts when a malformed entry is detected.

Your rule should include the ports where CVS runs. In addition, any ports where traffic may occur should be added to the list of ports for stream reassembly in your TCP policies so state can be maintained for CVS sessions. The TCP ports 2401 (`pserver`) and 514 (`rsh`) are included in the list of client ports where stream reassembly occurs. However, note that if your server runs as an `xinetd` server (i.e., pserver), it can run on any TCP port. Add any non-standard ports to the stream reassembly **Client Ports** list.

Related Topics

The byte_extract Keyword, on page 1582
TCP Stream Preprocessing Options, on page 1780

Active Response Keywords

The **resp** and **react** keywords provide two approaches to initiating active responses. An intrusion rule that contains either keyword initiates a single active response when a packet triggers the rule. Active response keywords initiate active responses to close TCP connections in response to triggered TCP rules or UDP sessions in response to triggered UDP rules. See Active Responses in Intrusion Drop Rules, on page 1759. Active responses are not intended to take the place of a firewall for a number of reasons, including that an attacker may have chosen to ignore or circumvent active responses.

Active responses are supported in inline, including routed or transparent, deployments. For example, in response to the `react` keyword in an inline deployment, the system can insert a TCP reset (RST) packet directly into the traffic for each end of the connection, which normally should close the connection. Active responses are not supported or suited for passive deployments.

Because active responses can be routed back, the system does not allow TCP resets to initiate TCP resets; this prevents an unending sequence of active responses. The system also does not allow ICMP unreachable packets to initiate ICMP unreachable packets in keeping with standard practice.

You can configure the TCP stream preprocessor to detect additional traffic on a TCP connection after an intrusion rule has triggered an active response. When the preprocessor detects additional traffic, it sends additional active responses up to a specified maximum to both ends of the connection or session. See **Maximum Active Responses** and **Minimum Response Seconds** in Advanced Transport/Network Preprocessor Options, on page 1759.

Related Topics

Active Responses in Intrusion Drop Rules, on page 1759

The resp Keyword

You can use the `resp` keyword to actively respond to TCP connections or UDP sessions, depending on whether you specify the TCP or UDP protocol in the rule header.

Keyword arguments allow you to specify the packet direction and whether to use TCP reset (RST) packets or ICMP unreachable packets as active responses.

You can use any of the TCP reset or ICMP unreachable arguments to close TCP connections. You should use only ICMP unreachable arguments to close UDP sessions.

Different TCP reset arguments also allow you to target active responses to the packet source, destination, or both. All ICMP unreachable arguments target the packet source and allow you to specify whether to use an ICMP network, host, or port unreachable packet, or all three.

The following table lists the arguments you can use with the `resp` keyword to specify exactly what you want the Firepower System to do when the rule triggers.

Table 179: resp Arguments

Argument	Description
reset_source	Directs a TCP reset packet to the endpoint that sent the packet that triggered the rule. Alternatively, you can specify `rst_snd`, which is supported for backward compatibility.
reset_dest	Directs a TCP reset packet to the intended destination endpoint of the packet that triggered the rule. Alternatively, you can specify `rst_rcv`, which is supported for backward compatibility.
reset_both	Directs a TCP reset packet to both the sending and receiving endpoints. Alternatively, you can specify `rst_all`, which is supported for backward compatibility.
icmp_net	Directs an ICMP network unreachable message to the sender.
icmp_host	Directs an ICMP host unreachable message to the sender.
icmp_port	Directs an ICMP port unreachable message to the sender. This argument is used to terminate UDP traffic.

Argument	Description
icmp_all	Directs the following ICMP messages to the sender: • network unreachable • host unreachable • port unreachable

For example, to configure a rule to reset both sides of a connection when a rule is triggered, use `reset_both` as the value for the `resp` keyword.

You can use a comma-separated list to specify multiple arguments as follows:

```
argument,argument,argument
```

Related Topics

The config response Command

The react Keyword

You can use the `react` keyword to send a default HTML page to the TCP connection client when a packet triggers the rule; after sending the HTML page, the system uses TCP reset packets to initiate active responses to both ends of the connection. The `react` keyword does not trigger active responses for UDP traffic.

Optionally, you can specify the following argument:

```
msg
```

When a packet triggers a `react` rule that uses the `msg` argument, the HTML page includes the rule event message.

If you do not specify the `msg` argument, the HTML page includes the following message:

```
You are attempting to access a forbidden site.
Consult your system administrator for details.
```

Note Because active responses can be routed back, ensure that the HTML response page does not trigger a `react` rule; this could result in an unending sequence of active responses. Cisco recommends that you test `react` rules extensively before activating them in a production environment.

Related Topics

Rule Anatomy, on page 1544
The config response Command

The detection_filter Keyword

You can use the `detection_filter` keyword to prevent a rule from generating events unless a specified number of packets trigger the rule within a specified time. This can stop the rule from prematurely generating

events. For example, two or three failed login attempts within a few seconds could be expected behavior, but a large number of attempts within the same time could indicate a brute force attack.

The `detection_filter` keyword requires arguments that define whether the system tracks the source or destination IP address, the number of times the detection criteria must be met before triggering an event, and how long to continue the count.

Use the following syntax to delay the triggering of events:

```
track by_src/by_dst, count count, seconds number_of_seconds
```

The `track` argument specifies whether to use the packet's source or destination IP address when counting the number of packets that meet the rule's detection criteria. Select from the argument values described in the following table to specify how the system tracks event instances.

Table 180: detection_filter Track Arguments

Argument	Description
by_src	Detection criteria count by source IP address.
by_dst	Detection criteria count by destination IP address.

The `count` argument specifies the number of packets that must trigger the rule for the specified IP address within the specified time before the rule generates an event.

The `seconds` argument specifies the number of seconds within which the specified number of packets must trigger the rule before the rule generates an event.

Consider the case of a rule that searches packets for the content `foo` and uses the `detection_filter` keyword with the following arguments:

```
track by_src, count 10, seconds 20
```

In the example, the rule will not generate an event until it has detected `foo` in 10 packets within 20 seconds from a given source IP address. If the system detects only 7 packets containing `foo` within the first 20 seconds, no event is generated. However, if `foo` occurs 40 times in the first 20 seconds, the rule generates 30 events and the count begins again when 20 seconds have elapsed.

Comparing the threshold and detection_filter Keywords

The `detection_filter` keyword replaces the deprecated `threshold` keyword. The `threshold` keyword is still supported for backward compatibility and operates the same as thresholds that you set within an intrusion policy.

The `detection_filter` keyword is a detection feature that is applied before a packet triggers a rule. The rule does not generate an event for triggering packets detected before the specified packet count and, in an inline deployment, does not drop those packets if the rule is set to drop packets. Conversely, the rule does generate events for packets that trigger the rule and occur after the specified packet count and, in an inline deployment, drops those packets if the rule is set to drop packets.

Thresholding is an event notification feature that does not result in a detection action. It is applied after a packet triggers an event. In an inline deployment, a rule that is set to drop packets drops all packets that trigger the rule, independent of the rule threshold.

Note that you can use the `detection_filter` keyword in any combination with the intrusion event thresholding, intrusion event suppression, and rate-based attack prevention features in an intrusion policy. Note also that

policy validation fails if you enable an imported local rule that uses the deprecated `threshold` keyword in combination with the intrusion event thresholding feature in an intrusion policy.

Related Topics

Intrusion Event Thresholds, on page 1505
Intrusion Policy Suppression Configuration, on page 1509
Setting a Dynamic Rule State from the Rules Page, on page 1512
Best Practices for Importing Local Intrusion Rules, on page 147

The tag Keyword

Use the `tag` keyword to tell the system to log additional traffic for the host or session. Use the following syntax when specifying the type and amount of traffic you want to capture using the `tag` keyword:

```
tagging_type, count, metric, optional_direction
```

The next three tables describe the other available arguments.

You can choose from two types of tagging. The following table describes the two types of tagging. Note that the session tag argument type causes the system to log packets from the same session as if they came from different sessions if you configure only rule header options in the intrusion rule. To group packets from the same session together, configure one or more rule options (such as a `flag` keyword or `content` keyword) within the same intrusion rule.

Table 181: Tag Arguments

Argument	Description
session	Logs packets in the session that triggered the rule.
host	Logs packets from the host that sent the packet that triggered the rule. You can add a directional modifier to log only the traffic coming from the host (`src`) or going to the host (`dst`).

To indicate how much traffic you want to log, use the following argument:

Table 182: Count Argument

Argument	Description
count	The number of packets or seconds you want to log after the rule triggers. This unit of measure is specified with the metric argument, which follows the count argument.

Select the metric you want to use to log by time or volume of traffic from those described in the following table.



Caution High-bandwidth networks can see thousands of packets per second, and tagging a large number of packets may seriously affect performance, so make sure you tune this setting for your network environment.

Table 183: Logging Metrics Arguments

Argument	Description
packets	Logs the number of packets specified by the count after the rule triggers.
seconds	Logs traffic for the number of seconds specified by the count after the rule triggers.

For example, when a rule with the following `tag` keyword value triggers:

```
host, 30, seconds, dst
```

all packets that are transmitted from the client to the host for the next 30 seconds are logged.

The flowbits Keyword

Use the `flowbits` keyword to assign state names to sessions. By analyzing subsequent packets in a session according to the previously named state, the system can detect and alert on exploits that span multiple packets in a single session.

The `flowbits` state name is a user-defined label assigned to packets in a specific part of a session. You can label packets with state names based on packet content to help distinguish malicious packets from those you do not want to alert on. You can define up to 1024 state names per managed device. For example, if you want to alert on malicious packets that you know only occur after a successful login, you can use the `flowbits` keyword to filter out the packets that constitute an initial login attempt so you can focus only on the malicious packets. You can do this by first creating a rule that labels all packets in the session that have an established login with a `logged_in` state, then creating a second rule where `flowbits` checks for packets with the state you set in the first rule and acts only on those packets.

An optional *group name* allows you to include a state name in a group of states. A state name can belong to several groups. States not associated with a group are not mutually exclusive, so a rule that triggers and sets a state that is not associated with a group does not affect other currently set states.

flowbits Keyword Options

The following table describes the various combinations of operators, states, and groups available to the `flowbits` keyword. Note that state names can contain alphanumeric characters, periods (.), underscores (_), and dashes (-).

Table 184: flowbits Options

Operator	State Option	Group	Description
`set`	`state_name`	optional	Sets the specified state for a packet. Sets the state in the specified group if a group is defined.
`set`	`state_name&state_name`	optional	Sets the specified states for a packet. Sets the states in the specified group if a group is defined.
`setx`	`state_name`	mandatory	Sets the specified state in the specified group for a packet, and unsets all other states in the group.
`setx`	`state_name&state_name`	mandatory	Sets the specified states in the specified group for a packet, and unsets all other states in the group.

Operator	State Option	Group	Description
`unset`	`state_name`	no group	Unsets the specified state for a packet.
`unset`	`state_name&state_name`	no group	Unsets the specified states for a packet.
`unset`	`all`	mandatory	Unsets all the states in the specified group.
`toggle`	`state_name`	no group	Unsets the specified state if it is set, and sets the specified state if it is unset.
`toggle`	`state_name&state_name`	no group	Unsets the specified states if they are set, and sets the specified states if they are unset.
`toggle`	`all`	mandatory	Unsets all states set in the specified group, and sets all states unset in the specified group.
`isset`	`state_name`	no group	Determines if the specified state is set in the packet.
`isset`	`state_name&state_name`	no group	Determines if the specified states are set in the packet.
`isset`	`state_name|state_name`	no group	Determines if any of the specified states are set in the packet.
`isset`	`any`	mandatory	Determines if any state is set in the specified group.
`isset`	`all`	mandatory	Determines if all states are set in the specified group.
`isnotset`	`state_name`	no group	Determines if the specified state is not set in the packet.
`isnotset`	`state_name&state_name`	no group	Determines if the specified states are not set in the packet.
`isnotset`	`state_name|state_name`	no group	Determines if any of the specified states is not set in the packet.
`isnotset`	`any`	mandatory	Determines if any state is not set in the packet.
`isnotset`	`all`	mandatory	Determines if all states are not set in the packet.
`reset`	(no state)	optional	Unsets all states for all packets. Unsets all states in a group if a group is specified.
`noalert`	(no state)	no group	Use this in conjunction with any other operator to suppress event generation.

Guidelines for Using the flowbits Keyword

Note the following when using the `flowbits` keyword:

- When using the `setx` operator, the specified state can only belong to the specified group, and not to any other group.
- You can define the `setx` operator multiple times, specifying different states and the same group with each instance.
- When you use the `setx` operator and specify a group, you cannot use the `set`, `toggle`, or `unset` operators on that specified group.
- The `isset` and `isnotset` operators evaluate for the specified state regardless of whether the state is in a group.
- During intrusion policy saves, intrusion policy reapplies, and access control policy applies (regardless of whether the access control policy references one intrusion policy or multiple intrusion policies), if you enable a rule that contains the `isset` or `isnotset` operator **without** a specified group, and you do not enable at least one rule that affects `flowbits` assignment (`set`, `setx`, `unset`, `toggle`) for the corresponding state name and protocol, all rules that affect `flowbits` assignment for the corresponding state name are enabled.
- During intrusion policy saves, intrusion policy reapplies, and access control policy applies (regardless of whether the access control policy references one intrusion policy or multiple intrusion policies), if you enable a rule that contains the `isset` or `isnotset` operator **with** a specified group, all rules that affect `flowbits` assignment (`set`, `setx`, `unset`, `toggle`) and define a corresponding group name are also enabled.

flowbits Keyword Examples

This section provides three examples that use the `flowbits` keyword.

flowbits Keyword Example: A Configuration Using state_name

This is an example of a flowbits configuration using `state_name`.

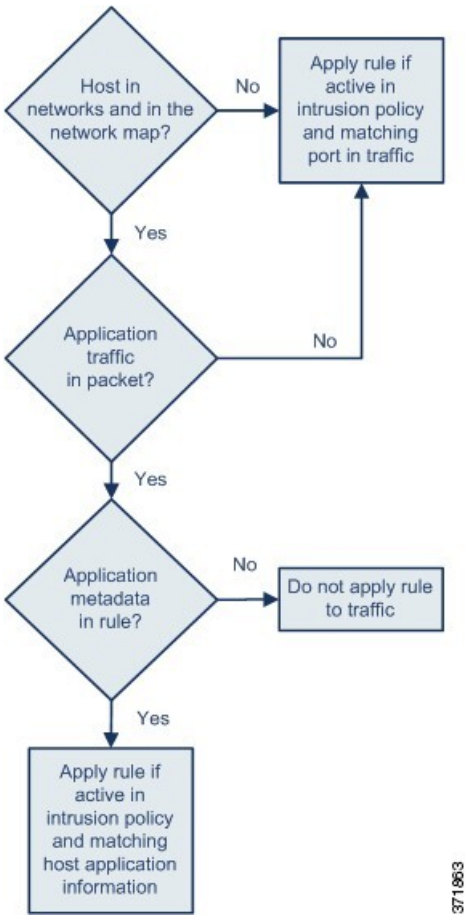
Consider the IMAP vulnerability described in CVE ID 2000-0284. This vulnerability exists in an implementation of IMAP, specifically in the LIST, LSUB, RENAME, FIND, and COPY commands. However, to take advantage of the vulnerability, the attacker must be logged into the IMAP server. Because the LOGIN confirmation from the IMAP server and the exploit that follows are necessarily in different packets, it is difficult to construct non-flow-based rules that catch this exploit. Using the `flowbits` keyword, you can cflowbitonstruct a series of rules that track whether the user is logged into the IMAP server and, if so, generate an event if one of the attacks is detected. If the user is not logged in, the attack cannot exploit the vulnerability and no event is generated.

The two rule fragments that follow illustrate this example. The first rule fragment looks for an IMAP login confirmation from the IMAP server:

```
alert tcp any 143 -> any any (msg:"IMAP login"; content:"OK

LOGIN"; flowbits:set,logged_in; flowbits:noalert;)
```

The following diagram illustrates the effect of the `flowbits` keyword in the preceding rule fragment:

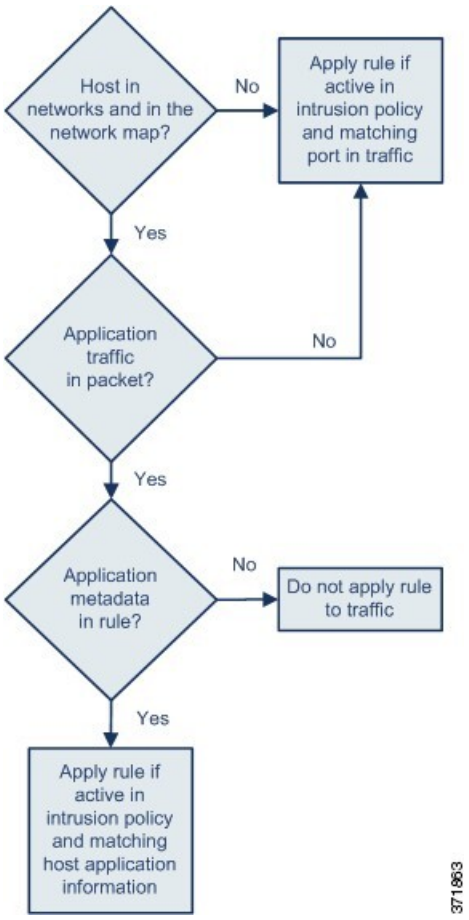


Note that `flowbits:set` sets a state of `logged_in`, while `flowbits:noalert` suppresses the alert because you are likely to see many innocuous login sessions on an IMAP server.

The next rule fragment looks for a LIST string, but does not generate an event unless the logged_in state has been set as a result of some previous packet in the session:

```
alert tcp any any -> any 143 (msg:"IMAP LIST";
content:"LIST"; flowbits:isset,logged_in;)
```

The following diagram illustrates the effect of the `flowbits` keyword in the preceding rule fragment:



In this case, if a previous packet has caused a rule containing the first fragment to trigger, then a rule containing the second fragment triggers and generates an event.

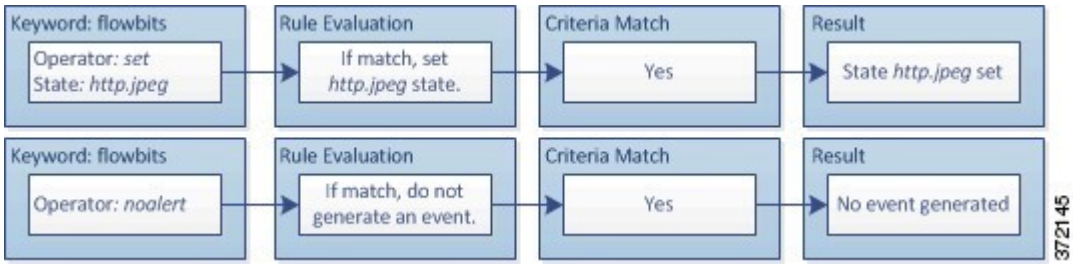
flowbits Keyword Example: A Configuration Resulting in False Positive Events

Including different state names that are set in different rules in a group can prevent false positive events that might otherwise occur when content in a subsequent packet matches a rule whose state is no longer valid. The following example illustrates how you can get false positives when you do not include multiple state names in a group.

Consider the case where the following three rule fragments trigger in the order shown during a single session:

```
(msg:"JPEG transfer";
content:"image/";pcre:"/^Content-?Type\x3a(\s*|\s*\r?\n\s+)image\x2fp?jpe?g/smi";
?flowbits:set,http.jpeg; flowbits:noalert;)
```

The following diagram illustrates the effect of the `flowbits` keyword in the preceding rule fragment:

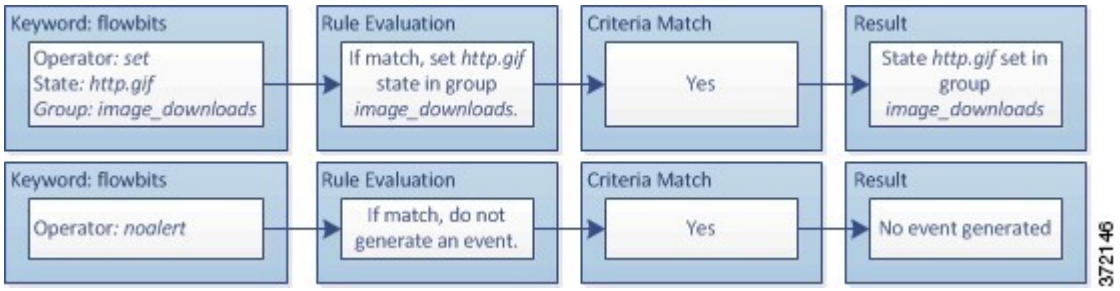


The `content` and `pcre` keywords in the first rule fragment match a JPEG file download, `flowbits:set,http.jpeg` sets the `http.jpeg flowbits` state, and `flowbits:noalert` stops the rule from generating events. No event is generated because the rule's purpose is to detect the file download and set the `flowbits` state so one or more companion rules can test for the state name in combination with malicious content and generate events when malicious content is detected.

The next rule fragment detects a GIF file download subsequent to the JPEG file download above:

```
(msg:"GIF transfer"; content:"image/";
pcre:"/^Content-?Type\x3a(\s*|\s*\r?\n\s+)image\x2fgif/smi";
?flowbits:set,http.jpg,image_downloads; flowbits:noalert;)
```

The following diagram illustrates the effect of the `flowbits` keyword in the preceding rule fragment:

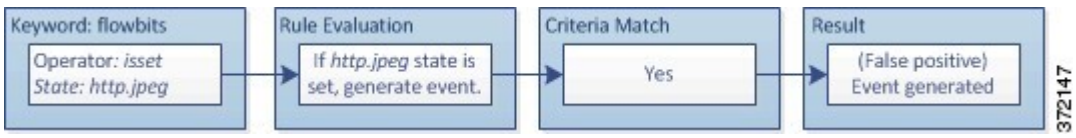


The `content` and `pcre` keywords in the second rule match the GIF file download, `flowbits:set,http.jpg` sets the `http.jpg` flowbit state, and `flowbits:noalert` stops the rule from generating an event. Note that the `http.jpeg` state set by the first rule fragment is still set even though it is no longer needed; this is because the JPEG download must have ended if a subsequent GIF download has been detected.

The third rule fragment is a companion to the first rule fragment:

```
(msg:"JPEG exploit";?flowbits:isset,http.jpeg;content:"|FF|";
pcre:"?/\xFF[\xE1\xE2\xED\xFE]\x00[\x00\x01]/";)
```

The following diagram illustrates the effect of the `flowbits` keyword in the preceding rule fragment:



In the third rule fragment, `flowbits:isset,http.jpeg` determines that the now-irrelevant `http.jpeg` state is set, and `content` and `pcre` match content that would be malicious in a JPEG file but not in a GIF file. The third rule fragment results in a false positive event for a nonexistent exploit in a JPEG file.

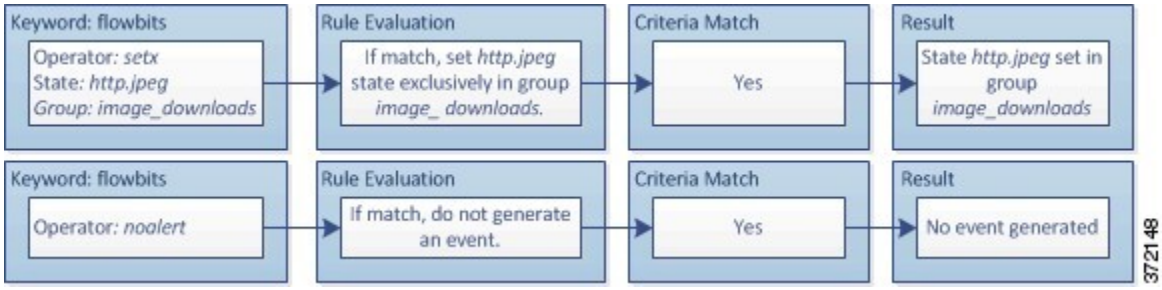
flowbits Keyword Example: A Configuration for Preventing False Positive Events

The following example illustrates how including state names in a group and using the `setx` operator can prevent false positives.

Consider the same case as the previous example, except that the first two rules now include their two different state names in the same state group.

```
(msg:"JPEG transfer";
content:"image/";pcre:"/^Content-?Type\x3a(\s*|\s*\r?\n\s+)image\x2fp?jpe?g/smi";
?flowbits:setx,http.jpeg,image_downloads; flowbits:noalert;)
```

The following diagram illustrates the effect of the `flowbits` keyword in the preceding rule fragment:

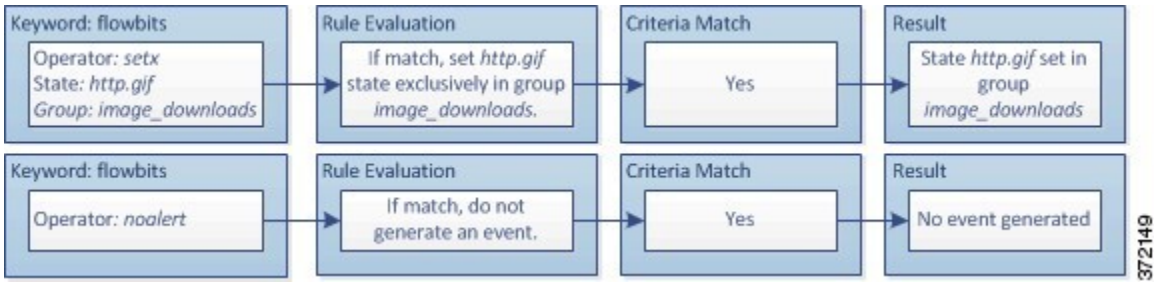


When the first rule fragment detects a JPEG file download, the `flowbits:setx,http.jpeg,image_downloads` keyword sets the `flowbits` state to `http.jpeg` and includes the state in the `image_downloads` group.

The next rule then detects a subsequent GIF file download:

```
(msg:"GIF transfer"; content:"image/";
pcre:"/^Content-?Type\x3a(\s*|\s*\r?\n\s+)image\x2fgif/smi";
?flowbits:setx,http.jpg,image_downloads; flowbits:noalert;)
```

The following diagram illustrates the effect of the `flowbits` keyword in the preceding rule fragment:

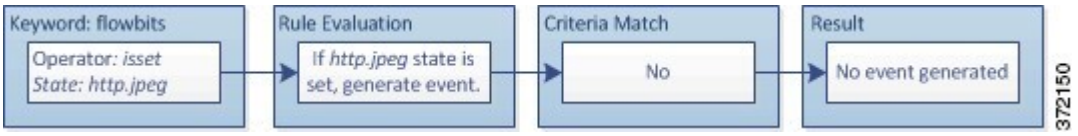


When the second rule fragment matches the GIF download, the `flowbits:setx,http.jpg,image_downloads` keyword sets the `http.jpg flowbits` state and unsets `http.jpeg`, the other state in the group.

The third rule fragment does not result in a false positive:

```
(msg:"JPEG exploit"; ?flowbits:isset,http.jpeg;content:"|FF|";
pcre:"/?\xFF[\xE1\xE2\xED\xFE]\x00[\x00\x01]/";)
```

The following diagram illustrates the effect of the `flowbits` keyword in the preceding rule fragment:



Because `flowbits:isset,http.jpeg` is false, the rules engine stops processing the rule and no event is generated, thus avoiding a false positive even in a case where content in the GIF file matches exploit content for a JPEG file.

The http_encode Keyword

You can use the `http_encode` keyword to generate events on the type of encoding in an HTTP request or response before normalization, either in the HTTP URI, in non-cookie data in an HTTP header, in cookies in HTTP requests headers, or set-cookie data in HTTP responses.

You must configure the HTTP Inspect preprocessor to inspect HTTP responses and HTTP cookies to return matches for rules using the `http_encode` keyword.

Also, you must enable both the decoding and alerting option for each specific encoding type in your HTTP Inspect preprocessor configuration so the `http_encode` keyword in an intrusion rule can trigger events on that encoding type.

The following table describes the encoding types this option can generate events for in HTTP URIs, headers, cookies, and set-cookies:

Table 185: http_encode Encoding Types

Encoding Type	Description
utf8	Detects UTF-8 encoding in the specified location when this encoding type is enabled for decoding by the HTTP Inspect preprocessor.
double_encode	Detects double encoding in the specified location when this encoding type is enabled for decoding by the HTTP Inspect preprocessor.
non_ascii	Detects non-ASCII characters in the specified location when non-ASCII characters are detected but the detected encoding type is not enabled.
uencode	Detects Microsoft %u encoding in the specified location when this encoding type is enabled for decoding by the HTTP Inspect preprocessor.
bare_byte	Detects bare byte encoding in the specified location when this encoding type is enabled for decoding by the HTTP Inspect preprocessor.

Related Topics

The HTTP Inspect Preprocessor, on page 1708
Server-Level HTTP Normalization Options, on page 1709

http_encode Keyword Syntax

Encoding Location

Specifies whether to search for the specified encoding type in an HTTP URI, header, or cookie, including a set-cookie.

Encoding Type

Specifies one or more encoding types using one of the following formats:

```
encode_type
encode_type|encode_type|encode_type...
```

where `encode_type` is one of the following:

```
utf8
double_encode
non_ascii
uencode
bare_byte.
```

Note that you cannot use the negation (!) and OR (|) operators together.

http_encode Keyword example: Using Two http_endcode Keywords to Search for Two Encodings

The following example uses two `http_encode` keywords in the same rule to search the HTTP URI for UTF-8 AND Microsoft IIS %u encoding:

First, the `http_encode` keyword:

- **Encoding Location**: `HTTP URI`
- **Encoding Type**: `utf8`

Then, the additional `http_encode` keyword:

- **Encoding Location**: `HTTP URI`
- **Encoding Type**: `uencode`

Overview: The file_type and file_group Keywords

The `file_type` and `file_group` keywords allow you to detect files transmitted via FTP, HTTP, SMTP, IMAP, POP3, and NetBIOS-ssn (SMB) based on their type and version. Do **not** use more than one `file_type` or `file_group` keyword in a single intrusion rule.



Tip Updating your vulnerability database (VDB) populates the intrusion rules editor with the most up-to-date file types, versions, and groups.

Note The system does not automatically enable preprocessors to accomodate the `file_type` and `file_group` keywords.

You **must** enable specific preprocessors if you want to generate events and, in an inline deployment, drop offending packets for traffic matching your `file_type` or `file_group` keywords.

Table 186: file_type and file_group Intrusion Event Generation

Protocol	Required Preprocessor or Preprocessor Option
FTP	FTP/Telnet preprocessor and the **Normalize TCP Payload** inline normalization preprocessor option
HTTP	HTTP Inspect preprocessor to generate intrusion events in HTTP traffic
SMTP	SMTP preprocessor to generate intrusion events in HTTP traffic
IMAP	IMAP preprocessor
POP3	POP preprocessor
Netbios-ssn (SMB)	The DCE/RPC preprocessor and the **SMB File Inspection** DCE/RPC preprocessor option

Related Topics

The FTP/Telnet Decoder, on page 1700
The Inline Normalization Preprocessor, on page 1762
The HTTP Inspect Preprocessor, on page 1708
The SMTP Preprocessor, on page 1736
The IMAP Preprocessor, on page 1730
The POP Preprocessor, on page 1733
The DCE/RPC Preprocessor, on page 1686

The file_type and file_group Keywords

file_type

The `file_type` keyword allows you to specify the file type and version of a file detected in traffic. File type arguments (for example, **JPEG** and **PDF**) identify the format of the file you want to find in traffic.

Note Do **not** use the `file_type` keyword with another `file_type` or `file_group` keyword in the same intrusion rule.

The system selects **Any Version** by default, but some file types allow you to select version options (for example, PDF version **1.7**) to identify specific file type versions you want to find in traffic.

file_group

The `file_group` keyword allows you to select a Cisco-defined group of similar file types to find in traffic (for example, **multimedia** or **audio**). File groups also include Cisco-defined versions for each file type in the group.



Note Do **not** use the `file_group` keyword with another `file_group` or `file_type` keyword in the same intrusion rule.

The file_data Keyword

The `file_data` keyword provides a pointer that serves as a reference for the positional arguments available for other keywords such as `content`, `byte_jump`, `byte_test`, and `pcre`. The detected traffic determines the type of data the `file_data` keyword points to. You can use the `file_data` keyword to point to the beginning of the following payload types:

- HTTP response body

  To inspect HTTP response packets, the HTTP Inspect preprocessor must be enabled and you must configure the preprocessor to inspect HTTP responses. The `file_data` keyword matches if the HTTP Inspect preprocessor detects HTTP response body data.

- Uncompressed gzip file data

  To inspect uncompressed gzip files in the HTTP response body, the HTTP Inspect preprocessor must be enabled and you must configure the preprocessor to inspect HTTP responses and to decompress gzip-compressed files in the HTTP response body. For more information, see the **Inspect HTTP Responses** and **Inspect Compressed Data** Server-Level HTTP Normalization options. The `file_data` keyword matches if the HTTP Inspect preprocessor detects uncompressed gzip data in the HTTP response body.

- Normalized JavaScript

  To inspect normalized JavaScript data, the HTTP Inspect preprocessor must be enabled and you must configure the preprocessor to inspect HTTP responses. The `file_data` keyword matches if the HTTP Inspect preprocessor detects JavaScript in response body data.

- SMTP payload

  To inspect the SMTP payload, the SMTP preprocessor must be enabled. The `file_data` keyword matches if the SMTP preprocessor detects SMTP data.

- Encoded email attachments in SMTP, POP, or IMAP traffic

  To inspect email attachments in SMTP, POP, or IMAP traffic, the SMTP, POP, or IMAP preprocessor, respectively, must be enabled, alone or in any combination. Then, for each enabled preprocessor, you must ensure that the preprocessor is configured to decode each attachment encoding type that you want decoded. The attachment decoding options that you can configure for each preprocessor are: **Base64 Decoding Depth**, **7-Bit/8-Bit/Binary Decoding Depth**, **Quoted-Printable Decoding Depth**, and **Unix-to-Unix Decoding Depth**.

  You can use multiple `file_data` keywords in a rule.

Related Topics

The HTTP Inspect Preprocessor, on page 1708
Server-Level HTTP Normalization Options, on page 1709
The SMTP Preprocessor, on page 1736
The IMAP Preprocessor, on page 1730

The pkt_data Keyword

The `pkt_data` keyword provides a pointer that serves as a reference for the positional arguments available for other keywords such as `content`, `byte_jump`, `byte_test`, and `pcre`.

When normalized FTP, telnet, or SMTP traffic is detected, the `pkt_data` keyword points to the beginning of the normalized packet payload. When other traffic is detected, the `pkt_data` keyword points to the beginning of the raw TCP or UDP payload.

The following normalization options must be enabled for the system to normalize the corresponding traffic for inspection by intrusion rules:

- Enable the FTP & Telnet preprocessor **Detect Telnet Escape codes within FTP commands** option to normalize FTP traffic for inspection.
- Enable the FTP & Telnet preprocessor **Normalize** telnet option to normalize telnet traffic for inspection.
- Enable the SMTP preprocessor **Normalize** option to normalize SMTP traffic for inspection.

You can use multiple `pkt_data` keywords in a rule.

Related Topics

Client-Level FTP Options, on page 1705
Telnet Options, on page 1701
SMTP Preprocessor Options, on page 1736

The base64_decode and base64_data Keywords

You can use the `base64_decode` and `base64_data` keywords in combination to instruct the rules engine to decode and inspect specified data as Base64 data. This can be useful, for example, for inspecting Base64-encoded HTTP Authentication request headers and Base64-encoded data in HTTP PUT and POST requests.

These keywords are particularly useful for decoding and inspecting Base64 data in HTTP requests. However, you can also use them with any protocol such as SMTP that uses the space and tab characters the same way HTTP uses these characters to extend a lengthy header line over multiple lines. When this line extension, which is known as folding, is not present in a protocol that uses it, inspection ends at any carriage return or line feed that is not followed with a space or tab.

base64_decode

The `base64_decode` keyword instructs the rules engine to decode packet data as Base64 data. Optional arguments let you specify the number of bytes to decode and where in the data to begin decoding.

You can use the `base64_decode` keyword once in a rule; it must precede at least one instance of the `base64_data` keyword.

Before decoding Base64 data, the rules engine unfolds lengthy headers that are folded across multiple lines. Decoding ends when the rules engine encounters any the following:

- the end of a header line
- the specified number of bytes to decode
- the end of the packet

The following table describes the arguments you can use with the `base64_decode` keyword.

Table 187: Optional base64_decode Arguments

Argument	Description
Bytes	Specifies the number of bytes to decode. When not specified, decoding continues to the end of a header line or the end of the packet payload, whichever comes first. You can specify a positive, non-zero value.
Offset	Determines the offset relative to the start of the packet payload or, when you also specify **Relative**, relative to the current inspection location. You can specify a positive, non-zero value.
Relative	Specifies inspection relative to the current inspection location.

base64_data

The `base64_data` keyword provides a reference for inspecting Base64 data decoded using the `base64_decode` keyword. The `base64_data` keyword sets inspection to begin at the start of the decoded Base64 data. Optionally, you can then use the positional arguments available for other keywords such as `content` or `byte_test` to further specify the location to inspect.

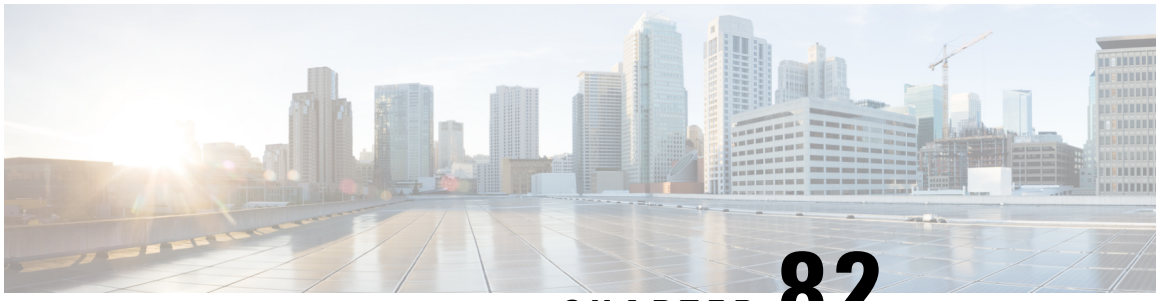
You must use the `base64_data` keyword at least once after using the `base64_decode` keyword; optionally, you can use `base64_data` multiple times to return to the beginning of the decoded Base64 data.

Note the following when inspecting Base64 data:

- You cannot use the fast pattern matcher.
- If you interrupt Base64 inspection in a rule with an intervening HTTP content argument, you must insert another `base64_data` keyword in the rule before further inspecting Base64 data.

Related Topics

Overview: HTTP content and protected_content Keyword Arguments, on page 1571
content Keyword Fast Pattern Matcher Arguments, on page 1575



CHAPTER 82

Intrusion Prevention Performance Tuning

The following topics describe how to refine intrusion prevention performance:

- About Intrusion Prevention Performance Tuning, on page 1655
- License Requirements for Intrusion Prevention Performance Tuning, on page 1656
- Requirements and Prerequisites for Intrusion Prevention Performance Tuning, on page 1656
- Limiting Pattern Matching for Intrusions, on page 1656
- Regular Expression Limits Overrides for Intrusion Rules, on page 1657
- Overriding Regular Expression Limits for Intrusion Rules, on page 1658
- Per Packet Intrusion Event Generation Limits, on page 1658
- Limiting Intrusion Events Generated Per Packet, on page 1659
- Packet and Intrusion Rule Latency Threshold Configuration, on page 1659
- Intrusion Performance Statistic Logging Configuration, on page 1665
- Configuring Intrusion Performance Statistic Logging, on page 1666

About Intrusion Prevention Performance Tuning

Cisco provides several features for improving the performance of your system as it analyzes traffic for attempted intrusions. You can:

- specify the number of packets to allow in the event queue. You can also, before and after stream reassembly, enable or disable inspection of packets that will be rebuilt into larger streams.
- override default match and recursion limits on PCRE that are used in intrusion rules to examine packet payload content.
- elect to have the rules engine log more than one event per packet or packet stream when multiple events are generated, allowing you to collect information beyond the reported event.
- balance security with the need to maintain device latency at an acceptable level with packet and rule latency thresholding.
- configure the basic parameters of how devices monitor and report their own performance. This allows you to specify the intervals at which the system updates performance statistics on your devices.

You configure these performance settings on a per-access-control-policy basis, and they apply to all intrusion policies invoked by that parent access control policy.

License Requirements for Intrusion Prevention Performance Tuning

FTD License

Threat

Classic License

Protection

Requirements and Prerequisites for Intrusion Prevention Performance Tuning

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Limiting Pattern Matching for Intrusions

Procedure

Step 1 In the access control policy editor, click **Advanced**.

Step 2 Click **Edit** () next to **Performance Settings**.

If **View** () appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings.If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Click **Pattern Matching Limits** in the **Performance Settings** pop-up window.

Step 4 Enter a value for the maximum number of events to queue in the **Maximum Pattern States to Analyze Per Packet** field.

Step 5 To disable the inspection of packets that will be rebuilt into larger streams of data before and after stream reassembly, check the **Disable Content Checks on Traffic Subject to Future Reassembly** check box. Inspection before and after reassembly requires more processing overhead and may decrease performance.

Step 6 Click **OK**.

Step 7 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Regular Expression Limits Overrides for Intrusion Rules

The default regular expression limits ensure a minimum level of performance. Overriding these limits could increase security, but could also significantly impact performance by permitting packet evaluation against inefficient regular expressions.

Caution Do not override default PCRE limits unless you are an experienced intrusion rule writer with knowledge of the impact of degenerative patterns.

Table 188: Regular Expression Constraint Options

Option	Description
Match Limit State	Specifies whether to override **Match Limit**. You have the following options: • select **Default** to use the value configured for **Match Limit** • select **Unlimited** to permit an unlimited number of attempts • select **Custom** to specify either a limit of 1 or greater for **Match Limit**, or to specify 0 to completely disable PCRE match evaluations
Match Limit	Specifies the number of times to attempt to match a pattern defined in a PCRE regular expression.
Match Recursion Limit State	Specifies whether to override **Match Recursion Limit**. You have the following options: • select **Default** to use the value configured for **Match Recursion Limit** • select **Unlimited** to permit an unlimited number of recursions • select **Custom** to specify either a limit of 1 or greater for **Match Recursion Limit**, or to specify 0 to completely disable PCRE recursions Note that for **Match Recursion Limit** to be meaningful, it must be smaller than **Match Limit**.

Option	Description
Match Recursion Limit	Specifies the number of recursions when evaluating a PCRE regular expression against the packet payload.

Related Topics

Overview: The pcre Keyword, on page 1587

Overriding Regular Expression Limits for Intrusion Rules

Procedure

Step 1 In the access control policy editor, click **Advanced**.

Step 2 Click **Edit** ( ) next to **Performance Settings**.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings.If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Click **Regular Expression Limits** in the **Performance Settings** pop-up window.

Step 4 You can modify any of the options as described in Regular Expression Limits Overrides for Intrusion Rules, on page 1657.

Step 5 Click **OK**.

Step 6 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Per Packet Intrusion Event Generation Limits

When the intrusion rules engine evaluates traffic against rules, it places the events generated for a given packet or packet stream in an event queue, then reports the top events in the queue to the user interface. When configuring the intrusion event logging limits, you can specify how many events can be placed in the queue and how many are logged, and select the criteria for determining event order within the queue.

Table 189: Intrusion Event Logging Limits Options

Option	Description
Maximum Events Stored Per Packet	The maximum number of events that can be stored for a given packet or packet stream.
Maximum Events Logged Per Packet	The number of events logged for a given packet or packet stream. This cannot exceed the **Maximum Events Stored Per Packet** value.

Option	Description
Prioritize Event Logging By	The value used to determine event ordering within the event queue. The highest ordered event is reported through the user interface. You can select from: • `priority`, which orders events in the queue by the event priority. • `content_length`, which orders events by the longest identified content match. When events are ordered by content length, rule events always take precedence over decoder and preprocessor events.

Limiting Intrusion Events Generated Per Packet

Procedure

Step 1 In the access control policy editor, click **Advanced**.

Step 2 Click **Edit** ( ) next to **Performance Settings**.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings.If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Click **Intrusion Event Logging Limits** in the **Performance Settings** pop-up window.

Step 4 You can modify any of the options in Per Packet Intrusion Event Generation Limits, on page 1658.

Step 5 Click **OK**.

Step 6 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Packet and Intrusion Rule Latency Threshold Configuration

Each access control policy has latency-based settings that use thresholding to manage packet and rule processing performance.

Packet latency thresholding measures the total elapsed time taken to process a packet by applicable decoders, preprocessors, and rules, and ceases inspection of the packet if the processing time exceeds a configurable threshold.

Rule latency thresholding measures the elapsed time each rule takes to process an individual packet, suspends the violating rule along with a group of related rules for a specified time if the processing time exceeds the rule latency threshold a configurable consecutive number of times, and restores the rules when the suspension expires.

Latency-Based Performance Settings

By default, the system takes latency-based performance settings from the latest intrusion rule update deployed on your system.

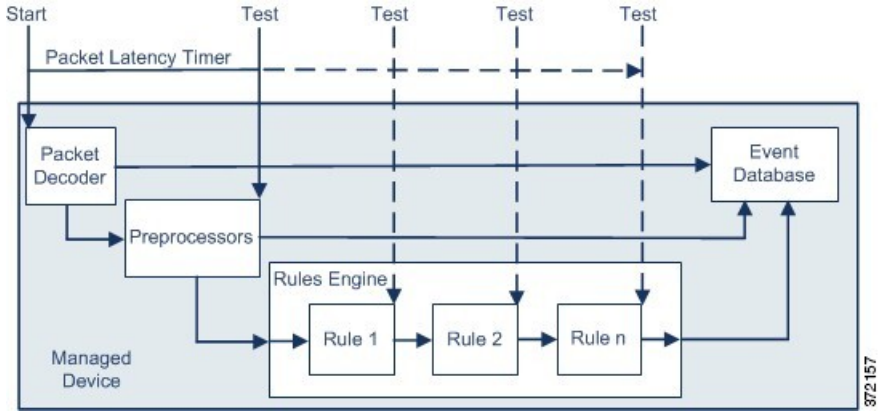
The latency settings that are actually applied depend on the security level of the network analysis policy (NAP) associated with the access control policy. Generally, this is the default NAP policy. However, if custom network analysis rules are configured, and if any of these specify a NAP policy that is more secure than the default NAP policy, then latency settings are based on the most secure NAP policy among the custom rules. If the default NAP policy or any custom rules invoke a custom NAP policy, then the security level used in the evaluation is the system-provided base policy on which each custom NAP policy is based.

The above is true regardless of whether the effective threshold and/or network analysis configurations are inherited or configured directly in the policy.

Packet Latency Thresholding

Packet latency thresholding measures elapsed time, not just processing time, in order to more accurately reflect the actual time required for the rule to process a packet. However, latency thresholding is a software-based latency implementation that does not enforce strict timing.

The trade-off for the performance and latency benefits derived from latency thresholding is that uninspected packets could contain attacks. A timer starts for each packet when decoder processing begins. Timing continues either until all processing ends for the packet or until the processing time exceeds the threshold at a timing test point.



As illustrated in the above figure, packet latency timing is tested at the following test points:

- after the completion of all decoder and preprocessor processing and before rule processing begins
- after processing by each rule

If the processing time exceeds the threshold at any test point, packet inspection ceases.

Tip Total packet processing time does not include routine TCP stream or IP fragment reassembly times.

Packet latency thresholding has no effect on events triggered by a decoder, preprocessor, or rule processing the packet. Any applicable decoder, preprocessor, or rule triggers normally until a packet is fully processed,

or until packet processing ends because the latency threshold is exceeded, whichever comes first. If a drop rule detects an intrusion in an inline deployment, the drop rule triggers an event and the packet is dropped.

Note No packets are evaluated against rules after processing for that packet ceases because of a packet latency threshold violation. A rule that would have triggered an event cannot trigger that event, and for drop rules, cannot drop the packet.

Packet latency thresholding can improve system performance in both passive and inline deployments, and can reduce latency in inline deployments, by stopping inspection of packets that require excessive processing time. These performance benefits might occur when, for example:

- for both passive and inline deployments, sequential inspection of a packet by multiple rules requires an excessive amount of time
- for inline deployments, a period of poor network performance, such as when someone downloads an extremely large file, slows packet processing

In a passive deployment, stopping the processing of packets might not contribute to restoring network performance because processing simply moves to the next packet.

Packet Latency Thresholding Notes

By default, latency-based performance settings for both packet and rule handling are automatically populated by the latest deployed intrusion rule update, and Cisco recommends that you do not change the default.

The information in this topic applies only if you choose to specify custom values.

Table 190: Packet Latency Thresholding Option

Option	Description
Threshold (microseconds)	Specifies the time, in microseconds, when inspection of a packet ceases.

You can enable rule 134:3 to generate an event and, in an inline deployment, drop offending packets when the system stops inspecting a packet because the packet latency threshold is exceeded. For more information, see Intrusion Rule State Options, on page 1503.

Configuring Packet Latency Thresholding

Note By default, latency-based performance settings for both packet and rule handling are automatically populated by the latest deployed intrusion rule update, and Cisco recommends that you do not change the default.

Procedure

Step 1 In the access control policy editor, click **Advanced**.

Step 2 Click **Edit** ( ) next to **Latency-Based Performance Settings**.

System > **Monitoring** > **Statistics**

Step 3 If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 4 Click **Packet Handling** in the **Latency-Based Performance Settings** pop-up window.

By default, **Installed Rule Update** is selected. Cisco recommends using this default.

The values displayed do not reflect the automated settings.

Step 5 If you choose to specify custom values:

- See Packet Latency Thresholding Notes, on page 1661 for recommended minimum **Threshold** settings.
- You must specify custom values in both the packet handling tab and the rule handling tab.

Step 6 Click **OK**.

Step 7 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Rule Latency Thresholding

Rule latency thresholding measures elapsed time, not just processing time, in order to more accurately reflect the actual time required for the rule to process a packet. However, latency thresholding is a software-based latency implementation that does not enforce strict timing.

The trade-off for the performance and latency benefits derived from latency thresholding is that uninspected packets could contain attacks. A timer measures the processing time each time a packet is processed against a group of rules. Any time the rule processing time exceeds a specified rule latency threshold, the system increments a counter. If the number of consecutive threshold violations reaches a specified number, the system takes the following actions:

- suspends the rules for the specified period
- triggers an event indicating the rules have been suspended
- re-enables the rules when the suspension expires
- triggers an event indicating the rules have been re-enabled

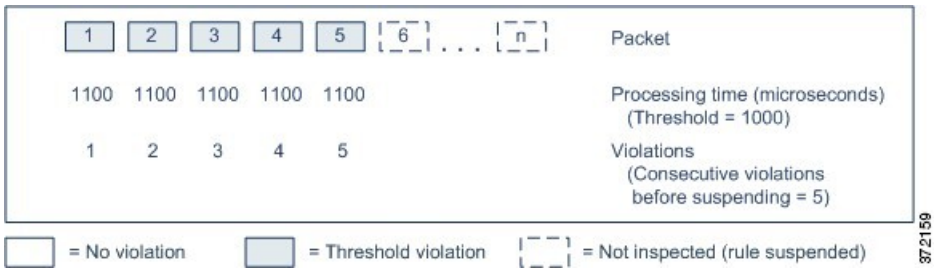
The system zeroes the counter when the group of rules has been suspended, or when rule violations are not consecutive. Permitting some consecutive violations before suspending rules lets you ignore occasional rule violations that might have negligible impact on performance and focus instead on the more significant impact of rules that repeatedly exceed the rule latency threshold.

The following example shows five consecutive rule processing times that do not result in rule suspension.



In the above example, the time required to process each of the first three packets violates the rule latency threshold of 1000 microseconds, and the violations counter increments with each violation. Processing of the fourth packet does not violate the threshold, and the violations counter resets to zero. The fifth packet violates the threshold and the violations counter restarts at one.

The following example shows five consecutive rule processing times that do result in rule suspension.



In the second example, the time required to process each of the five packets violates the rule latency threshold of 1000 microseconds. The group of rules is suspended because the rule processing time of 1100 microseconds for each packet violates the threshold of 1000 microseconds for the specified five consecutive violations. Any subsequent packets, represented in the figure as packets 6 through n, are not examined against suspended rules until the suspension expires. If more packets occur after the rules are re-enabled, the violations counter begins again at zero.

Rule latency thresholding has no effect on intrusion events triggered by the rules processing the packet. A rule triggers an event for any intrusion detected in the packet, regardless of whether the rule processing time exceeds the threshold. If the rule detecting the intrusion is a drop rule in an inline deployment, the packet is dropped. When a drop rule detects an intrusion in a packet that results in the rule being suspended, the drop rule triggers an intrusion event, the packet is dropped, and that rule and all related rules are suspended.

Note Packets are not evaluated against suspended rules. A suspended rule that would have triggered an event cannot trigger that event and, for drop rules, cannot drop the packet.

Rule latency thresholding can improve system performance in both passive and inline deployments, and can reduce latency in inline deployments, by suspending rules that take the most time to process packets. Packets are not evaluated again against suspended rules until a configurable time expires, giving the overloaded device time to recover. These performance benefits might occur when, for example:

- hastily written, largely untested rules require an excessive amount of processing time
- a period of poor network performance, such as when someone downloads an extremely large file, causes slow packet inspection

Rule Latency Thresholding Notes

By default, latency-based performance settings for both packet and rule handling are automatically populated by the latest deployed intrusion rule update, and Cisco recommends that you do not change the default.

The information in this topic applies only if you choose to specify custom values.

Rule latency thresholding suspends rules for the time specified by **Suspension Time** when the time rules take to process a packet exceeds **Threshold** for the consecutive number of times specified by **Consecutive Threshold Violations Before Suspending Rule**.

You can enable rule 134:1 to generate an event when rules are suspended, and rule 134:2 to generate an event when suspended rules are enabled. See Intrusion Rule State Options, on page 1503.

Table 191: Rule Latency Thresholding Options

Option	Description
Threshold	Specifies the time in microseconds that rules should not exceed when examining a packet.
Consecutive Threshold Violations Before Suspending Rule	Specifies the consecutive number of times rules can take longer than the time set for **Threshold** to inspect packets before rules are suspended.
Suspension Time	Specifies the number of seconds to suspend a group of rules.

Configuring Rule Latency Thresholding

Note By default, latency-based performance settings for both packet and rule handling are automatically populated by the latest deployed intrusion rule update, and Cisco recommends that you do not change the default.

Procedure

Step 1 In the access control policy editor, click **Advanced**.

Step 2 Click **Edit** ( ) next to **Latency-Based Performance Settings**.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Click **Rule Handling** in the **Latency-Based Performance Settings** pop-up window.

By default, **Installed Rule Update** is selected. Cisco recommends using this default.

The values displayed do not reflect the automated settings.

Step 4 If you choose to specify custom values:

- You can configure any of the options in Rule Latency Thresholding Notes, on page 1664.
- You must specify custom values in both the packet handling tab and the rule handling tab.

Step 5 Click **OK**.

Step 6 Click **Save** to save the policy.

What to do next

- If you want to generate events, enable latency rules 134:1 and 134:2. For more information, see Intrusion Rule State Options, on page 1503.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Intrusion Performance Statistic Logging Configuration

Sample time (seconds) and Minimum number of packets

When the number of seconds specified elapses between performance statistics updates, the system verifies it has analyzed the specified number of packets. If it has, the system updates performance statistics. Otherwise, the system waits until it analyzes the specified number of packets.

Troubleshooting Options: Log Session/Protocol Distribution

Support might ask you during a troubleshooting call to log protocol distribution, packet length, and port statistics.

Caution Do not enable **Log Session/Protocol Distribution** unless instructed to by Support. Note that enabling or disabling **Log Session/Protocol Distribution** restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Troubleshooting Options: Summary

Support might ask you during a troubleshooting call to configure the system to calculate the performance statistics only when the Snort process is shut down or restarted. To enable this option, you must also enable the **Log Session/Protocol Distribution** troubleshooting option.

Caution Do not enable **Summary** unless instructed to do so by Support.

Configuring Intrusion Performance Statistic Logging

Procedure

Step 1 In the access control policy editor, click **Advanced**, then click **Edit** ( ) next to **Performance Settings**.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings.If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 Click **Performance Statistics** in the pop-up window that appears.

Step 3 Modify the **Sample time** or **Minimum number of packets** as described above.

Step 4 Optionally, expand the **Troubleshoot Options** section and modify those options only if asked to do so by Support.

Caution Enabling or disabling **Log Session/Protocol Distribution** restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Step 5 Click **OK**.

What to do next

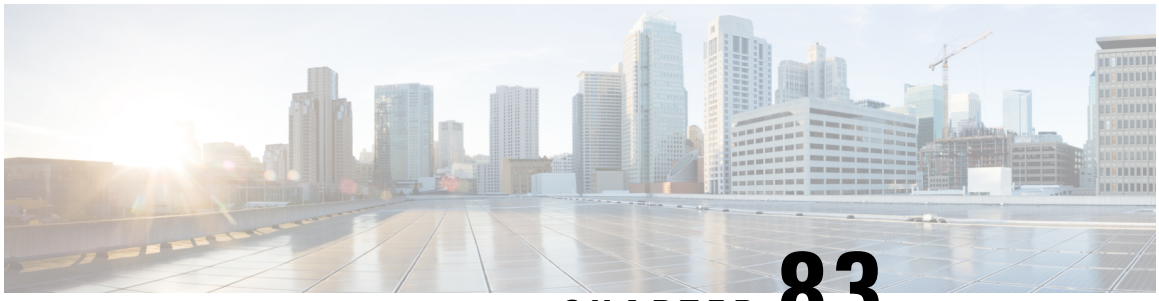
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.



PART XIX

Advanced Network Analysis and Preprocessing

- Advanced Access Control Settings for Network Analysis and Intrusion Policies, on page 1669
- Getting Started with Network Analysis Policies, on page 1677
- Application Layer Preprocessors, on page 1685
- SCADA Preprocessors, on page 1751
- Transport & Network Layer Preprocessors, on page 1757
- Detecting Specific Threats, on page 1791
- Adaptive Profiles, on page 1811



CHAPTER 83

Advanced Access Control Settings for Network Analysis and Intrusion Policies

The following topics describe how to configure advanced settings for network analysis and intrusion policies:

- About Advanced Access Control Settings for Network Analysis and Intrusion Policies, on page 1669
- Requirements and Prerequisites for Advanced Access Control Settings for Network Analysis and Intrusion Policies, on page 1669
- Inspection of Packets That Pass Before Traffic Is Identified, on page 1670
- Advanced Settings for Network Analysis Policies, on page 1672

About Advanced Access Control Settings for Network Analysis and Intrusion Policies

Many of the advanced settings in an access control policy govern intrusion detection and prevention configurations that require specific expertise to configure. Advanced settings typically require little or no modification and are not common to every deployment.

Requirements and Prerequisites for Advanced Access Control Settings for Network Analysis and Intrusion Policies

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Access Admin

- Network Admin

Inspection of Packets That Pass Before Traffic Is Identified

For some features, including URL filtering, application detection, rate limiting, and Intelligent Application Bypass, a few packets must pass in order for the connection to be established, and to enable the system to identify the traffic and determine which access control rule (if any) will handle that traffic.

You must explicitly configure your access control policy to inspect these packets, prevent them from reaching their destination, and generate any events. See Specify a Policy to Handle Packets That Pass Before Traffic Identification, on page 1670.

As soon as the system identifies the access control rule or default action that should handle the connection, the remaining packets in the connection are handled and inspected accordingly.

Best Practices for Handling Packets That Pass Before Traffic Identification

- The default action specified for an access control policy is NOT applied to these packets.
- Instead, use the following guidelines to choose a value for the **Intrusion Policy used before Access Control rule is determined** setting in the Advanced settings of the access control policy.
  - You can choose a system-created or custom intrusion policy. For example, you can choose **Balanced Security and Connectivity**.
  - For performance reasons, unless you have good reason to do otherwise, this setting should match the default action set for your access control policy.
  - If your system does not perform intrusion inspection (for example, in a discovery-only deployment), select **No Rules Active**. The system will not inspect these initial packets, and they will be allowed to pass.
  - By default, this setting uses the default variable set. Ensure that this is suitable for your purposes. For information, see Variable Sets, on page 412.
  - The network analysis policy associated with the first matching network analysis rule preprocesses traffic for the policy you select. If there are no network analysis rules, or none match, the default network analysis policy is used.

Specify a Policy to Handle Packets That Pass Before Traffic Identification



Note This setting is sometimes referred to as the *default intrusion policy*. (This is distinct from the default action for an access control policy.)

Caution Changing the total number of intrusion policies used by an access control policy restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information. You change the the total number of intrusion policies by adding an intrusion policy that is not currently used, or by removing the last instance of an intrusion policy. You can use an intrusion policy in an access control rule, as the default action, or as the default intrusion policy.

Before you begin

Review best practices for these settings. See Best Practices for Handling Packets That Pass Before Traffic Identification, on page 1670.

Procedure

Step 1 In the access control policy editor, click **Advanced**, then click **Edit** ( ) next to the **Network Analysis** and **Intrusion Policies** section.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 Select an intrusion policy from the **Intrusion Policy used before Access Control rule is determined** drop-down list.

If you choose a user-created policy, you can click **Edit** ( ) to edit the policy in a new window. You cannot edit system-provided policies.

Step 3 Optionally, select a different variable set from the **Intrusion Policy Variable Set** drop-down list. You can also select **Edit** ( ) next to the variable set to create and edit variable sets. If you do not change the variable set, the system uses a default set.

Step 4 Click **OK**.

Step 5 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Variable Sets, on page 412

Advanced Settings for Network Analysis Policies

Network analysis policies govern how traffic is decoded and preprocessed so that it can be further evaluated, especially for anomalous traffic that might signal an intrusion attempt. This traffic preprocessing occurs after Security Intelligence matching and traffic decryption, but before intrusion policies inspect packets in detail. By default, the system-provided Balanced Security and Connectivity network analysis policy is the default network analysis policy.

Tip The system-provided Balanced Security and Connectivity network analysis policy and the Balanced Security and Connectivity intrusion policy work together and can both be updated in intrusion rule updates. However, the network analysis policy governs mostly preprocessing options, whereas the intrusion policy governs mostly intrusion rules.

A simple way to tune preprocessing is to create and use a custom network analysis policy as the default. For advanced users with complex deployments, you can create multiple network analysis policies, each tailored to preprocess traffic differently. Then, you can configure the system to use those policies to govern the preprocessing of traffic using different security zones, networks, or VLANs.

To accomplish this, you add custom *network analysis rules* to your access control policy. A network analysis rule is simply a set of configurations and conditions that specifies how you preprocess traffic that matches those qualifications. You create and edit network analysis rules in the advanced options in an existing access control policy. Each rule belongs to only one policy.

Each rule has:

- a set of rule conditions that identifies the specific traffic you want to preprocess
- an associated network analysis policy that you want to use to preprocess traffic that meets all the rules' conditions

When it is time for the system to preprocess traffic, it matches packets to network analysis rules in top-down order by rule number. Traffic that does not match any network analysis rules is preprocessed by the default network analysis policy.

Setting the Default Network Analysis Policy

You can choose a system- or user-created policy.

Note If you disable a preprocessor but the system needs to evaluate preprocessed packets against an enabled intrusion or preprocessor rule, the system automatically enables and uses the preprocessor although it remains disabled in the network analysis policy web interface. Tailoring preprocessing, especially using multiple custom network analysis policies, is an **advanced** task. Because preprocessing and intrusion inspection are so closely related, you **must** be careful that you allow the network analysis and intrusion policies examining a single packet to complement each other.

Procedure

Step 1 In the access control policy editor, click **Advanced**, then click **Edit** ( ) next to the Network Analysis and Intrusion Policies section.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 From the **Default Network Analysis Policy** drop-down list, select a default network analysis policy.

If you choose a user-created policy, you can click **Edit** ( ) to edit the policy in a new window. You cannot edit system-provided policies.

Caution Changing the total number of network analysis policies used by an access control policy restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.You change the total number of network analysis policies by adding a policy that is not currently used, or by removing the last instance of a network analysis policy. You can use a network analysis policy with network analysis rules or as the default network analysis policy.

Step 3 Click **OK**.

Step 4 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Limitations of Custom Policies, on page 1457

Network Analysis Rules

Within your access control policy's advanced settings, you can use network analysis rules to tailor preprocessing configurations to network traffic.

Network analysis rules are numbered, starting at 1. When it is time for the system to preprocess traffic, it matches packets to network analysis rules in top-down order by ascending rule number, and preprocesses traffic according to the first rule where all the rule's conditions match.

You can add zone, network, and VLAN tag conditions to a rule. If you do not configure a particular condition for a rule, the system does not match traffic based on that criterion. For example, a rule with a network condition but no zone condition evaluates traffic based on its source or destination IP address, regardless of its ingress or egress interface. Traffic that does not match any network analysis rules is preprocessed by the default network analysis policy.

Configuring Network Analysis Rules

Procedure

Step 1 In the access control policy editor, click **Advanced**, then click **Edit** () next to the Network Analysis and Intrusion Policies section.

If **View** () appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings.If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Tip Click **Network Analysis Policy List** to view and edit existing custom network analysis policies.

Step 2 Next to **Network Analysis Rules**, click the statement that indicates how many custom rules you have.

Step 3 Click **Add Rule**.

Step 4 Configure the rule's conditions by clicking the conditions you want to add; see Rule Condition Types, on page 363.

Step 5 Click **Network Analysis** and choose the **Network Analysis Policy** you want to use to preprocess the traffic matching this rule.

Click **Edit** () to edit a custom policy in a new window. You cannot edit system-provided policies.

Caution Changing the total number of network analysis policies used by an access control policy restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.You change the total number of network analysis policies by adding a policy that is not currently used, or by removing the last instance of a network analysis policy. You can use a network analysis policy with network analysis rules or as the default network analysis policy.

Step 6 Click **Add**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Managing Network Analysis Rules

A network analysis rule is simply a set of configurations and conditions that specifies how you preprocess traffic that matches those qualifications. You create and edit network analysis rules in the advanced options in an existing access control policy. Each rule belongs to only one policy.

Procedure

Step 1 In the access control policy editor, click **Advanced**, then click **Edit** () next to the Intrusion and Network Analysis Policies section.

If **View** () appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings.If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 Next to **Network Analysis Rules**, click the statement that indicates how many custom rules you have.

Step 3 Edit your custom rules. You have the following options:

- To edit a rule's conditions, or change the network analysis policy invoked by the rule, click **Edit** () next to the rule.
- To change a rule's order of evaluation, click and drag the rule to the correct location. To select multiple rules, use the Shift and Ctrl keys.
- To delete a rule, click **Delete** () next to the rule.

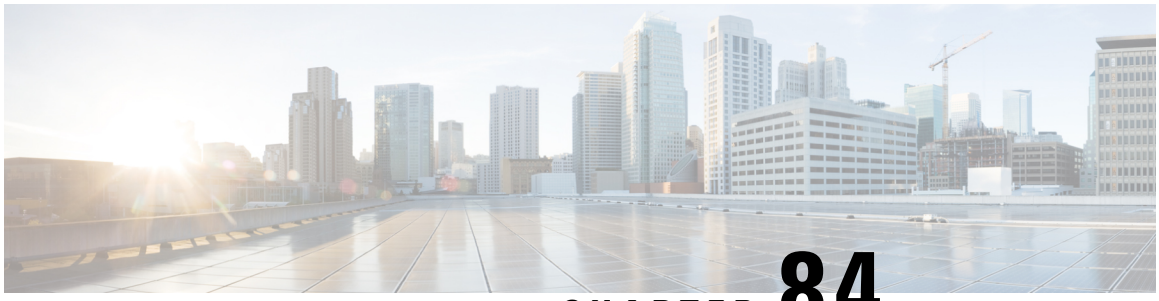
Tip Right-clicking a rule displays a context menu that allows you to cut, copy, paste, edit, delete, and add new network analysis rules.

Step 4 Click **OK**.

Step 5 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.



CHAPTER 84

Getting Started with Network Analysis Policies

The following topics describe how to get started with network analysis policies:

- Network Analysis Policy Basics, on page 1677
- License Requirements for Network Analysis Policies, on page 1678
- Requirements and Prerequisites for Network Analysis Policies, on page 1678
- Managing Network Analysis Policies, on page 1678

Network Analysis Policy Basics

Network analysis policies govern many traffic preprocessing options, and are invoked by advanced settings in your access control policy. Network analysis-related preprocessing occurs after Security Intelligence matching and SSL decryption, but before intrusion or file inspection begins.

By default, the system uses the *Balanced Security and Connectivity* network analysis policy to preprocess all traffic handled by an access control policy. However, you can choose a different default network analysis policy to perform this preprocessing. For your convenience, the system provides a choice of several non-modifiable network analysis policies, which are tuned for a specific balance of security and connectivity by the Cisco Talos Intelligence Group (Talos). You can also create a custom network analysis policy with custom preprocessing settings.

Tip System-provided intrusion and network analysis policies are similarly named but contain different configurations. For example, the Balanced Security and Connectivity network analysis policy and the Balanced Security and Connectivity intrusion policy work together and can both be updated in intrusion rule updates. However, the network analysis policy governs mostly preprocessing options, whereas the intrusion policy governs mostly intrusion rules. Network analysis and intrusion policies work together to examine your traffic.

You can also tailor traffic preprocessing options to specific security zones, networks, and VLANs by creating multiple custom network analysis policies, then assigning them to preprocess different traffic. (Note that ASA FirePOWER cannot restrict preprocessing by VLAN.)

License Requirements for Network Analysis Policies

FTD License

Threat

Classic License

Protection

Requirements and Prerequisites for Network Analysis Policies

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Intrusion Admin

Managing Network Analysis Policies

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Manage your network analysis policy:

- Compare—Click **Compare Policies**; see Comparing Policies, on page 355.
- Create — If you want to create a new network analysis policy, click **Create Policy**.

- Delete — If you want to delete a network analysis policy, click **Delete** (), then confirm that you want to delete the policy. You cannot delete a network analysis policy if an access control policy references it.

  If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

- Deploy—Click **Deploy**; see Deploy Configuration Changes, on page 346.

- Edit — If you want to edit an existing network analysis policy, click **Edit** () and proceed as described in Network Analysis Policy Settings and Cached Changes, on page 1681.

  If **View** () appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

- Report—Click **Report** () ; see Generating Current Policy Reports, on page 356.

Custom Network Analysis Policy Creation for Snort 2

When you create a new network analysis policy you must give it a unique name, specify a base policy, and choose an *inline mode*.

The base policy defines the network analysis policy's default settings. Modifying a setting in the new policy overrides—but does not change—the settings in the base policy. You can use either a system-provided or custom policy as your base policy.

The network analysis policy's inline mode allows preprocessors to modify (normalize) and drop traffic to minimize the chances of attackers evading detection. Note that in passive deployments, the system cannot affect traffic flow regardless of the inline mode.

Related Topics

The Base Layer, on page 1465
Preprocessor Traffic Modification in Inline Deployments, on page 1683
Creating a Custom Network Analysis Policy, on page 1679
Editing Network Analysis Policies, on page 1681

Creating a Custom Network Analysis Policy

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Create Policy**. If you have unsaved changes in another policy, click **Cancel** when prompted to return to the **Network Analysis Policy** page.

Step 3 Enter a unique **Name**.

In a multidomain deployment, policy names must be unique within the domain hierarchy. The system may identify a conflict with the name of a policy you cannot view in your current domain.

Step 4 Optionally, enter a **Description**.

Step 5 Choose the initial **Base Policy**. You can use either a system-provided or custom policy as your base policy.

Step 6 If you want to allow preprocessors to affect traffic in an inline deployment, enable **Inline Mode**.

Step 7 To create the policy:

- Click **Create Policy** to create the new policy and return to the **Network Analysis Policy** page. The new policy has the same settings as its base policy.
- Click **Create and Edit Policy** to create the policy and open it for editing in the advanced network analysis policy editor.

Related Topics

Creating Custom User Roles, on page 57

Network Analysis Policy Management for Snort 2

On the Network Analysis Policy page (or **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**), you can view your current custom network analysis policies, along with the following information:

- the time and date the policy was last modified (in local time) and the user who modified it
- whether the **Inline Mode** setting is enabled, which allows preprocessors to affect traffic
- which access control policies and devices are using the network analysis policy to preprocess traffic
- whether a policy has unsaved changes, as well as information about who (if anyone) is currently editing the policy

In addition to custom policies that you create, the system provides two custom policies: Initial Inline Policy and Initial Passive Policy. These two network analysis policies use the Balanced Security and Connectivity network analysis policy as their base. The only difference between them is their inline mode, which allows preprocessors to affect traffic in the inline policy and disables it in the passive policy. You can edit and use these system-provided custom policies.

Note that you can create and edit network analysis as well as intrusion policies if your Firepower System user account's role is restricted to Intrusion Policy or Modify Intrusion Policy.

Related Topics

Creating a Custom Network Analysis Policy, on page 1679
Editing Network Analysis Policies, on page 1681

Network Analysis Policy Settings and Cached Changes

When you create a new network analysis policy, it has the same settings as its base policy.

When tailoring a network analysis policy, especially when disabling preprocessors, keep in mind that some preprocessors and intrusion rules require that traffic first be decoded or preprocessed in a certain way. If you disable a required preprocessor, the system automatically uses it with its current settings, although the preprocessor remains disabled in the network analysis policy web interface.

Note Because preprocessing and intrusion inspection are so closely related, the network analysis and intrusion policies examining a single packet **must** complement each other. Tailoring preprocessing, especially using multiple custom network analysis policies, is an **advanced** task.

The system caches one network analysis policy per user. While editing a network analysis policy, if you select any menu or other path to another page, your changes stay in the system cache even if you leave the page.

Related Topics

How Policies Examine Traffic For Intrusions, on page 1448
Limitations of Custom Policies, on page 1457

Editing Network Analysis Policies

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the network analysis policy you want to configure.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Edit your network analysis policy:

- Change the base policy — If you want to change the base policy, choose a base policy from the **Base Policy** drop-down list on the Policy Information page.
- Manage policy layers — If you want to manage policy layers, click **Policy Layers** in the navigation panel.
- Modify a preprocessor — If you want to enable, disable, or edit the settings for a preprocessor, click **Settings** in the navigation panel.
- Modify traffic — If you want to allow preprocessors to modify or drop traffic, check the **Inline Mode** check box on the Policy Information page.

- View settings — If you want to view the settings in the base policy, click **Manage Base Policy** on the Policy Information page.

Step 4 To save changes you made in this policy since the last policy commit, choose **Policy Information**, then click **Commit Changes**. If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want a preprocessor to generate events and, in an inline deployment, drop offending packets, enable rules for the preprocessor. For more information, see Setting Intrusion Rule States, on page 1504.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

The Base Layer, on page 1465
Changing the Base Policy, on page 1467
Preprocessor Configuration in a Network Analysis Policy for Snort 2, on page 1682
Preprocessor Traffic Modification in Inline Deployments, on page 1683
Managing Layers, on page 1470
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

Preprocessor Configuration in a Network Analysis Policy for Snort 2

Preprocessors prepare traffic to be further inspected by normalizing traffic and identifying protocol anomalies. Preprocessors can generate preprocessor events when packets trigger preprocessor options that you configure. The base policy for your network analysis policy determines which preprocessors are enabled by default and the default configuration for each.

Note In most cases, preprocessors require specific expertise to configure and typically require little or no modification. Tailoring preprocessing, especially using multiple custom network analysis policies, is an **advanced** task. Because preprocessing and intrusion inspection are so closely related, the network analysis and intrusion policies examining a single packet **must** complement each other.

Modifying a preprocessor configuration requires an understanding of the configuration and its potential impact on your network.

Note that some advanced transport and network preprocessor settings apply globally to all networks, zones, and VLANs where you deploy your access control policy. You configure these advanced settings in an access control policy rather than in a network analysis policy.

Note also that you configure the sensitive data preprocessor, which detects sensitive data such as credit card numbers and Social Security numbers in ASCII text, in intrusion policies.

Related Topics

The DCE/RPC Preprocessor, on page 1686
The DNP3 Preprocessor, on page 1754
The DNS Preprocessor, on page 1697

The FTP/Telnet Decoder, on page 1700
The GTP Preprocessor, on page 1728
The HTTP Inspect Preprocessor, on page 1708
The IMAP Preprocessor, on page 1730
The Inline Normalization Preprocessor, on page 1762
The IP Defragmentation Preprocessor, on page 1769
The Modbus Preprocessor, on page 1752
The Packet Decoder, on page 1774
The POP Preprocessor, on page 1733
Sensitive Data Detection Basics, on page 1523
The SIP Preprocessor, on page 1724
The SMTP Preprocessor, on page 1736
The SSH Preprocessor, on page 1742
The SSL Preprocessor, on page 1746
The Sun RPC Preprocessor, on page 1722
TCP Stream Preprocessing, on page 1778
UDP Stream Preprocessing, on page 1788
Limitations of Custom Policies, on page 1457

Preprocessor Traffic Modification in Inline Deployments

In an inline deployment (that is, where relevant configurations are deployed to devices using routed, switched, or transparent interfaces, or inline interface pairs), some preprocessors can modify and block traffic. For example:

- The inline normalization preprocessor normalizes packets to prepare them for analysis by other preprocessors and the intrusion rules engine. You can also use the preprocessor's **Allow These TCP Options** and **Block Unresolvable TCP Header Anomalies** options to block certain packets.
- The system can drop packets with invalid checksums.
- The system can drop packets matching rate-based attack prevention settings.

For a preprocessor configured in the network analysis policy to affect traffic, you must enable and correctly configure the preprocessor, as well as correctly deploy managed devices inline. Finally, you must enable the network analysis policy's **Inline Mode** setting.

Preprocessor Configuration in a Network Analysis Policy Notes

When you select **Settings** in the navigation panel of a network analysis policy, the policy lists its preprocessors by type. On the Settings page, you can enable or disable preprocessors in your network analysis policy, as well as access preprocessor configuration pages.

A preprocessor must be enabled for you to configure it. When you enable a preprocessor, a sublink to the configuration page for the preprocessor appears beneath the **Settings** link in the navigation panel, and an **Edit** link to the configuration page appears next to the preprocessor on the Settings page.

Tip To revert a preprocessor's configuration to the settings in the base policy, click **Revert to Defaults** on a preprocessor configuration page. When prompted, confirm that you want to revert.

When you disable a preprocessor, the sublink and **Edit** link no longer appear, but your configurations are retained. Note that to perform their particular analysis, many preprocessors and intrusion rules require that traffic first be decoded or preprocessed in a certain way. If you disable a required preprocessor, the system automatically uses it with its current settings, although the preprocessor remains disabled in the network analysis policy web interface.

If you want to assess how your configuration would function in an inline deployment without actually modifying traffic, you can disable inline mode. In passive deployments or inline deployments in tap mode, the system cannot affect traffic regardless of the inline mode setting.



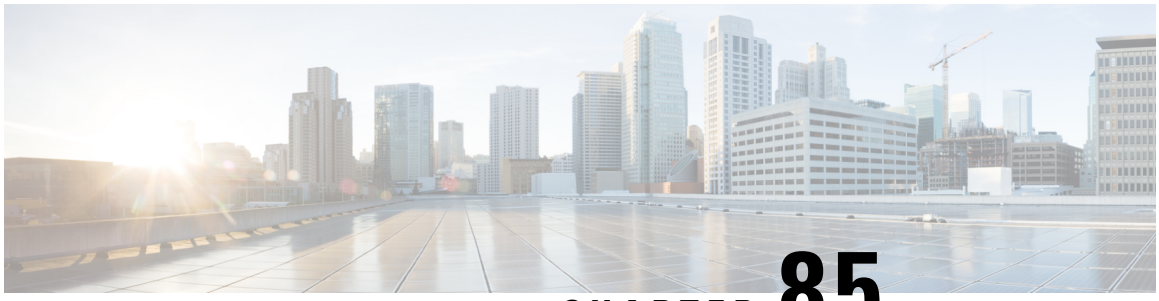
Note Disabling inline mode can affect intrusion event performance statistics graphs. With inline mode enabled in an inline deployment, the Intrusion Event Performance page (**Overview** > **Summary** > **Intrusion Event Performance**) displays graphs that represent normalized and blocked packets. If you disable inline mode, or in a passive deployment, many of the graphs display data about the traffic the system would have normalized or dropped.



Note In an inline deployment, Cisco recommends that you enable inline mode and configure the inline normalization preprocessor with the **Normalize TCP Payload** option enabled. In a passive deployment, Cisco recommends that you use adaptive profile updates.

Related Topics

Advanced Transport/Network Preprocessor Settings, on page 1758
Checksum Verification, on page 1761
The Inline Normalization Preprocessor, on page 1762
Intrusion Event Performance Statistics Graph Types, on page 2290



CHAPTER 85

Application Layer Preprocessors

The following topics explain application layer preprocessors and how to configure them:

- Introduction to Application Layer Preprocessors, on page 1685
- License Requirements for Application Layer Preprocessors, on page 1686
- Requirements and Prerequisites for Application Layer Preprocessors, on page 1686
- The DCE/RPC Preprocessor, on page 1686
- The DNS Preprocessor, on page 1697
- The FTP/Telnet Decoder, on page 1700
- The HTTP Inspect Preprocessor, on page 1708
- The Sun RPC Preprocessor, on page 1722
- The SIP Preprocessor, on page 1724
- The GTP Preprocessor, on page 1728
- The IMAP Preprocessor, on page 1730
- The POP Preprocessor, on page 1733
- The SMTP Preprocessor, on page 1736
- The SSH Preprocessor, on page 1742
- The SSL Preprocessor, on page 1746

Introduction to Application Layer Preprocessors

Application layer protocols can represent the same data in a variety of ways. The Firepower System provides application layer protocol decoders that normalize specific types of packet data into formats that the intrusion rules engine can analyze. Normalizing application-layer protocol encodings allows the rules engine to effectively apply the same content-related rules to packets whose data is represented differently and obtain meaningful results.

When an intrusion rule or rule argument requires a disabled preprocessor, the system automatically uses it with its current configuration even though it remains disabled in the network analysis policy's web interface.

Note that preprocessors do not generate events in most cases unless you enable the accompanying preprocessor rules in an intrusion policy.

License Requirements for Application Layer Preprocessors

FTD License

Threat

Classic License

Protection

Requirements and Prerequisites for Application Layer Preprocessors

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Intrusion Admin

The DCE/RPC Preprocessor

The DCE/RPC protocol allows processes on separate network hosts to communicate as if the processes were on the same host. These inter-process communications are commonly transported between hosts over TCP and UDP. Within the TCP transport, DCE/RPC might also be further encapsulated in the Windows Server Message Block (SMB) protocol or in Samba, an open-source SMB implementation used for inter-process communication in a mixed environment comprised of Windows and UNIX- or Linux-like operating systems. In addition, Windows IIS web servers on your network might use IIS RPC over HTTP, which provides distributed communication through a firewall, to proxy TCP-transported DCE/RPC traffic.

Note that descriptions of DCE/RPC preprocessor options and functionality include the Microsoft implementation of DCE/RPC known as MSRPC; descriptions of SMB options and functionality refer to both SMB and Samba.

Although most DCE/RPC exploits occur in DCE/RPC client requests targeted for DCE/RPC servers, which could be practically any host on your network that is running Windows or Samba, exploits can also occur in server responses. The DCE/RPC preprocessor detects DCE/RPC requests and responses encapsulated in TCP, UDP, and SMB transports, including TCP-transported DCE/RPC using version 1 RPC over HTTP. The preprocessor analyzes DCE/RPC data streams and detects anomalous behavior and evasion techniques in DCE/RPC traffic. It also analyzes SMB data streams and detects anomalous SMB behavior and evasion techniques.

The DCE/RPC preprocessor also desegments SMB and defragments DCE/RPC in addition to the IP defragmentation provided by the IP defragmentation preprocessor and the TCP stream reassembly provided by the TCP stream preprocessor.

Finally, the DCE/RPC preprocessor normalizes DCE/RPC traffic for processing by the rules engine.

Connectionless and Connection-Oriented DCE/RPC Traffic

DCE/RPC messages comply with one of two distinct DCE/RPC Protocol Data Unit (PDU) protocols:

connection-oriented DCE/RPC PDU protocol

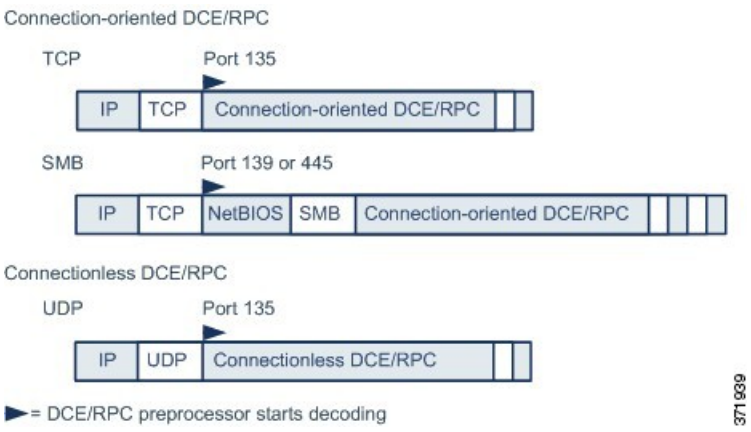
The DCE/RPC preprocessor detects connection-oriented DCE/RPC in the TCP, SMB, and RPC over HTTP transports.

connectionless DCE/RPC PDU protocol

The DCE/RPC preprocessor detects connectionless DCE/RPC in the UDP transport.

The two DCE/RPC PDU protocols have their own unique headers and data characteristics. For example, the connection-oriented DCE/RPC header length is typically 24 bytes and the connectionless DCE/RPC header length is fixed at 80 bytes. Also, correct fragment order of fragmented connectionless DCE/RPC cannot be handled by a connectionless transport and, instead, must be ensured by connectionless DCE/RPC header values; in contrast, the transport protocol ensures correct fragment order for connection-oriented DCE/RPC. The DCE/RPC preprocessor uses these and other protocol-specific characteristics to monitor both protocols for anomalies and other evasion techniques, and to decode and defragment traffic before passing it to the rules engine.

The following diagram illustrates the point at which the DCE/RPC preprocessor begins processing DCE/RPC traffic for the different transports.



Note the following in the figure:

- The well-known TCP or UDP port 135 identifies DCE/RPC traffic in the TCP and UDP transports.
- The figure does not include RPC over HTTP.

  For RPC over HTTP, connection-oriented DCE/RPC is transported directly over TCP as shown in the figure after an initial setup sequence over HTTP.
- The DCE/RPC preprocessor typically receives SMB traffic on the well-known TCP port 139 for the NetBIOS Session Service or the similarly implemented well-known Windows port 445.

Because SMB has many functions other than transporting DCE/RPC, the preprocessor first tests whether the SMB traffic is carrying DCE/RPC traffic and stops processing if it is not or continues processing if it is.

- IP encapsulates all DCE/RPC transports.
- TCP transports all connection-oriented DCE/RPC.
- UDP transports connectionless DCE/RPC.

DCE/RPC Target-Based Policies

Windows and Samba DCE/RPC implementations differ significantly. For example, all versions of Windows use the DCE/RPC context ID in the first fragment when defragmenting DCE/RPC traffic, and all versions of Samba use the context ID in the last fragment. As another example, Windows Vista uses the opnum (operation number) header field in the first fragment to identify a specific function call, and Samba and all other Windows versions use the opnum field in the last fragment.

There are also significant differences in Windows and Samba SMB implementations. For example, Windows recognizes the SMB OPEN and READ commands when working with named pipes, but Samba does not recognize these commands.

When you enable the DCE/RPC preprocessor, you automatically enable a default target-based policy. Optionally, you can add target-based policies that target other hosts running different Windows or Samba versions. The default target-based policy applies to any host not included in another target-based policy.

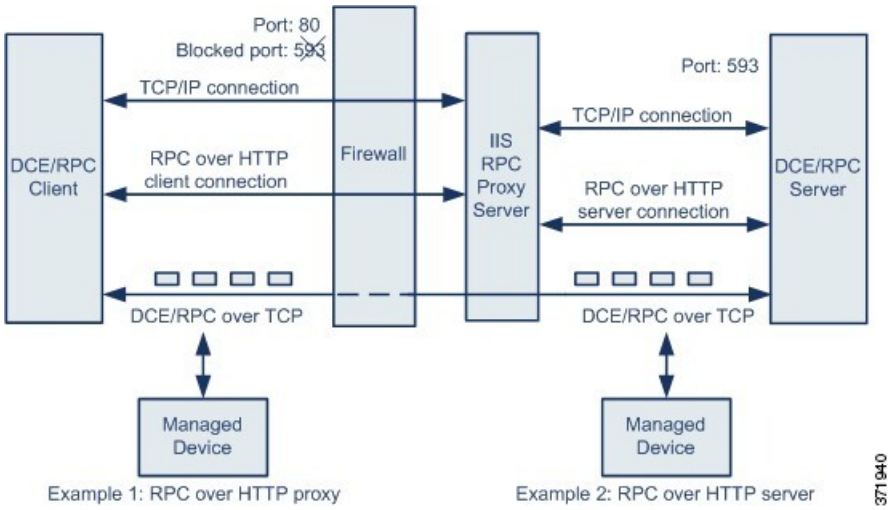
In each target-based policy, you can:

- enable one or more transports and specify *detection ports* for each
- enable and specify *auto-detection ports*
- set the preprocessor to detect when there is an attempt to connect to one or more shared SMB resources that you identify
- configure the preprocessor to detect files in SMB traffic and to inspect a specified number of bytes in a detected file
- modify an advanced option that should be modified only by a user with SMB protocol expertise; this option lets you set the preprocessor to detect when a number of chained SMB AndX commands exceed a specified maximum number

In addition to enabling SMB traffic file detection in the DCE/RPC preprocessor, you can configure a file policy to optionally capture and block these files, or submit them to the Cisco AMP cloud for dynamic analysis. Within that policy, you must create a file rule with an **Action** of **Detect Files** or **Block Files** and a selected **Application Protocol** of **Any** or **NetBIOS-ssn (SMB)**.

RPC over HTTP Transport

Microsoft RPC over HTTP allows you to tunnel DCE/RPC traffic through a firewall as shown in the following diagram. The DCE/RPC preprocessor detects version 1 of Microsoft RPC over HTTP.



The Microsoft IIS proxy server and the DCE/RPC server can be on the same host or on different hosts. Separate proxy and server options provide for both cases. Note the following in the figure:

- The DCE/RPC server monitors port 593 for DCE/RPC client traffic, but the firewall blocks port 593.

  Firewalls typically block port 593 by default.

- RPC over HTTP transports DCE/RPC over HTTP using well-known HTTP port 80, which firewalls are likely to permit.
- Example 1 shows that you would choose the **RPC over HTTP proxy** option to monitor traffic between the DCE/RPC client and the Microsoft IIS RPC proxy server.
- Example 2 shows that you would choose the **RPC over HTTP server** option when the Microsoft IIS RPC proxy server and the DCE/RPC server are located on different hosts and the device monitors traffic between the two servers.
- Traffic is comprised solely of connection-oriented DCE/RPC over TCP after RPC over HTTP completes the proxied setup between the DCE/RPC client and server.

DCE/RPC Global Options

Global DCE/RPC preprocessor options control how the preprocessor functions. Note that, except for the **Memory Cap Reached** and **Auto-Detect Policy on SMB Session** options, modifying these options could have a negative impact on performance or detection capability. You should not modify them unless you have a thorough understanding of the preprocessor and the interaction between the preprocessor and enabled DCE/RPC rules.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Maximum Fragment Size

When **Enable Defragmentation** is selected, specifies the maximum DCE/RPC fragment length allowed. The preprocessor truncates larger fragments for processing purposes to the specified size before defragmenting but does not alter the actual packet. A blank field disables this option.

Make sure that the **Maximum Fragment Size** option is greater than or equal to the depth to which the rules need to detect.

Reassembly Threshold

When **Enable Defragmentation** is selected, 0 disables this option, or specifies a minimum number of fragmented DCE/RPC bytes and, if applicable, segmented SMB bytes to queue before sending a reassembled packet to the rules engine. A low value increases the likelihood of early detection but could have a negative impact on performance. You should test for performance impact if you enable this option.

Make sure that the **Reassembly Threshold** option is greater than or equal to the depth to which the rules need to detect.

Enable Defragmentation

Specifies whether to defragment fragmented DCE/RPC traffic. When disabled, the preprocessor still detects anomalies and sends DCE/RPC data to the rules engine, but at the risk of missing exploits in fragmented DCE/RPC data.

Although this option provides the flexibility of not defragmenting DCE/RPC traffic, most DCE/RPC exploits attempt to take advantage of fragmentation to hide the exploit. Disabling this option would bypass most known exploits, resulting in a large number of false negatives.

Memory Cap Reached

Detects when the maximum memory limit allocated to the preprocessor is reached or exceeded. When the maximum memory cap is reached or exceeded, the preprocessor frees all pending data associated with the session that caused the memory cap event and ignores the rest of that session.

You can enable rule 133:1 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Auto-Detect Policy on SMB Session

Detects the Windows or Samba version that is identified in SMB `Session Setup AndX` requests and responses. When the detected version is different from the Windows or Samba version configured for the **Policy** configuration option, the detected version overrides the configured version for that session only.

For example, if you set **Policy** to Windows XP and the preprocessor detects Windows Vista, the preprocessor uses a Windows Vista policy for that session. Other settings remain in effect.

When the DCE/RPC transport is not SMB (that is, when the transport is TCP or UDP), the version cannot be detected and the policy cannot be automatically configured.

To enable this option, choose one of the following from the drop-down list:

- Choose **Client** to inspect server-to-client traffic for the policy type.
- Choose **Server** to inspect client-to-server traffic for the policy type.
- Choose **Both** to inspect server-to-client and client-to-server traffic for the policy type.

Legacy SMB Inspection Mode

Specifies which SMB versions to inspect. When **Legacy SMB Inspection Mode** is enabled, the DCE/RPC preprocessor inspects only SMB Version 1 traffic. When this option is disabled, the DCE/RPC preprocessor inspects traffic that uses SMB Versions 1, 2, and 3.

Related Topics

Basic content and protected_content Keyword Arguments, on page 1567
Overview: The byte_jump and byte_test Keywords

DCE/RPC Target-Based Policy Options

In each target-based policy, you can enable one or more of the TCP, UDP, SMB, and RPC over HTTP transports. When you enable a transport, you must also specify one or more *detection ports*, that is, ports that are known to carry DCE/RPC traffic.

Cisco recommends that you use the default detection ports, which are either well-known ports or otherwise commonly-used ports for each protocol. You would add detection ports only if you detected DCE/RPC traffic on a non-default port.

You can specify ports for one or more transports in any combination in a Windows target-based policy to match the traffic on your network, but you can only specify ports for the SMB transport in a Samba target-based policy.



Note You must enable at least one DCE/RPC transport in the default target-based policy except when you have added a DCE/RPC target-based policy that has at least one transport enabled. For example, you might want to specify the hosts for all DCE/RPC implementations and not have the default target-based policy deploy to unspecified hosts, in which case you would not enable a transport for the default target-based policy.

Optionally, you can also enable and specify *auto-detection ports*, that is, ports that the preprocessor tests first to determine if they carry DCE/RPC traffic and continues processing only when it detects DCE/RPC traffic.

When you enable auto-detection ports, ensure that they are set to the port range from 1025 to 65535 to cover the entire ephemeral port range.

Note that auto-detection occurs only for ports not already identified by transport detection ports.

It is unlikely that you would enable or specify auto-detection ports for the RPC over HTTP Proxy Auto-Detect Ports option or the SMB Auto-Detect Ports option because there is little likelihood that traffic for either would occur or even be possible except on the specified default detection ports.

Each target-based policy allows you to specify the various options below. If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Networks

The host IP addresses where you want to deploy the DCE/RPC target-based server policy. Also named the **Server Address** field in the Add Target pop-up window when you add a target-based policy.

You can specify a single IP address or address block, or a comma-separated list of either or both. You can configure up to 255 total profiles including the default policy.

Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Note that the `default` setting in the default policy specifies all IP addresses on your monitored network segment that are not covered by another target-based policy. Therefore, you cannot and do not need to specify an IP address or CIDR block/prefix length for the default policy, and you cannot leave this setting blank in another policy or use address notation to represent `any` (for example, 0.0.0.0/0 or ::/0).

Policy

The Windows or Samba DCE/RPC implementation used by the targeted host or hosts on your monitored network segment.

Note that you can enable the **Auto-Detect Policy on SMB Session** global option to automatically override the setting for this option on a per session basis when SMB is the DCE/RPC transport.

SMB Invalid Shares

Identifies one or more SMB shared resources the preprocessor will detect when there is an attempt to connect to a shared resource that you specify. You can specify multiple shares in a comma-separated list and, optionally, you can enclose shares in quotes, which was required in previous software versions but is no longer required; for example:

```
"C$", D$, "admin", private
```

The preprocessor detects invalid shares in SMB traffic when you have enabled **SMB Ports**.

Note that in most cases you should append a dollar sign to a drive named by Windows that you identify as an invalid share. For example, identify drive C as C$ or "C$".

Note also that to detect SMB invalid shares, you must also enable **SMB Ports** or **SMB Auto-Detect Ports.**

You can enable rule 133:26 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

SMB Maximum AndX Chain

The maximum number of chained SMB AndX commands to permit. Typically, more than a few chained AndX commands represent anomalous behavior and could indicate an evasion attempt. Specify 1 to permit no chained commands or 0 to disable detecting the number of chained commands.

Note that the preprocessor first counts the number of chained commands and generates an event if accompanying SMB preprocessor rules are enabled and the number of chained commands equals or exceeds the configured value. It then continues processing.

Caution Only someone who is expert in the SMB protocol should modify the setting for the **SMB Maximum AndX Chains** option.

You can enable rule 133:20 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

RPC proxy traffic only

Enabling **RPC over HTTP Proxy Ports** indicates whether detected client-side RPC over HTTP traffic is proxy traffic only or might include other web server traffic. For example, port 80 could carry both proxy and other web server traffic.

When this option is disabled, both proxy and other web server traffic are expected. Enable this option, for example, if the server is a dedicated proxy server. When enabled, the preprocessor tests traffic to determine if it carries DCE/RPC, ignores the traffic if it does not, and continues processing if it does. Note that enabling this option adds functionality only if the **RPC over HTTP Proxy Ports** check box is also enabled.

RPC over HTTP Proxy Ports

Enables detection of DCE/RPC traffic tunneled by RPC over HTTP over each specified port when your managed device is positioned between the DCE/RPC client and the Microsoft IIS RPC proxy server.

When enabled, you can add any ports where you see DCE/RPC traffic, although this is unlikely to be necessary because web servers typically use the default port for both DCE/RPC and other traffic. When enabled, you would not enable **RPC over HTTP Proxy Auto-Detect Ports**, but you would enable the **RPC Proxy Traffic Only** when detected client-side RPC over HTTP traffic is proxy traffic only and does not include other web server traffic.



Note You would rarely, if ever, select this option.

RPC over HTTP Server Ports

Enables detection of DCE/RPC traffic tunneled by RPC over HTTP on each specified port when the Microsoft IIS RPC proxy server and the DCE/RPC server are located on different hosts and the device monitors traffic between the two servers.

Typically, when you enable this option you should also enable **RPC over HTTP Server Auto-Detect Ports** with a port range from 1025 to 65535 for that option even if you are not aware of any proxy web servers on your network. Note that the RPC over HTTP server port is sometimes reconfigured, in which case you should add the reconfigured server port to port list for this option.

TCP Ports

Enables detection of DCE/RPC traffic in TCP on each specified port.

Legitimate DCE/RPC traffic and exploits might use a wide variety of ports, and other ports above port 1024 are common. Typically, when this option is enabled you should also enable **TCP Auto-Detect Ports** with a port range from 1025 to 65535 for that option.

UDP Ports

Enables detection of DCE/RPC traffic in UDP on each specified port.

Legitimate DCE/RPC traffic and exploits might use a wide variety of ports, and other ports above port 1024 are common. Typically, when this option is enabled you should also enable **UDP Auto-Detect Ports** with a port range from 1025 to 65535 for that option.

SMB Ports

Enables detection of DCE/RPC traffic in SMB on each specified port.

You could encounter SMB traffic using the default detection ports. Other ports are rare. Typically, use the default settings.

Note that you can enable the **Auto-Detect Policy on SMB Session** global option to automatically override the policy type configured for a targeted policy on a per session basis when SMB is the DCE/RPC transport.

RPC over HTTP Proxy Auto-Detect Ports

Enables auto-detection of DCE/RPC traffic tunneled by RPC over HTTP on the specified ports when your managed device is positioned between the DCE/RPC client and the Microsoft IIS RPC proxy server.

When enabled, you would typically specify a port range from 1025 to 65535 to cover the entire range of ephemeral ports.

RPC over HTTP Server Auto-Detect Ports

Enables auto-detection of DCE/RPC traffic tunneled by RPC over HTTP on the specified ports when the Microsoft IIS RPC proxy server and the DCE/RPC server are located on different hosts and the device monitors traffic between the two servers.

TCP Auto-Detect Ports

Enables auto-detection of DCE/RPC traffic in TCP on the specified ports.

UDP Auto-Detect Ports

Enables auto-detection of DCE/RPC traffic in UDP on each specified port.

SMB Auto-Detect Ports

Enables auto-detection of DCE/RPC traffic in SMB.



Note You would rarely, if ever, select this option.

SMB File Inspection

Enables inspection of SMB traffic for file detection. You have the following options:

- Select **Off** to disable file inspection.
- Select **Only** to inspect file data without inspecting the DCE/RPC traffic in SMB. Selecting this option can improve performance over inspecting both files and DCE/RPC traffic.
- Select **On** to inspect both files and the DCE/RPC traffic in SMB. Selecting this option can impact performance.

Inspection of SMB traffic for the following is not supported:

- files transferred in an established TCP or SMB session before this option is enabled and the policy applied
- files transferred concurrently in a single TCP or SMB session

- files transferred across multiple TCP or SMB sessions
- files transferred with non-contiguous data, such as when message signing is negotiated
- files transferred with different data at the same offset, overlapping the data
- files opened on a remote client for editing that the client saves to the file server

SMB File Inspection Depth

If **SMB File Inspection** is set to **Only** or **On**, the number of bytes inspected when a file is detected in SMB traffic. Specify one of the following:

- a positive value
- `0` to inspect the entire file
- `-1` to disable file inspection

Enter a value in this field equal to or smaller than the one defined in the File and Malware Settings section of the Advanced tab in your access control policy. If you set a value for this option larger than the one defined for **Limit the number of bytes inspected when doing file type detection**, the system uses the access control policy setting as the functional maximum.

If **SMB File Inspection** is set to **Off**, this field is disabled.

Related Topics

Firepower System IP Address Conventions, on page 15

Traffic-Associated DCE/RPC Rules

Most DCE/RPC preprocessor rules trigger against anomalies and evasion techniques detected in SMB, connection-oriented DCE/RPC, or connectionless DCE/RPC traffic. The following table identifies the rules that you can enable for each type of traffic.

Table 192: Traffic-Associated DCE/RPC Rules

Traffic	Preprocessor Rule GID:SID
SMB	133:2 through 133:26, and 133:48 through 133:59
Connection-Oriented DCE/RPC	133:27 through 133:39
Detect Connectionless DCE/RPC	133:40 through 133:43

Configuring the DCE/RPC Preprocessor

You configure the DCE/RPC preprocessor by modifying any of the global options that control how the preprocessor functions, and by specifying one or more target-based server policies that identify the DCE/RPC servers on your network by IP address and by either the Windows or Samba version running on them. Target-based policy configuration also includes enabling transport protocols, specifying the ports carrying DCE/RPC traffic to those hosts, and setting other server-specific options.

The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Before you begin

- Confirm that networks you want to identify in a custom target-based policy match or are a subset of the networks, zones, and VLANs handled by its parent network analysis policy. See Advanced Settings for Network Analysis Policies, on page 1672 for more information.

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation panel on the left.

Step 4 If **DCE/RPC Configuration** under **Application Layer Preprocessors** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **DCE/RPC Configuration**.

Step 6 Modify the options in the **Global Settings** section; see DCE/RPC Global Options, on page 1689.

Step 7 You have the following choices:

- Add a server profile — Click **Add** ( ) next to **Servers**. Specify one or more IP addresses in the **Server Address** field, then click **OK**.
- Delete a server profile — Click **Delete** ( ) next to the policy.
- Edit a server profile — Click the configured address for the profile under **Servers**, or click **default**. You can modify any of the settings in the **Configuration** section; see DCE/RPC Target-Based Policy Options, on page 1691.

Step 8 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, cached changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want to generate intrusion events, enable DCE/RPC preprocessor rules (GID 132 or 133). For more information, see Setting Intrusion Rule States, on page 1504, DCE/RPC Global Options, on page

1689, DCE/RPC Target-Based Policy Options, on page 1691, and Traffic-Associated DCE/RPC Rules, on page 1695.

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Firepower System IP Address Conventions, on page 15
File and Malware Inspection Performance and Storage Options, on page 1441
DCE/RPC Keywords, on page 1614
Managing Layers, on page 1470
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

The DNS Preprocessor

The DNS preprocessor inspects DNS name server responses for the following specific exploits:

- Overflow attempts on RData text fields
- Obsolete DNS resource record types
- Experimental DNS resource record types

The most common type of DNS name server response provides one or more IP addresses that correspond to domain names in the query that prompted the response. Other types of server responses provide, for example, the destination of an email message or the location of a name server that can provide information not available from the server originally queried.

A DNS response is comprised of:

- a message header
- a Question section that contains one or more requests
- three sections that respond to requests in the Question section
  - Answer
  - Authority
  - Additional Information.

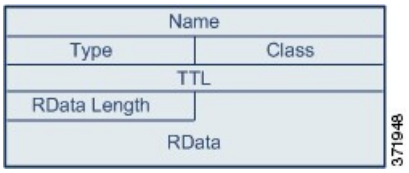
Responses in these three sections reflect the information in *resource records* (RR) maintained on the name server. The following table describes these three sections.

Table 193: DNS Name Server RR Responses

This section...	Includes...	For example...
Answer	Optionally, one or more resource records that provide a specific answer to a query	The IP address corresponding to a domain name
Authority	Optionally, one or more resource records that point to an authoritative name server	The name of an authoritative name server for the response

This section...	Includes...	For example...
Additional Information	Optionally, one or more resource records that provided additional information related to the Answer sections	The IP address of another server to query

There are many types of resource records, all adhering to the following structure:



Theoretically, any type of resource record can be used in the Answer, Authority, or Additional Information section of a name server response message. The DNS preprocessor inspects any resource record in each of the three response sections for the exploits it detects.

The Type and RData resource record fields are of particular importance to the DNS preprocessor. The Type field identifies the type of resource record. The RData (resource data) field provides the response content. The size and content of the RData field differ depending on the type of resource record.

DNS messages typically use the UDP transport protocol but also use TCP when the message type requires reliable delivery or the message size exceeds UDP capabilities. The DNS preprocessor inspects DNS server responses in both UDP and TCP traffic.

The DNS preprocessor does not inspect TCP sessions picked up in midstream, and ceases inspection if a session loses state because of dropped packets.

DNS Preprocessor Options

Ports

This field specifies the source port or ports the DNS preprocessor should monitor for DNS server responses. Separate multiple ports with commas.

The typical port to configure for the DNS preprocessor is well-known port 53, which DNS name servers use for DNS messages in both UDP and TCP.

Detect Overflow attempts on RData Text fields

When the resource record type is TXT (text), the RData field is a variable-length ASCII text field.

When selected, this option detects a specific vulnerability identified by entry CVE-2006-3441 in MITRE's Current Vulnerabilities and Exposures database. This is a known vulnerability in Microsoft Windows 2000 Service Pack 4, Windows XP Service Pack 1 and Service Pack 2, and Windows Server 2003 Service Pack 1. An attacker can exploit this vulnerability and take complete control of a host by sending or otherwise causing the host to receive a maliciously crafted name server response that causes a miscalculation in the length of an RData text field, resulting in a buffer overflow.

You should enable this option when your network might include hosts running operating systems that have not been upgraded to correct this vulnerability.

You can enable rule 131:3 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Detect Obsolete DNS RR Types

RFC 1035 identifies several resource record types as obsolete. Because these are obsolete record types, some systems do not account for them and may be open to exploits. You would not expect to encounter these record types in normal DNS responses unless you have purposely configured your network to include them.

You can configure the system to detect known obsolete resource record types. The following table lists and describes these record types.

Table 194: Obsolete DNS Resource Record Types

RR Type	Code	Description
3	MD	a mail destination
4	MF	a mail forwarder

You can enable rule 131:1 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Detecting Experimental DNS RR Types

RFC 1035 identifies several resource record types as experimental. Because these are experimental record types, some systems do not account for them and may be open to exploits. You would not expect to encounter these record types in normal DNS responses unless you have purposely configured your network to include them.

You can configure the system to detect known experimental resource record types. The following table lists and describes these record types.

Table 195: Experimental DNS Resource Record Types

RR Type	Code	Description
7	MB	a mailbox domain name
8	MG	a mail group member
9	MR	a mail rename domain name
10	NUL	a null resource record

You can enable rule 131:2 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Configuring the DNS Preprocessor

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation panel.

Step 4 If **DNS Configuration** under **Application Layer Preprocessors** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **DNS Configuration**.

Step 6 Modify the settings as described in DNS Preprocessor Options, on page 1698.

Step 7 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, cached changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want to generate intrusion events, enable DNS preprocessor rules (GID 131). For more information, see Setting Intrusion Rule States, on page 1504 and DNS Preprocessor Options, on page 1698.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Layers in Intrusion and Network Analysis Policies, on page 1463
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

The FTP/Telnet Decoder

The FTP/Telnet decoder analyzes FTP and telnet data streams, normalizing FTP and telnet commands before processing by the rules engine.

Global FTP and Telnet Options

You can set global options to determine whether the FTP/Telnet decoder performs stateful or stateless inspection of packets, whether the decoder detects encrypted FTP or telnet sessions, and whether the decoder continues to check a data stream after it encounters encrypted data.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Stateful Inspection

When selected, causes the FTP/Telnet decoder to save state and provide session context for individual packets and only inspect reassembled sessions. When cleared, analyzes each individual packet without session context.

To check for FTP data transfers, this option must be selected.

Detect Encrypted Traffic

Detects encrypted telnet and FTP sessions.

You can enable rules 125:7 and 126:2 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Continue to Inspect Encrypted Data

Instructs the preprocessor to continue checking a data stream after it is encrypted, looking for eventual decrypted data that can be processed.

Telnet Options

You can enable or disable normalization of telnet commands by the FTP/Telnet decoder, enable or disable a specific anomaly case, and set the threshold number of Are You There (AYT) attacks to permit.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Ports

Indicates the ports whose telnet traffic you want to normalize. Telnet typically connects to TCP port 23. In the interface, list multiple ports separated by commas.



Caution Because encrypted traffic (SSL) cannot be decoded, adding port 22 (SSH) could yield unexpected results.

Normalize

Normalizes telnet traffic to the specified ports.

Detect Anomalies

Enables detection of Telnet SB (subnegotiation begin) without the corresponding SE (subnegotiation end).

Telnet supports subnegotiation, which begins with SB (subnegotiation begin) and must end with an SE (subnegotiation end). However, certain implementations of Telnet servers will ignore the SB without a corresponding SE. This is anomalous behavior that could be an evasion case. Because FTP uses the Telnet protocol on the control connection, it is also susceptible to this behavior.

You can enable rule 126:3 to generate an event and, in an inline deployment, drop offending packets when this anomaly is detected in Telnet traffic, and rule 125:9 when it is detected on the FTP command channel. See Setting Intrusion Rule States, on page 1504.

Are You There Attack Threshold Number

Detects when the number of consecutive AYT commands exceeds the specified threshold. Cisco recommends that you set the AYT threshold to a value no higher than the default value.

You can enable rule 126:1 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Server-Level FTP Options

You can set options for decoding on multiple FTP servers. Each server profile you create contains the server IP address and the ports on the server where traffic should be monitored. You can specify which FTP commands to validate and which to ignore for a particular server, and set maximum parameter lengths for commands. You can also set the specific command syntax the decoder should validate against for particular commands and set alternate maximum command parameter lengths.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Networks

Use this option to specify one or more IP addresses of FTP servers.

You can specify a single IP address or address block, or a comma-separated list comprised of either or both. You can configure up to 1024 characters, and you can specify up to 255 profiles including the default profile.

Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Note that the `default` setting in the default policy specifies all IP addresses on your monitored network segment that are not covered by another target-based policy. Therefore, you cannot and do not need to specify an IP address or address block for the default policy, and you cannot leave this setting blank in another policy or use address notation to represent `any` (for example, 0.0.0.0/0 or ::/0).

Ports

Use this option to specify the ports on the FTP server where the managed device should monitor traffic. In the interface, list multiple ports separated by commas. Port 21 is the well-known port for FTP traffic.

File Get Commands

Use this option to define the FTP commands used to transfer files from server to client. Do not change these values unless directed to do so by Support.

Caution Do not modify the **File Get Commands** field unless directed to by Support.

File Put Commands

Use this option to define the FTP commands used to transfer files from client to server. Do not change these values unless directed to do so by Support.



Caution Do not modify the **File Put Commands** field unless directed to by Support.

Additional FTP Commands

Use this line to specify the additional commands that the decoder should detect. Separate additional commands by spaces.

Additional commands you may want to add include `XPWD`, `XCWD`, `XCUP`, `XMKD`, and `XRMD`. For more information on these commands, see RFC 775, the Directory oriented FTP commands specification by the Network Working Group.

Default Max Parameter Length

Use this option to detect the maximum parameter length for commands where an alternate maximum parameter length has not been set. You can add as many alternative maximum parameter lengths as needed.

You can enable rule 125:3 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Alternate Max Parameter Length

Use this option to specify commands where you want to detect a different maximum parameter length, and to specify the maximum parameter length for those commands. Click **Add** to add lines where you can specify a different maximum parameter length to detect for particular commands.

Check Commands for String Format Attacks

Use this option to check the specified commands for string format attacks.

You can enable rule 125:5 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Command Validity

Use this option to enter a valid format for a specific command. Click **Add** to add a command validation line.

You can enable rules 125:2 and 125:4 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Ignore FTP Transfers

Use this option to improve performance on FTP data transfers by disabling all inspection other than state inspection on the data transfer channel.



Note To inspect data transfers, the global FTP/Telnet **Stateful Inspection** option must be selected.

Detect Telnet Escape Codes within FTP Commands

Use this option to detect when telnet commands are used over the FTP command channel.

You can enable rule 125:1 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Ignore Erase Commands during Normalization

When **Detect Telnet Escape Codes within FTP Commands** is selected, use this option to ignore telnet character and line erase commands when normalizing FTP traffic. The setting should match how the FTP server handles telnet erase commands. Note that newer FTP servers typically ignore telnet erase commands, while older servers typically process them.

Troubleshooting Option: Log FTP Command Validation Configuration

Support might ask you during a troubleshooting call to configure your system to print the configuration information for each FTP command listed for the server.

Caution Do not enable **Log FTP Command Validation Configuration** unless instructed to do so by Support.

FTP Command Validation Statements

When setting up a validation statement for an FTP command, you can specify a group of alternative parameters by separating the parameters with spaces. You can also create a binary OR relationship between two parameters by separating them with a pipe character (`|`) in the validation statement. Surrounding parameters by square brackets (`[]`) indicates that those parameters are optional. Surrounding parameters with curly brackets (`{}`) indicates that those parameters are required.

You can create FTP command parameter validation statements to validate the syntax of a parameter received as part of an FTP communication.

Any of the parameters listed in the following table can be used in FTP command parameter validation statements.

Table 196: FTP Command Parameters

If you use...	The following validation occurs...
`int`	The represented parameter must be an integer.
`number`	The represented parameter must be an integer between 1 and 255.
`char` _chars	The represented parameter must be a single character and a member of the characters specified in the _chars argument. For example, defining the command validity for `MODE` with the validation statement `char` SBC checks that the parameter for the `MODE` command comprises the character `S` (representing Stream mode), the character `B` (representing Block mode), or the character `C` (representing Compressed mode).

If you use...	The following validation occurs...
`date` _datefmt	If _datefmt contains `#`, the represented parameter must be a number. If _datefmt contains `c`, the represented parameter must be a character. If _datefmt contains literal strings, the represented parameter must match the literal string.
`string`	The represented parameter must be a string.
`host_port`	The represented parameter must be a valid host port specifier as defined by RFC 959, the File Transfer Protocol specification by the Network Working Group.

You can combine the syntax in the table above as needed to create parameter validation statements that correctly validate each FTP command where you need to validate traffic.

Note When you include a complex expression in a TYPE command, surround it by spaces. Also, surround each operand within the expression by spaces. For example, type `char A | B` , not `char A|B`.

Related Topics

Server-Level FTP Options, on page 1702
Firepower System IP Address Conventions, on page 15
FTP Command Validation Statements, on page 1704

Client-Level FTP Options

Use these options to configure custom FTP client profiles. If an option description does not include a preprocessor rule, the option is not associated with a preprocessor rule.

Networks

Use this option to specify one or more IP addresses of FTP clients.

You can specify a single IP address or address block, or a comma-separated list comprised of either or both. You can specify up to 1024 characters, and you can specify up to 255 profiles including the default profile.

Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Note that the `default` setting in the default policy specifies all IP addresses on your monitored network segment that are not covered by another target-based policy. Therefore, you cannot and do not need to specify an IP address or address block for the default policy, and you cannot leave this setting blank in another policy or use address notation to represent `any` (for example, 0.0.0.0/0 or ::/0).

Max Response Length

Use this option to specify the maximum allowed response length to an FTP command accepted by the client. This can detect basic buffer overflows.

You can enable rule 125:6 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Detect FTP Bounce Attempts

Use this option to detect FTP bounce attacks.

You can enable rule 125:8 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Allow FTP Bounce to

Use this option to configure a list of additional hosts and ports on those hosts on which FTP PORT commands should not be treated as FTP bounce attacks.

Detect Telnet Escape Codes within FTP Commands

Use this option to detect when telnet commands are used over the FTP command channel.

You can enable rule 125:1 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Ignore Erase Commands During Normalization

When **Detect Telnet Escape Codes within FTP Commands** is selected, use this option to ignore telnet character and line erase commands when normalizing FTP traffic.The setting should match how the FTP client handles telnet erase commands. Note that newer FTP clients typically ignore telnet erase commands, while older clients typically process them.

Related Topics

Firepower System IP Address Conventions, on page 15

Configuring the FTP/Telnet Decoder

You can configure client profiles for FTP clients to monitor FTP traffic from clients.

The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Before you begin

- Confirm that any networks you want to identify in a custom target-based policy match or are a subset of the networks, zones, and VLANs handled by its parent network analysis policy. See Advanced Settings for Network Analysis Policies, on page 1672 for more information.

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation panel.

Step 4 If **FTP and Telnet Configuration** under **Application Layer Preprocessors** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **FTP and Telnet Configuration**.

Step 6 Set options in the **Global Settings** section as described in Global FTP and Telnet Options, on page 1700.

Step 7 Set options in the **Telnet Settings** section as described in Telnet Options, on page 1701.

Step 8 Manage FTP server profiles:

- Add a server profile — Click **Add** ( ) next to **FTP Server**. Specify one or more IP addresses for the client in the **Server Address** field and click **OK**. You can specify a single IP address or address block, or a comma-separated list of either or both. You can specify up to 1024 characters, and you can configure up to 255 policies, including the default policy.
- Edit a server profile — Click the configured address for a custom profile under **FTP Server**, or click **default**. You can modify the settings in the **Configuration** section; see Server-Level FTP Options, on page 1702.
- Delete a server profile — Click **Delete** ( ) next to the profile.

Step 9 Manage FTP client profiles:

- Add a client profile — Click **Add** ( ) next to **FTP Client**. Specify one or more IP addresses for the client in the **Client Address** field and click **OK**. You can specify a single IP address or address block, or a comma-separated list of either or both. You can specify up to 1024 characters, and you can configure up to 255 policies, including the default policy.
- Edit a client profile — Click the configured address for a profile you have added under **FTP Client**, or click **default**. You can modify the settings in the Configuration page area; see Client-Level FTP Options, on page 1705.
- Delete a client profile — Click **Delete** ( ) next to a custom profile.

Step 10 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, cached changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want to generate intrusion events, enable FTP and telnet preprocessor rules (GID 125 and 126). For more information, see Setting Intrusion Rule States, on page 1504.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Firepower System IP Address Conventions, on page 15
Managing Layers, on page 1470
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

The HTTP Inspect Preprocessor

The HTTP Inspect preprocessor is responsible for:

- decoding and normalizing HTTP requests sent to and HTTP responses received from web servers on your network
- separating messages sent to web servers into URI, non-cookie header, cookie header, method, and message body components to improve performance of HTTP-related intrusion rules
- separating messages received from web servers into status code, status message, non-set-cookie header, cookie header, and response body components to improve performance of HTTP-related intrusion rules
- detecting possible URI-encoding attacks
- making the normalized data available for additional rule processing

HTTP traffic can be encoded in a variety of formats, making it difficult for rules to appropriately inspect. HTTP Inspect decodes 14 types of encoding, ensuring that your HTTP traffic gets the best inspection possible.

You can configure HTTP Inspect options globally, on a single server, or for a list of servers.

Note that the preprocessor engine performs HTTP normalization *statelessly*. That is, it normalizes HTTP strings on a packet-by-packet basis, and can only process HTTP strings that have been reassembled by the TCP stream preprocessor.

Global HTTP Normalization Options

The global HTTP options provided for the HTTP Inspect preprocessor control how the preprocessor functions. Use these options to enable or disable HTTP normalization when ports not specified as web server ports receive HTTP traffic.

Note the following:

- If you enable **Unlimited Decompression**, the **Maximum Compressed Data Depth** and **Maximum Decompressed Data Depth** options are automatically set to 65535 when you commit your changes.
- The highest value is used when the values for **Maximum Compressed Data Depth** or **Maximum Decompressed Data Depth** are different in:
  - the default network analysis policy

- any other custom network analysis policy invoked by network analysis rules in the same access control policy

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Detect Anomalous HTTP Servers

Detects HTTP traffic sent to or received by ports not specified as web server ports.



Note If you turn this option on, be sure to list all ports that do receive HTTP traffic in a server profile on the HTTP Configuration page. If you do not, and you enable this option and the accompanying preprocessor rule, normal traffic to and from the server will generate events. The default server profile contains all ports normally used for HTTP traffic, but if you modified that profile, you may need to add those ports to another profile to prevent events from being generated.

You can enable rule 120:1 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Detect HTTP Proxy Servers

Detects HTTP traffic using proxy servers not defined by the **Allow HTTP Proxy Use** option.

You can enable rule 119:17 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Maximum Compressed Data Depth

Sets the maximum size of compressed data to decompress when **Inspect Compressed Data** (and, optionally, **Decompress SWF File (LZMA)**, **Decompress SWF File (Deflate)**, or **Decompress PDF File (Deflate)**) is enabled.

Maximum Decompressed Data Depth

Sets the maximum size of the normalized decompressed data when **Inspect Compressed Data** (and, optionally, **Decompress SWF File (LZMA)**, **Decompress SWF File (Deflate)**, or **Decompress PDF File (Deflate)**) is enabled.

Server-Level HTTP Normalization Options

You can set server-level options for each server you monitor, globally for all servers, or for a list of servers. Additionally, you can use a predefined server profile to set these options, or you can set them individually to meet the needs of your environment. Use these options, or one of the default profiles that set these options, to specify the HTTP server ports whose traffic you want to normalize, the amount of server response payload you want to normalize, and the types of encoding you want to normalize.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Networks

Use this option to specify the IP address of one or more servers. You can specify a single IP address or address block, or a comma-separated list comprised of either or both.

In addition to a limit of up to 255 total profiles, including the default profile, you can include up to 496 characters, or approximately 26 entries, in an HTTP server list, and specify a total of 256 address entries for all server profiles.



Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Note that the `default` setting in the default policy specifies all IP addresses on your monitored network segment that are not covered by another target-based policy. Therefore, you cannot and do not need to specify an IP address or CIDR block/prefix length for the default policy, and you cannot leave this setting blank in another policy or use address notation to represent `any` (for example, 0.0.0.0/0 or ::/0).

Ports

The ports whose HTTP traffic the preprocessor engine normalizes. Separate multiple port numbers with commas.

Oversize Dir Length

Detects URL directories longer than the specified value.

You can enable rule 119:15 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504..

Client Flow Depth

Specifies the number of bytes for rules to inspect in raw HTTP packets, including header and payload data, in client-side HTTP traffic defined in **Ports**. Client flow depth does not apply when HTTP content rule options within a rule inspect specific parts of a request message.

Specify any of the following:

- A positive value inspects the specified number of bytes in the first packet. If the first packet contains fewer bytes than specified, inspect the entire packet. Note that the specified value applies to both segmented and reassembled packets.

  Note also that a value of 300 typically eliminates inspection of large HTTP Cookies that appear at the end of many client request headers.

- 0 inspects all client-side traffic, including multiple packets in a session and exceeding the upper byte limit if necessary. Note that this value is likely to affect performance.

- -1 ignores all client-side traffic.

Server Flow Depth

Specifies the number of bytes for rules to inspect in raw HTTP packets in server-side HTTP traffic specified by **Ports**. Inspection includes the raw header and payload when **Inspect HTTP Responses** disabled and only the raw response body when **Inspect HTTP Response** is enabled.

Server flow depth specifies the number of bytes of raw server response data in a session for rules to inspect in server-side HTTP traffic defined in **Ports**. You can use this option to balance performance and the level of inspection of HTTP server response data. Server flow depth does not apply when HTTP content options within a rule inspect specific parts of a response message.

Unlike client flow depth, server flow depth specifies the number of bytes per HTTP response, not per HTTP request packet, for rules to inspect.

You can specify any of the following:

- A positive value:

  When **Inspect HTTP Responses** is **enabled**, inspects only the raw HTTP response body, and not raw HTTP headers; also inspects decompressed data when **Inspect Compressed Data** is enabled.

  When **Inspect HTTP Responses** is **disabled**, inspects the raw packet header and payload.

  If the session includes fewer response bytes than specified, rules fully inspect all response packets in a given session, across multiple packets as needed. If the session includes more response bytes than specified, rules inspect only the specified number of bytes for that session, across multiple packets as needed.

  Note that a small flow depth value may cause false negatives from rules that target server-side traffic defined in **Ports**. Most of these rules target either the HTTP header or content that is likely to be in the first hundred or so bytes of non-header data. Headers are usually under 300 bytes long, but header size may vary.

  Note also that the specified value applies to both segmented and reassembled packets.

- 0 inspects the entire packet for all HTTP server-side traffic defined in **Ports**, including response data in a session that exceeds 65535 bytes.

  Note that this value is likely to affect performance.

- -1:

  When **Inspect HTTP Responses** is **enabled**, inspects only raw HTTP headers and not the raw HTTP response body.

  When **Inspect HTTP Responses** is **disabled**, ignores all server-side traffic defined in **Ports.**

Maximum Header Length

Detects a header field longer than the specified maximum number of bytes in an HTTP request; also in HTTP responses when **Inspect HTTP Responses** is enabled. A value of 0 disables this option. Specify a positive value to enable it.

You can enable rule 119:19 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504..

Maximum Number of Headers

Detects when the number of headers exceeds this setting in an HTTP request. A value of 0 disables this option. Specify a positive value to enable it.

You can enable rule 119:20 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504..

Maximum Number of Spaces

Detects when the number of white spaces in a folded line equals or exceeds this setting in an HTTP request. A value of 0 disables this option. Specify a positive value to enable it.

You can enable rule 119:26 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504..

HTTP Client Body Extraction Depth

Specifies the number of bytes to extract from the message body of an HTTP client request. You can use an intrusion rule to inspect the extracted data by selecting the `content` or `protected_content` keyword **HTTP Client Body** option.

Specify -1 to ignore the client body. Specify 0 to extract the entire client body. Note that identifying specific bytes to extract can improve system performance. Note also that you must specify a value greater than or equal to 0 for the **HTTP Client Body** option to function in an intrusion rule.

Small Chunk Size

Specifies the maximum number of bytes at which a chunk is considered small. Specify a positive value. A value of 0 disables detection of anomalous consecutive small segments. See the **Consecutive Small Chunks** option for more information.

Consecutive Small Chunks

Specifies how many consecutive small chunks represent an abnormally large number in client or server traffic that uses chunked transfer encoding. The **Small Chunk Size** option specifies the maximum size of a small chunk.

For example, set **Small Chunk Size** to 10 and **Consecutive Small Chunks** to 5 to detect 5 consecutive chunks of 10 bytes or less.

You can enable preprocessor rule 119:27 to generate events and, in an inline deployment, drop offending packets on excessive small chunks in client traffic, and rule 120:7 in server traffic. When **Small Chunk Size** is enabled and this option is set to 0 or 1, enabling these rules would trigger an event on every chunk of the specified size or less.

HTTP Methods

Specifies HTTP request methods in addition to GET and POST that you expect the system to encounter in traffic. Use a comma to separate multiple values.

Intrusion rules use the `content` or `protected_content` keyword with the **HTTP Method** argument to search for content in HTTP methods. You can enable rule 119:31to generate events and, in an inline deployment, drop offending packets when a method other than GET, POST, or a method configured for this option is encountered in traffic. See Setting Intrusion Rule States, on page 1504.

No Alerts

Disables intrusion events when accompanying preprocessor rules are enabled.



Note This option does **not** disable HTTP standard text rules and shared object rules.

Normalize HTTP Headers

When **Inspect HTTP Responses** is enabled, enables normalization of non-cookie data in request and response headers. When **Inspect HTTP Responses** is **not** enabled, enables normalization of the entire HTTP header, including cookies, in request and response headers.

Inspect HTTP Cookies

Enables extraction of cookies from HTTP request headers. Also enables extraction of set-cookie data from response headers when **Inspect HTTP Responses** is enabled. Disabling this option when cookie extraction is not required can improve performance.

Note that the `Cookie:` and `Set-Cookie:` header names, leading spaces on the header line, and the `CRLF` that terminates the header line are inspected as part of the header and not as part of the cookie.

Normalize Cookies in HTTP headers

Enables normalization of cookies in HTTP request headers. When **Inspect HTTP Responses** is enabled, also enables normalization of set-cookie data in response headers. You must select **Inspect HTTP Cookies** before selecting this options.

Allow HTTP Proxy Use

Allows the monitored web server to be used as an HTTP proxy. This option is used only in the inspection of HTTP requests.

Inspect URI Only

Inspects only the URI portion of the normalized HTTP request packet.

Inspect HTTP Responses

Enables extended inspection of HTTP responses so, in addition to decoding and normalizing HTTP request messages, the preprocessor extracts response fields for inspection by the rules engine. Enabling this option causes the system to extract the response header, body, status code, and so on, and also extracts set-cookie data when **Inspect HTTP Cookies** is enabled.

You can enable rules 120:2 and 120:3 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504..

Normalize UTF Encodings to UTF-8

When **Inspect HTTP Responses** is enabled, detects UTF-16LE, UTF-16BE, UTF-32LE, and UTF32-BE encodings in HTTP responses and normalizes them to UTF-8.

You can enable rule 120:4 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Inspect Compressed Data

When **Inspect HTTP Responses** is enabled, enables decompression of gzip and deflate-compatible compressed data in the HTTP response body, and inspection of the normalized decompressed data. The system inspects chunked and non-chunked HTTP response data. The system inspects decompressed data packet by packet across multiple packets as needed; that is, the system does not combine the decompressed data from different packets for inspection. Decompression ends when **Maximum Compressed Data Depth**, **Maximum Decompressed Data Depth**, or the end of the compressed data is reached. Inspection of decompressed data ends when **Server Flow Depth** is reached unless you also select **Unlimited Decompression**. You can use the `file_data` rule keyword to inspect decompressed data.

You can enable rule 120:6 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504..

Unlimited Decompression

When **Inspect Compressed Data** (and, optionally, **Decompress SWF File (LZMA)**, **Decompress SWF File (Deflate)**, or **Decompress PDF File (Deflate)**) is enabled, overrides **Maximum Decompressed Data Depth** across multiple packets; that is, this option enables unlimited decompression across multiple packets. Note that enabling this option does not affect **Maximum Compressed Data Depth** or **Maximum Decompressed Data Depth** within a single packet. Note also that enabling this option sets **Maximum Compressed Data Depth** and **Maximum Decompressed Data Depth** to 65535 when you commit your changes.

Normalize Javascript

When **Inspect HTTP Responses** is enabled, enables detection and normalization of Javascript within the HTTP response body. The preprocessor normalizes obfuscated Javascript data such as the unescape and decodeURI functions and the String.fromCharCode method. The preprocessor normalizes the following encodings within the unescape, decodeURI, and decodeURIComponent functions:

- %XX
- %uXXXX
- 0xXX
- \xXX
- \uXXXX

The preprocessor detects consecutive white spaces and normalizes them into a single space. When this option is enabled, a configuration field allows you to specify the maximum number of consecutive white spaces to permit in obfuscated Javascript data. You can enter a value from 1 to 65535. The value 0 disables event generation, regardless of whether the preprocessor rule (120:10) associated with this field is enabled.

The preprocessor also normalizes the Javascript plus (+) operator and concatenates strings using the operator.

You can use the `file_data` intrusion rule keyword to point intrusion rules to the normalized Javascript data.

You can enable rules 120:9, 120:10, and 120:11 to generate events and, in an inline deployment, drop offending packets, as follows:

Table 197: Normalize Javascript Option Rules

This rule...	Triggers when...
120:9	the obfuscation level within the preprocessor is greater than or equal to 2.

This rule...	Triggers when...
120:10	the number of consecutive white spaces in the Javascript obfuscated data is greater than or equal to the value configured for the maximum number of consecutive white spaces allowed.
120:11	escaped or encoded data includes more than one type of encoding.

Decompress SWF File (LZMA) and Decompress SWF File (Deflate)

When **HTTP Inspect Responses** is enabled, these options decompress the compressed portions of files located within the HTTP response body of HTTP requests.

Note You can **only** decompress the compressed portions of files found in HTTP GET responses.

- **Decompress SWF File (LZMA)** decompresses the LZMA-compatible compressed portions of Adobe ShockWave Flash (`.swf`) files
- **Decompress SWF File (Deflate)** decompresses the deflate-compatible compressed portions of Adobe ShockWave Flash (`.swf`) files

Decompression ends when **Maximum Compressed Data Depth**, **Maximum Decompressed Data Depth**, or the end of the compressed data is reached. Inspection of decompressed data ends when **Server Flow Depth** is reached unless you also select **Unlimited Decompression**. You can use the `file_data` intrusion rule keyword to inspect decompressed data.

You can enable rules 120:12 and 120:13 to generate events and, in an inline deployment, drop offending packets, as follows:

Table 198: Decompress SWF File Option Rules

This rule...	Triggers when...
120:12	deflate file decompression fails.
120:13	LZMA file decompression fails.

Decompress PDF File (Deflate)

When **HTTP Inspect Responses** is enabled, **Decompress PDF File (Deflate)** decompresses the deflate-compatible compressed portions of Portable Document Format (`.pdf`) files located within the HTTP response body of HTTP requests. The system can only decompress PDF files with the `/FlateDecode` stream filter. Other stream filters (including `/FlateDecode /FlateDecode`) are unsupported.

Note You can **only** decompress the compressed portions of files found in HTTP GET responses.

Decompression ends when **Maximum Compressed Data Depth**, **Maximum Decompressed Data Depth**, or the end of the compressed data is reached. Inspection of decompressed data ends when **Server Flow Depth**

is reached unless you also select **Unlimited Decompression**. You can use the `file_data` intrusion rule keyword to inspect decompressed data.

You can enable rules 120:14, 120:15, 120:16, and 120:17 to generate events and, in an inline deployment, drop offending packets, as follows:

Table 199: Decompress PDF File (Deflate) Option Rules

This rule...	Triggers when...
120:14	file decompression fails.
120:15	file decompression fails due to an unsupported compression type.
120:16	file decompression fails due to an unsupported PDF stream filter.
120:17	file parsing fails.

Extract Original Client IP Address

Enables the examination of original client IP addresses during intrusion inspection. The system extracts the original client IP address from the X-Forwarded-For (XFF), True-Client-IP, or custom HTTP headers you define in the **XFF Header Priority** option. You can view the extracted original client IP address in the intrusion events table.

You can enable rules 119:23, 119:29, and 119:30 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504..

XFF Header Priority

Specifies the order in which the system processes original client IP headers when multiple headers are present in an HTTP request. By default, the system examines X-Forwarded-For (XFF) headers, then True-Client-IP headers. Use the up and down arrow icons beside each header type to adjust its priority.

This option also allows you to specify original client IP headers other than XFF or True-Client-IP for extraction and evaluation. Click **Add** to add custom header names to the priority list. The system only supports custom headers that use the same syntax as an XFF or True-Client-IP header.

Keep in mind the following when configuring this option:

- The system uses this priority order when evaluating original client IP address headers for both access control and intrusion inspection.
- If multiple original client IP headers are present, the system processes only the header with the highest priority.
- The XFF header contains a list of IP addresses, which represent the proxy servers through which the request has passed. To prevent spoofing, the system uses the last IP address in the list (that is, the address appended by the trusted proxy) as the original client IP address.

Log URI

Enables extraction of the raw URI, if present, from HTTP request packets and associates the URI with all intrusion events generated for the session.

When this option is enabled, you can display the first fifty characters of the extracted URI in the HTTP URI column of the intrusion events table view. You can display the complete URI, up to 2048 bytes, in the packet view.

Log Hostname

Enables extraction of the host name, if present, from the HTTP request Host header and associates the host name with all intrusion events generated for the session. When multiple Host headers are present, extracts the host name from the first header.

When this option is enabled, you can display the first fifty characters of the extracted host name in the HTTP Hostname column of the intrusion events table view. You can display the complete host name, up to 256 bytes, in the packet view.

You can enable rule 119:25 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Note that, when enabled, rule 119:24 triggers if it detects multiple Host headers in an HTTP request, regardless of the setting for this option.

Profile

Specifies the types of encoding that are normalized for HTTP traffic. The system provides a default profile appropriate for most servers, default profiles for Apache servers and IIS servers, and custom default settings that you can tailor to meet the needs of your monitored traffic:

- Select **All** to use the standard default profile, appropriate for all servers.
- Select **IIS** to use the system-provided IIS profile.
- Select **Apache** to use the system-provided Apache profile.
- Select **Custom** to create your own server profile.

Server-Level HTTP Normalization Encoding Options

When you set the HTTP server-level **Profile** option to `Custom`, you can specify the types of encoding that are normalized for HTTP traffic, and enable HTTP preprocessor rules to generate events against traffic containing the different encoding types.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

ASCII Encoding

Decodes encoded ASCII characters and specifies whether the rules engine generates an event on ASCII-encoded URIs.

You can enable rule 119:1 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

UTF-8 Encoding

Decodes standard UTF-8 Unicode sequences in the URI.

You can enable rule 119:6 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Microsoft %U Encoding

Decodes the IIS %u encoding scheme that uses %u followed by four characters where the 4 characters are a hex encoded value that correlates to an IIS Unicode codepoint.



Tip Legitimate clients rarely use %u encodings, so Cisco recommends decoding HTTP traffic encoded with %u encodings.

You can enable rule 119:3 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Bare Byte UTF-8 Encoding

Decodes bare byte encoding, which uses non-ASCII characters as valid values in decoding UTF-8 values.



Tip Bare byte encoding allows the user to emulate an IIS server and interpret non-standard encodings correctly. Cisco recommends enabling this option because no legitimate clients encode UTF-8 this way.

You can enable rule 119:4 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Microsoft IIS Encoding

Decodes using Unicode codepoint mapping.



Tip Cisco recommends enabling this option, because it is seen mainly in attacks and evasion attempts.

You can enable rule 119:7 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Double Encoding

Decodes IIS double encoded traffic by making two passes through the request URI performing decodes in each one. Cisco recommends enabling this option because it is usually found only in attack scenarios.

You can enable rule 119:2 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Multi-Slash Obfuscation

Normalizes multiple slashes in a row into a single slash.

You can enable rule 119:8 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

IIS Backslash Obfuscation

Normalizes backslashes to forward slashes.

You can enable rule 119:9 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Directory Traversal

Normalizes directory traversals and self-referential directories. If you enable the accompanying preprocessor rules to generate events against this type of traffic, it may generate false positives because some web sites refer to files using directory traversals.

You can enable rules 119:10 and 119:11 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Tab Obfuscation

Normalizes the non-RFC standard of using a tab for a space delimiter. Apache and other non-IIS web servers use the tab character (0x09) as a delimiter in URLs.



Note Regardless of the configuration for this option, the HTTP Inspect preprocessor treats a tab as white space if a space character (0x20) precedes it.

You can enable rule 119:12 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Invalid RFC Delimiter

Normalizes line breaks (\n) in URI data.

You can enable rule 119:13 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Webroot Directory Traversal

Detects directory traversals that traverse past the initial directory in the URL.

You can enable rule 119:18 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Tab URI Delimiter

Turns on the use of the tab character (0x09) as a delimiter for a URI. Apache, newer versions of IIS, and some other web servers use the tab character as a delimiter in URLs.



Note Regardless of the configuration for this option, the HTTP Inspect preprocessor treats a tab as white space if a space character (0x20) precedes it.

Non-RFC characters

Detects the non-RFC character list you add in the corresponding field when it appears within incoming or outgoing URI data. When modifying this field, use the hexadecimal format that represents the byte character.

If and when you configure this option, set the value with care. Using a character that is very common may overwhelm you with events.

You can enable rule 119:14 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Max Chunk Encoding Size

Detects abnormally large chunk sizes in URI data.

You can enable rules 119:16 and 119:22 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Disable Pipeline Decoding

Disables HTTP decoding for pipelined requests. When this option is disabled, performance is enhanced because HTTP requests waiting in the pipeline are not decoded or analyzed, and are only inspected using generic pattern matching.

Non-Strict URI Parsing

Enables non-strict URI parsing. Use this option only on servers that will accept non-standard URIs in the format "GET /index.html abc xo qr \n". Using this option, the decoder assumes that the URI is between the first and second space, even if there is no valid HTTP identifier after the second space.

Extended ASCII Encoding

Enables parsing of extended ASCII characters in an HTTP request URI. Note that this option is available in custom server profiles only, and not in the default profiles provided for Apache, IIS, or all servers.

Related Topics

Overview: HTTP content and protected_content Keyword Arguments, on page 1571
Firepower System IP Address Conventions, on page 15

Configuring The HTTP Inspect Preprocessor

The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Before you begin

- Confirm that any networks you want to identify in a custom target-based policy match or are a subset of the networks, zones, and VLANs handled by its parent network analysis policy. See Advanced Settings for Network Analysis Policies, on page 1672 for more information.

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation panel.

Step 4 If **HTTP Configuration** under **Application Layer Preprocessors** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **HTTP Configuration**.

Step 6 Modify the options in the Global Settings page area; see Global HTTP Normalization Options, on page 1708.

Step 7 You have three choices:

- Add a server profile — Click **Add** ( ) in the **Servers** section. Specify one or more IP addresses for the client in the **Server Address** field, and click **OK**. You can specify a single IP address or address block, or a comma-separated list of either or both. You can include up to 496 characters in a list, specify a total of 256 address entries for all server profiles, and create a total of 255 profiles including the default profile.
- Edit a server profile — Click the configured address for a profile you have added under **Servers**, or click **default**. You can modify any of the settings in the **Configuration** section; see Server-Level HTTP Normalization Options, on page 1709. If you choose **Custom** for the **Profile** value, you can also modify the encoding options described in Server-Level HTTP Normalization Encoding Options, on page 1717.
- Delete a server profile — Click **Delete** ( ) next to a custom profile.

Step 8 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, cached changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want generate events and, in an inline deployment, drop offending packets, enable HTTP preprocessor rules (GID 119). For more information, see Setting Intrusion Rule States, on page 1504.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Managing Layers, on page 1470
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

Additional HTTP Inspect Preprocessor Rules

You can enable the rules in the **Preprocessor Rule GID:SID** column of the following table to generate events for HTTP Inspect preprocessor rules that are not associated with specific configuration options.

Table 200: Additional HTTP Inspect Preprocessor Rules

Preprocessor Rule GID:SID	Triggers when...
119:21	an HTTP request header has more than one `content-length` field.
119:24	an HTTP request has more than one Host header.
119:28	an HTTP POST method has neither a `content-length` header nor chunked `transfer-encoding`.
119:32	HTTP version 0.9 is encountered in traffic. Note that the TCP stream configuration must also be enabled.
119:33	an HTTP URI includes an unescaped space.
119:34	a TCP connection contains 24 or more pipelined HTTP requests.
120:5	UTF-7 encoding is encountered in HTTP response traffic; UTF-7 should only appear where 7-bit parity is required, such as in SMTP traffic.
120:8	the `content-length` or chunk size is invalid.

The Sun RPC Preprocessor

Remote Procedure Call (RPC) normalization takes fragmented RPC records and normalizes them to a single record so the rules engine can inspect the complete record. For example, an attacker may attempt to discover the port where RPC `admind` runs. Some UNIX hosts use RPC `admind` to perform remote distributed system tasks. If the host performs weak authentication, a malicious user could take control of remote administration. The standard text rule (GID: 1) with the Snort ID (SID) 575 detects this attack by searching for content in specific locations to identify inappropriate `portmap GETPORT` requests.

Sun RPC Preprocessor Options

Ports

Specify the ports whose traffic you want to normalize. In the interface, list multiple ports separated by commas. Typical RPC ports are 111 and 32771. If your network sends RPC traffic to other ports, consider adding them.

Detect fragmented RPC records

Detects RPC fragmented records.

You can enable rules 106:1 and 106:5 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Detect multiple records in one packet

Detects more than one RPC request per packet (or reassembled packet).

You can enable rule 106:2 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Detect fragmented record sums which exceed one fragment

Detects reassembled fragment record lengths that exceed the current packet length.

You can enable rule 106:3 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Detect single fragment records which exceed the size of one packet

Detects partial records

You can enable rule 106:4 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Configuring the Sun RPC Preprocessor

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation panel.

Step 4 If **Sun RPC Configuration** under **Application Layer Preprocessors** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **Sun RPC Configuration**.

Step 6 Modify the settings described in Sun RPC Preprocessor Options, on page 1722.

Step 7 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, cached changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want to generate events and, in an inline deployment, drop offending packets, enable Sun RPC preprocessor rules (GID 106). For more information, see Setting Intrusion Rule States, on page 1504.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Managing Layers, on page 1470
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

The SIP Preprocessor

The Session Initiation Protocol (SIP) provides call setup, modification, and teardown of one or more sessions for one or more users of client applications such as Internet telephony, multimedia conferencing, instant messaging, online gaming, and file transfer. A *method* field in each SIP request identifies the purpose of the request, and a Request-URI specifies where to send the request. A status code in each SIP response indicates the outcome of the requested action.

After calls are set up using SIP, the Real-time Transport Protocol (RTP) is responsible for subsequent audio and video communication; this part of the session is sometimes referred to as the call channel, the data channel, or the audio/video data channel. RTP uses the Session Description Protocol (SDP) within the SIP message body for data-channel parameter negotiation, session announcement, and session invitation.

The SIP preprocessor is responsible for:

- decoding and analyzing SIP 2.0 traffic
- extracting the SIP header and message body, including SDP data when present, and passing the extracted data to the rules engine for further inspection
- generating events when the following conditions are detected and the corresponding preprocessor rules are enabled:
  - anomalies and known vulnerabilities in SIP packets
  - out-of-order and invalid call sequences
- optionally, ignoring the call channel

The preprocessor identifies the RTP channel based on the port identified in the SDP message, which is embedded in the SIP message body, but the preprocessor does not provide RTP protocol inspection.

Note the following when using the SIP preprocessor:

- UDP typically carries media sessions supported by SIP. UDP stream preprocessing provides SIP session tracking for the SIP preprocessor.
- SIP rule keywords allow you to point to the SIP packet header or message body and to limit detection to packets for specific SIP methods or status codes.

SIP Preprocessor Options

For the following options, you can specify a positive value from 1 to 65535 bytes, or 0 to disable event generation for the option regardless of whether the associated rule is enabled.

- **Maximum Request URI Length**
- **Maximum Call ID Length**
- **Maximum Request Name Length**

- **Maximum From Length**
- **Maximum To Length**
- **Maximum Via Length**
- **Maximum Contact Length**
- **Maximum Content Length**

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Ports

Specifies the ports to inspect for SIP traffic. You can specify an integer from 0 to 65535. Separate multiple port numbers with commas.

Methods to Check

Specifies SIP methods to detect. You can specify any of the following currently defined SIP methods:

```
ack, benotify, bye, cancel, do, info, invite, join, message,
notify, options, prack, publish, quath, refer, register,
service, sprack, subscribe, unsubscribe, update
```

Methods are case-insensitive. The method name can include alphabetic characters, numbers, and the underscore character. No other special characters are permitted. Separate multiple methods with commas.

Because new SIP methods might be defined in the future, your configuration can include an alphabetic string that is not currently defined. The system supports up to 32 methods, including the 21 currently defined methods and an additional 11 methods. The system ignores any undefined methods that you might configure.

Note that, in addition to any methods you specify for this option, the 32 total methods includes methods specified using the `sip_method` keyword in intrusion rules.

Maximum Dialogs within a Session

Specifies the maximum number of dialogs allowed within a stream session. If more dialogs than this number are created, the oldest dialogs are dropped until the number of dialogs does not exceed the maximum number specified. You can specify an integer from 1 to 4194303.

You can enable rule 140:27 to generate events and, in an inline deployment, drop offending packets for this option. See Setting Intrusion Rule States, on page 1504.

Maximum Request URI Length

Specifies the maximum number of bytes to allow in the Request-URI header field. A Longer URI generates an event and, in an inline deployment, drops offending packets when rule 140:3 is enabled. The request URI field indicates the destination path or page for the request.

Maximum Call ID Length

Specifies the maximum number of bytes to allow in the request or response Call-ID header field. A longer Call-ID generates an event and, in an inline deployment, drops offending packets when rule 140:5 is enabled. The Call-ID field uniquely identifies the SIP session in requests and responses.

Maximum Request Name Length

Specifies the maximum number of bytes to allow in the request name, which is the name of the method specified in the CSeq transaction identifier. A longer request name generates an event and, in an inline deployment, drops offending packets when rule 140:7 is enabled.

Maximum From Length

Specifies the maximum number of bytes to allow in the request or response From header field. A longer From generates an event and, in an inline deployment, drops offending packets when rule 140:9 is enabled. The From field identifies the message initiator.

Maximum To Length

Specifies the maximum number of bytes to allow in the request or response To header field. A longer To generates an event and, in an inline deployment, drops offending packets when rule 140:11 is enabled. The To field identifies the message recipient.

Maximum Via Length

Specifies the maximum number of bytes to allow in the request or response Via header field. A longer Via generates an event and, in an inline deployment, drops offending packets when rule 140:13 is enabled. The Via field provides the path followed by the request and, in a response, receipt information.

Maximum Contact Length

Specifies the maximum number of bytes to allow in the request or response Contact header field. A longer Contactgenerates an event and, in an inline deployment, drops offending packets when rule 140:15 is enabled. The Contact field provides a URI that specifies the location to contact with subsequent messages.

Maximum Content Length

Specifies the maximum number of bytes to allow in the content of the request or response message body. Longer content generates an event and, in an inline deployment, drops offending packets when rule 140:16 is enabled.

Ignore Audio/Video Data Channel

Enables and disables inspection of data channel traffic. Note that the preprocessor continues inspection of other non-data-channel SIP traffic when you enable this option.

Related Topics

SIP Keywords, on page 1617

Configuring the SIP Preprocessor

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation panel.

Step 4 If **SIP Configuration** under **Application Layer Preprocessors** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **SIP Configuration**.

Step 6 Modify the options described in SIP Preprocessor Options, on page 1724.

Step 7 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, cached changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want to generate events and, in an inline deployment, drop offending packets, enable SIP preprocessor rules (GID 140). For more information, see Setting Intrusion Rule States, on page 1504.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Managing Layers, on page 1470
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

Additional SIP Preprocessor Rules

The SIP preprocessor rules in the following table are not associated with specific configuration options. As with other SIP preprocessor rules, you must enable these rules if you want them to generate events and, in an inline deployment, drop offending packets.

Table 201: Additional SIP Preprocessor Rules

Preprocessor Rule GID:SID	Triggers when...
140:1	the preprocessor is monitoring the maximum number of SIP sessions allowed by the system.
140:2	the required Request_URI field is empty in a SIP request.
140:4	the Call-ID header field is empty in a SIP request or response.
140:6	the value for the sequence number in the SIP request or response CSeq field is not a 32-bit unsigned integer less than 231.

Preprocessor Rule GID:SID	Triggers when...
140:8	the From header field is empty in a SIP request or response.
140:10	the To header field is empty in a SIP request or response.
140:12	the Via header field is empty in a SIP request or response
140:14	the required Contact header field is empty in a SIP request or response.
140:17	a single SIP request or response packet in UDP traffic contains multiple messages. Note that older SIP versions supported multiple messages, but SIP 2.0 supports only one message per packet.
140:18	the actual length of the message body in a SIP request or response in UDP traffic does not match the value specified in the Content-Length header field in a SIP request or response.
140:19	the preprocessor does not recognize a method name in the CSeq field of a SIP response.
140:20	the SIP server does not challenge an authenticated invite message. Note that this occurs in the case of the InviteReplay billing attack.
140:21	session information changes before the call is set up. Note that this occurs in the case of the FakeBusy billing attack.
140:22	the response status code is not a three-digit number.
140:23	the Content-Type header field does not specify a content type and the message body contains data.
140:24	the SIP version is not 1, 1.1, or 2.0.
140:25	the method specified in the CSeq header and the method field do not match in a SIP request.
140:26	the preprocessor does not recognize the method named in the SIP request method field.

The GTP Preprocessor

The General Service Packet Radio (GPRS) Tunneling Protocol (GTP) provides communication over a GTP core network. The GTP preprocessor detects anomalies in GTP traffic and forwards command channel signaling messages to the rules engine for inspection. You can use the `gtp_version`, `gtp_type`, and `gtp_info` rule keywords to inspect GTP command channel traffic for exploits.

A single configuration option allows you to modify the default setting for the ports that the preprocessor inspects for GTP command channel messages.

GTP Preprocessor Rules

You must enable the GTP preprocessor rules in the following table if you want them to generate events and, in an inline deployment, drop offending packets.

Table 202: GTP Preprocessor Rules

Preprocessor Rule GID:SID	Description
143:1	Generates an event when the preprocessor detects an invalid message length.
143:2	Generates an event when the preprocessor detects an invalid information element length.
143:3	Generates an event when the preprocessor detects information elements that are out of order.

Configuring the GTP Preprocessor

You can use this procedure to modify the ports the GTP preprocessor monitors for GTP command messages.

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation panel on the left.

Step 4 If **GTP Command Channel Configuration** under **Application Layer Preprocessors** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **GTP Command Channel Configuration**.

Step 6 Enter a **Ports** value.

Separate multiple ports with commas.

Step 7 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, cached changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want to enable intrusion events, enable GTP preprocessor rules (GID 143). For more information, see Setting Intrusion Rule States, on page 1504.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

The IMAP Preprocessor

The Internet Message Application Protocol (IMAP) is used to retrieve email from a remote IMAP server. The IMAP preprocessor inspects server-to-client IMAP4 traffic and, when associated preprocessor rules are enabled, generates events on anomalous traffic. The preprocessor can also extract and decode email attachments in client-to-server IMAP4 traffic and send the attachment data to the rules engine. You can use the `file_data` keyword in an intrusion rule to point to the attachment data.

Extraction and decoding include multiple attachments, when present, and large attachments that span multiple packets.

IMAP Preprocessor Options

Note that decoding, or extraction when the MIME email attachment does not require decoding, includes multiple attachments when present, and large attachments that span multiple packets.

Note also that the highest value is used when the values for the **Base64 Decoding Depth**, **7-Bit/8-Bit/Binary Decoding Depth**, **Quoted-Printable Decoding Depth**, or **Unix-to-Unix Decoding Depth** options are different in:

- the default network analysis policy
- any other custom network analysis policy invoked by network analysis rules in the same access control policy



Caution Changing the value for **Base64 Decoding Depth**, **7-Bit/8-Bit/Binary Decoding Depth**, **Quoted-Printable Decoding Depth**, or **Unix-to-Unix Decoding Depth** restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Ports

Specifies the ports to inspect for IMAP traffic. You can specify an integer from 0 to 65535. Separate multiple port numbers with commas.

Base64 Decoding Depth

Specifies the maximum number of bytes to extract and decode from each Base64 encoded MIME email attachment. You can specify a positive value, or specify 0 to decode all the Base64 data. Specify -1 to ignore Base64 data.

Note that positive values not divisible by 4 are rounded up to the next multiple of 4 except for the values 65533, 65534, and 65535, which are rounded down to 65532.

When this option is enabled, you can enable rule 141:4 generate events and, in an inline deployment, drop offending packets when decoding fails; decoding could fail, for example, because of incorrect encoding or corrupted data.

7-Bit/8-Bit/Binary Decoding Depth

Specifies the maximum bytes of data to extract from each MIME email attachment that does not require decoding. These attachment types include 7-bit, 8-bit, binary, and various multipart content types such as plain text, jpeg images, mp3 files, and so on. You can specify a positive value, or specify 0 to extract all data in the packet. Specify -1 to ignore non-decoded data.

When this option is enabled, you can enable rule 141:6 to generate events and, in an inline deployment, drop offending packets when extraction fails; extraction could fail, for example, because of corrupted data.

Quoted-Printable Decoding Depth

Specifies the maximum number of bytes to extract and decode from each quoted-printable (QP) encoded MIME email attachment. You can specify a positive value, or specify 0 to decode all QP encoded data in the packet. Specify -1 to ignore QP encoded data.

When this option is enabled, you can enable rule 141:5 to generate events and, in an inline deployment, drop offending packets when decoding fails; decoding could fail, for example, because of incorrect encoding or corrupted data.

Unix-to-Unix Decoding Depth

Specifies the maximum number of bytes to extract and decode from each Unix-to-Unix encoded (uuencoded) email attachment. You can specify a positive value, or specify 0 to decode all uuencoded data in the packet. Specify -1 to ignore uuencoded data.

When this option is enabled, you can enable rule 141:7 to generate events and, in an inline deployment, drop offending packets when decoding fails; decoding could fail, for example, because of incorrect encoding or corrupted data.

Related Topics

The file_data Keyword, on page 1651

Configuring the IMAP Preprocessor

Caution Changing the value for **Base64 Decoding Depth**, **7-Bit/8-Bit/Binary Decoding Depth**, **Quoted-Printable Decoding Depth**, or **Unix-to-Unix Decoding Depth** restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation panel.

Step 4 If **IMAP Configuration** under **Application Layer Preprocessors** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **IMAP Configuration**.

Step 6 Modify the settings described in IMAP Preprocessor Options, on page 1730.

Step 7 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, cached changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want to enable intrusion events, enable IMAP preprocessor rules (GID 141); see Setting Intrusion Rule States, on page 1504.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Layers in Intrusion and Network Analysis Policies, on page 1463
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

Additional IMAP Preprocessor Rules

The IMAP preprocessor rules in the following table are not associated with specific configuration options. As with other IMAP preprocessor rules, you must enable these rules if you want them to generate events and, in an inline deployment, drop offending packets.

Table 203: Additional IMAP Preprocessor Rules

Preprocessor Rule GID:SID	Description
141:1	Generates an event when the preprocessor detects a client command that is not defined in RFC 3501.

Preprocessor Rule GID:SID	Description
141:2	Generates an event when the preprocessor detects a server response that is not defined in RFC 3501.
141:3	Generates an event when the preprocessor is using the maximum amount of memory allowed by the system. At this point, the preprocessor stops decoding until memory becomes available.

The POP Preprocessor

The Post Office Protocol (POP) is used to retrieve email from a remote POP mail server. The POP preprocessor inspects server-to-client POP3 traffic and, when associated preprocessor rules are enabled, generates events on anomalous traffic. The preprocessor can also extract and decode email attachments in client-to-server POP3 traffic and send the attachment data to the rules engine. You can use the `file_data` keyword in an intrusion rule to point to attachment data.

Extraction and decoding include multiple attachments, when present, and large attachments that span multiple packets.

POP Preprocessor Options

Note that decoding, or extraction when the MIME email attachment does not require decoding, includes multiple attachments when present, and large attachments that span multiple packets.

Note also that the highest value is used when the values for the **Base64 Decoding Depth**, **7-Bit/8-Bit/Binary Decoding Depth**, **Quoted-Printable Decoding Depth**, or **Unix-to-Unix Decoding Depth** options are different in:

- the default network analysis policy
- any other custom network analysis policy invoked by network analysis rules in the same access control policy



Caution Changing the value for **Base64 Decoding Depth**, **7-Bit/8-Bit/Binary Decoding Depth**, **Quoted-Printable Decoding Depth**, or **Unix-to-Unix Decoding Depth** restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Ports

Specifies the ports to inspect for POP traffic. You can specify an integer from 0 to 65535. Separate multiple port numbers with commas.

Base64 Decoding Depth

Specifies the maximum number of bytes to extract and decode from each Base64 encoded MIME email attachment. You can specify a positive value, or specify 0 to decode all the Base64 data. Specify -1 to ignore Base64 data.

Note that positive values not divisible by 4 are rounded up to the next multiple of 4 except for the values 65533, 65534, and 65535, which are rounded down to 65532.

When this option is enabled, you can enable rule 142:4 to generate an event and, in an inline deployment, drop offending packets when decoding fails; decoding could fail, for example, because of incorrect encoding or corrupted data.

7-Bit/8-Bit/Binary Decoding Depth

Specifies the maximum bytes of data to extract from each MIME email attachment that does not require decoding. These attachment types include 7-bit, 8-bit, binary, and various multipart content types such as plain text, jpeg images, mp3 files, and so on. You can specify a positive value, or specify 0 to extract all data in the packet. Specify -1 to ignore non-decoded data.

When this option is enabled, you can enable rule 142:6 to generate an event and, in an inline deployment, drop offending packets when extraction fails; extraction could fail, for example, because of corrupted data.

Quoted-Printable Decoding Depth

Specifies the maximum number of bytes to extract and decode from each quoted-printable (QP) encoded MIME email attachment. You can specify a positive value, or specify 0 to decode all QP encoded data in the packet. Specify -1 to ignore QP encoded data.

When this option is enabled, you can enable rule 142:5 to generate an event and, in an inline deployment, drop offending packets when decoding fails; decoding could fail, for example, because of incorrect encoding or corrupted data.

Unix-to-Unix Decoding Depth

Specifies the maximum number of bytes to extract and decode from each Unix-to-Unix encoded (uuencoded) email attachment. You can specify a positive value, or specify 0 to decode all uuencoded data in the packet. Specify -1 to ignore uuencoded data.

When this option is enabled, you can enable rule 142:7 to generate an event and, in an inline deployment, drop offending packets when decoding fails; decoding could fail, for example, because of incorrect encoding or corrupted data.

Related Topics

Managing Layers, on page 1470
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461
The file_data Keyword, on page 1651

Configuring the POP Preprocessor

Caution Changing the value for **Base64 Decoding Depth**, **7-Bit/8-Bit/Binary Decoding Depth**, **Quoted-Printable Decoding Depth**, or **Unix-to-Unix Decoding Depth** restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation panel.

Step 4 If **POP Configuration** under **Application Layer Preprocessors** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **POP Configuration**.

Step 6 Modify the settings described in POP Preprocessor Options, on page 1733.

Step 7 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, cached changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want to enable intrusion events, enable POP preprocessor rules (GID 142). For more information, see Setting Intrusion Rule States, on page 1504.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Managing Layers, on page 1470
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

Additional POP Preprocessor Rules

The POP preprocessor rules in the following table are not associated with specific configuration options. As with other POP preprocessor rules, you must enable these rules if you want them to generate events and, in an inline deployment, drop offending packets.

Table 204: Additional POP Preprocessor Rules

Preprocessor Rule GID:SID	Description
142:1	Generates an event when the preprocessor detects a client command that is not defined in RFC 1939.
142:2	Generates an event when the preprocessor detects a server response that is not defined in RFC 1939.
142:3	Generates an event when the preprocessor is using the maximum amount of memory allowed by the system. At this point, the preprocessor stops decoding until memory becomes available.

The SMTP Preprocessor

The SMTP preprocessor instructs the rules engine to normalize SMTP commands. The preprocessor can also extract and decode email attachments in client-to-server traffic and, depending on the software version, extract email file names, addresses, and header data to provide context when displaying intrusion events triggered by SMTP traffic.

SMTP Preprocessor Options

You can enable or disable normalization, and you can configure options to control the types of anomalous traffic the SMTP decoder detects.

Note that decoding, or extraction when the MIME email attachment does not require decoding, includes multiple attachments when present, and large attachments that span multiple packets.

Note also that the highest value is used when the values for the **Base64 Decoding Depth**, **7-Bit/8-Bit/Binary Decoding Depth**, **Quoted-Printable Decoding Depth**, or **Unix-to-Unix Decoding Depth** options are different in:

- the default network analysis policy
- any other custom network analysis policy invoked by network analysis rules in the same access control policy



Caution Changing the value for **Base64 Decoding Depth**, **7-Bit/8-Bit/Binary Decoding Depth**, **Quoted-Printable Decoding Depth**, or **Unix-to-Unix Decoding Depth** restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Ports

Specifies the ports whose SMTP traffic you want to normalize. You can specify a value greater than or equal to 0. Separate multiple ports with commas.

Stateful Inspection

When selected, causes SMTP decoder to save state and provide session context for individual packets and only inspects reassembled sessions. When cleared, analyzes each individual packet without session context.

Normalize

When set to `All`, normalizes all commands. Checks for more than one space character after a command.

When set to `None`, normalizes no commands.

When set to `Cmds`, normalizes the commands listed in **Custom Commands**.

Custom Commands

When **Normalize** is set to `Cmds`, normalizes the listed commands.

Specify commands which should be normalized in the text box. Checks for more than one space character after a command.

The space (ASCII 0x20) and tab (ASCII 0x09) characters count as space characters for normalization purposes.

Ignore Data

Does not process mail data; processes only MIME mail header data.

Ignore TLS Data

Does not process data encrypted under the Transport Layer Security protocol.

No Alerts

Disables intrusion events when accompanying preprocessor rules are enabled.

Detect Unknown Commands

Detects unknown commands in SMTP traffic.

You can enable rule 124:5 to generate events and, in an inline deployment, drop offending packets for this option.

Max Command Line Len

Detects when an SMTP command line is longer than this value. Specify `0` to never detect command line length.

RFC 2821, the Network Working Group specification on the Simple Mail Transfer Protocol, recommends 512 as a maximum command line length.

You can enable rule 124:1 to generate events and, in an inline deployment, drop offending packets for this option.

Max Header Line Len

Detects when an SMTP data header line is longer than this value. Specify `0` to never detect data header line length.

You can enable rules 124:2 and 124:7 to generate events and, in an inline deployment, drop offending packets for this option.

Max Response Line Len

Detects when an SMTP response line is longer than this value. Specify `0` to never detect response line length.

RFC 2821 recommends 512 as a maximum response line length.

You can enable rule 124:3 to generate events and, in an inline deployment, drop offending packets for this option and also for **Alt Mac Command Line Len**, when that option is enabled.

Alt Max Command Line Len

Detects when the SMTP command line for any of the specified commands is longer than this value. Specify `0` to never detect command line length for the specified commands. Different default line lengths are set for numerous commands.

This setting overrides the Max Command Line Len setting for the specified commands.

You can enable rule 124:3 to generate events and, in an inline deployment, drop offending packets for this option and also for **Max Response Line Len** when that option is enabled.

Invalid Commands

Detects if these commands are sent from the client side.

You can enable rule 124:6 to generate events and, in an inline deployment, drop offending packets for this option and also for **Invalid Commands**.

Valid Commands

Permits commands in this list.

Even if this list is empty, the preprocessor permits the following valid commands: ATRN AUTH BDAT DATA DEBUG EHLO EMAL ESAM ESND ESOM ETRN EVFY EXPN HELO HELP IDENT MAIL NOOP ONEX QUEU QUIT RCPT RSET SAML SEND SIZE SOML STARTTLS TICK TIME TURN TURNME VERB VRFY XADR XAUTH XCIR XEXCH50 X-EXPS XGEN XLICENSE X-LINK2STATE XQUE XSTA XTRN XUSR



Note RCPT TO and MAIL FROM are SMTP commands. The preprocessor configuration uses command names of RCPT and MAIL, respectively. Within the code, the preprocessor maps RCPT and MAIL to the correct command name.

You can enable rule 124:4 to generate events and, in an inline deployment, drop offending packets for this option and also for **Invalid Commands** when that option is configured.

Data Commands

Lists commands that initiate sending data in the same way the SMTP DATA command sends data per RFC 5321. Separate multiple commands with spaces.

Binary Data Commands

Lists commands that initiate sending data in a way that is similar to how the BDAT command sends data per RFC 3030. Separate multiple commands with spaces.

Authentication Commands

Lists commands that initiate an authentication exchange between client and server. Separate multiple commands with spaces.

Detect xlink2state

Detects packets that are part of X-Link2State Microsoft Exchange buffer data overflow attacks. In inline deployments, the system can also drop those packets.

You can enable rule 124:8 to generate events and, in an inline deployment, drop offending packets for this option.

Base64 Decoding Depth

When **Ignore Data** is disabled, specifies the maximum number of bytes to extract and decode from each Base64 encoded MIME email attachment. You can specify from a positive value, or specify 0 to decode all the Base64 data. Specify -1 to ignore Base64 data. The preprocessor will not decode data when **Ignore Data** is selected.

Note that positive values not divisible by 4 are rounded up to the next multiple of 4 except for the values 65533, 65534, and 65535, which are rounded down to 65532.

When this option is enabled, you can enable rule 124:10 to ggenerate events and, in an inline deployment, drop offending packets when decoding fails; decoding could fail, for example, because of incorrect encoding or corrupted data.

Note that this option replaces the deprecated options **Enable MIME Decoding** and **Maximum MIME Decoding Depth**, which are still supported in existing intrusion policies for backward compatibility.

7-Bit/8-Bit/Binary Decoding Depth

When **Ignore Data** is disabled, specifies the maximum bytes of data to extract from each MIME email attachment that does not require decoding. These attachment types include 7-bit, 8-bit, binary, and various multipart content types such as plain text, jpeg images, mp3 files, and so on. You can specify a positive value, or specify 0 to extract all data in the packet. Specify -1 to ignore non-decoded data. The preprocessor will not extract data when **Ignore Data** is selected.

Quoted-Printable Decoding Depth

When **Ignore Data** is disabled, specifies the maximum number of bytes to extract and decode from each quoted-printable (QP) encoded MIME email attachment.

You can specify from 1 to 65535 bytes, or specify 0 to decode all QP encoded data in the packet. Specify -1 to ignore QP encoded data. The preprocessor will not decode data when **Ignore Data** is selected.

When this option is enabled, you can enable rule 124:11 to generate events and, in an inline deployment, drop offending packets when decoding fails; decoding could fail, for example, because of incorrect encoding or corrupted data.

Unix-to-Unix Decoding Depth

When **Ignore Data** is disabled, specifies the maximum number of bytes to extract and decode from each Unix-to-Unix encoded (uuencoded) email attachment. You can specify from 1 to 65535 bytes, or specify 0 to decode all uuencoded data in the packet. Specify -1 to ignore uuencoded data. The preprocessor will not decode data when **Ignore Data** is selected.

When this option is enabled, you can enable rule 124:13 to generate events and, in an inline deployment, drop offending packets when decoding fails; decoding could fail, for example, because of incorrect encoding or corrupted data.

Log MIME Attachment Names

Enables extraction of MIME attachment file names from the MIME Content-Disposition header and associates the file names with all intrusion events generated for the session. Multiple file names are supported.

When this option is enabled, you can view file names associated with events in the Email Attachment column of the intrusion events table view.

Log To Addresses

Enables extraction of recipient email addresses from the SMTP RCPT TO command and associates the recipient addresses with all intrusion events generated for the session. Multiple recipients are supported.

When this option is enabled, you can view recipients associated with events in the Email Recipient column of the intrusion events table view.

Log From Addresses

Enables extraction of sender email addresses from the SMTP MAIL FROM command and associates the sender addresses with all intrusion events generated for the session. Multiple sender addresses are supported.

When this option is enabled, you can view senders associated with events in the Email Sender column of the intrusion events table view.

Log Headers

Enables extraction of email headers. The number of bytes to extract is determined by the value specified for **Header Log Depth**.

You can use the `content` or `protected_content` keyword to write intrusion rules that use email header data as a pattern. You can also view the extracted email header in the intrusion event packet view.

Header Log Depth

Specifies the number of bytes of the email header to extract when **Log Headers** is enabled. You can specify 0 to 20480 bytes. A value of 0 disables **Log Headers**.

Related Topics

Basic content and protected_content Keyword Arguments, on page 1567

Configuring SMTP Decoding

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Caution Changing the value for **Base64 Decoding Depth**, **7-Bit/8-Bit/Binary Decoding Depth**, **Quoted-Printable Decoding Depth**, or **Unix-to-Unix Decoding Depth** restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation pane.

Step 4 If **SMTP Configuration** under **Application Layer Preprocessors** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **SMTP Configuration**.

Step 6 Modify the options described in SMTP Preprocessor Options, on page 1736.

Step 7 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, cached changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want to generate events and, in an inline deployment, drop offending packets, enable SMTP preprocessor rules (GID 124). For more information, see Setting Intrusion Rule States, on page 1504.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Managing Layers, on page 1470
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

The SSH Preprocessor

The SSH preprocessor detects:

- The Challenge-Response Buffer Overflow exploit
- The CRC-32 exploit
- The SecureCRT SSH Client Buffer Overflow exploit
- Protocol mismatches
- Incorrect SSH message direction
- Any version string other than version 1 or 2

Challenge-Response Buffer Overflow and CRC-32 attacks occur after the key exchange and are, therefore, encrypted. Both attacks send an uncharacteristically large payload of more than 20 KBytes to the server immediately after the authentication challenge. CRC-32 attacks apply only to SSH Version 1; Challenge-Response Buffer Overflow exploits apply only to SSH Version 2. The version string is read at the beginning of the session. Except for the difference in the version string, both attacks are handled in the same way.

The SecureCRT SSH exploit and protocol mismatch attacks occur when attempting to secure a connection, before the key exchange. The SecureCRT exploit sends an overly long protocol identifier string to the client that causes a buffer overflow. A protocol mismatch occurs when either a non-SSH client application attempts to connect to a secure SSH server or the server and client version numbers do not match.

You can configure the SSH preprocessor to inspect traffic on a specified port or list of ports, or to automatically detect SSH traffic. It will continue to inspect SSH traffic until either a specified number of encrypted packets has passed within a specified number of bytes, or until a specified maximum number of bytes is exceeded within the specified number of packets. If the maximum number of bytes is exceeded, it is assumed that a CRC-32 (SSH Version 1) or a Challenge-Response Buffer Overflow (SSH Version 2) attack has occurred. Note that without configuration the preprocessor detects any version string value other than version 1 or 2.

Also note the that the SSH preprocessor does not handle brute force attacks.

SSH Preprocessor Options

The preprocessor stops inspecting traffic for a session when either of the following occurs:

- a valid exchange between the server and the client has occurred for this number of encrypted packets; the connection continues.
- the **Number of Bytes Sent Without Server Response** is reached before the number of encrypted packets to inspect is reached; the assumption is made that there is an attack.

Each valid server response during **Number of Encrypted Packets to Inspect** resets the **Number of Bytes Sent Without Server Response** and the packet count continues.

Consider the following example SSH preprocessor configuration:

- **Server Ports**: 22
- **Autodetect Ports**: off

- **Maximum Length of Protocol Version String**: 80
- **Number of Encrypted Packets to Inspect**: 25
- **Number of Bytes Sent Without Server Response**: 19,600
- All detect options are enabled.

In the example, the preprocessor inspects traffic only on port 22. That is, auto-detection is disabled, so it inspects only on the specified port.

Additionally, the preprocessor in the example stops inspecting traffic when either of the following occurs:

- The client sends 25 encrypted packets which contain no more than 19,600 bytes, cumulative. The assumption is there is no attack.
- The client sends more than 19,600 bytes within 25 encrypted packets. In this case, the preprocessor considers the attack to be the Challenge-Response Buffer Overflow exploit because the session in the example is an SSH Version 2 session.

The preprocessor in the example will also detect any of the following that occur while it is processing traffic:

- a server overflow, triggered by a version string greater than 80 bytes and indicating a SecureCRT exploit
- a protocol mismatch
- a packet flowing in the wrong direction

Finally, the preprocessor will automatically detect any version string other than version 1 or version 2.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Server Ports

Specifies on which ports the SSH preprocessor should inspect traffic.

You can configure a single port or a comma-separated list of ports.

Autodetect Ports

Sets the preprocessor to automatically detect SSH traffic.

When this option is selected, the preprocessor inspects all traffic for an SSH version number. It stops processing when neither the client nor the server packet contains a version number. When disabled, the preprocessor inspects only the traffic identified by the **Server Ports** option.

Number of Encrypted Packets to Inspect

Specifies the number of stream reassembled encrypted packets to examine per session.

Setting this option to zero will allow all traffic to pass.

Reducing the number of encrypted packets to inspect may result in some attacks escaping detection. Raising the number of encrypted packets to inspect may negatively affect performance.

Number of Bytes Sent Without Server Response

Specifies the maximum number of bytes an SSH client may send to a server without getting a response before assuming there is a Challenge-Response Buffer Overflow or CRC-32 attack.

Increase the value for this option if the preprocessor generates false positives on the Challenge-Response Buffer Overflow or CRC-32 exploit.

Maximum Length of Protocol Version String

Specifies the maximum number of bytes allowed in the server's version string before considering it to be a SecureCRT exploit.

Detect Challenge-Response Buffer Overflow Attack

Enables or disables detecting the Challenge-Response Buffer Overflow exploit.

You can enable rule 128:1 to generate events and, in an inline deployment, drop offending packets for this option. Note that an SFTP session can occasionally trigger rule 128:1.

Detect SSH1 CRC-32 Attack

Enables or disables detecting the CRC-32 exploit.

You can enable rule 128:2 to generate events and, in an inline deployment, drop offending packets for this option.

Detect Server Overflow

Enables or disables detecting the SecureCRT SSH Client Buffer Overflow exploit.

You can enable rule 128:3 togenerate events and, in an inline deployment, drop offending packets for this option.

Detect Protocol Mismatch

Enables or disables detecting protocol mismatches.

You can enable rule 128:4 to generate events and, in an inline deployment, drop offending packets for this option.

Detect Bad Message Direction

Enables or disables detecting when traffic flows in the wrong direction (that is, if the presumed server generates client traffic, or if a client generates server traffic).

You can enable rule 128:5 to generate events and, in an inline deployment, drop offending packets for this option.

Detect Payload Size Incorrect for the Given Payload

Enables or disables detecting packets with an incorrect payload size such as when the length specified in the SSH packet is not consistent with the total length specified in the IP header or the message is truncated, that is, there is not enough data for a full SSH header.

You can enable rule 128:6 to generate events and, in an inline deployment, drop offending packets for this option.

Detect Bad Version String

Note that, when enabled, the preprocessor detects without configuration any version string other than version 1 or 2.

You can enable rule 128:7 to generate events and, in an inline deployment, drop offending packets for this option.

Configuring the SSH Preprocessor

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation panel.

Step 4 If **SSH Configuration** under **Application Layer Preprocessors** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **SSH Configuration**.

Step 6 Modify the options described in SSH Preprocessor Options, on page 1742.

Step 7 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, cached changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want to enable intrusion events, enable SSH preprocessor rules (GID 128). For more information, see Setting Intrusion Rule States, on page 1504.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Managing Layers, on page 1470
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

The SSL Preprocessor

The SSL preprocessor allows you to configure SSL inspection, which can block encrypted traffic, decrypt it, or inspect the traffic with access control. Whether or not you configure SSL inspection, the SSL preprocessor also analyzes SSL handshake messages when detected in traffic and determines when a session becomes encrypted. Identifying encrypted traffic allows the system to stop intrusion and file inspection of encrypted payloads, which helps reduce false positives and improve performance.

The SSL preprocessor can also examine encrypted traffic to detect attempts to exploit the Heartbleed bug, and generate events when it detects such exploits.

You can suspend inspecting traffic for intrusions and malware once the session is encrypted. If you configure SSL inspection, the SSL preprocessor also identifies encrypted traffic you can block, decrypt, or inspect with access control.

Using the SSL preprocessor to decrypt encrypted traffic does not require a license. All other SSL preprocessor functionality, including halting inspection of encrypted payloads for malware and intrusions, and detecting Heartbleed bug exploits, requires a Protection license.

How SSL Preprocessing Works

The SSL preprocessor stops intrusion and file inspection of encrypted data, and inspects encrypted traffic with an SSL policy if you configure SSL inspection. This can help to eliminate false positives. The SSL preprocessor maintains state information as it inspects the SSL handshake, tracking both the state and SSL version for that session. When the preprocessor detects that a session state is encrypted, the system marks the traffic in that session as encrypted. You can configure the system to stop processing on all packets in an encrypted session when encryption is established, as well as generate an event when it detects an attempt to exploit the Heartbleed bug.

For each packet, the SSL preprocessor verifies that the traffic contains an IP header, a TCP header, and a TCP payload, and that it occurs on the ports specified for SSL preprocessing. For qualifying traffic, the following scenarios determine whether the traffic is encrypted:

- The system observes all packets in a session, **Server side data is trusted** is not enabled, and the session includes a Finished message from both the server and the client and at least one packet from each side with an Application record and without an Alert record.
- The system misses some of the traffic, **Server side data is trusted** is not enabled, and the session includes at least one packet from each side with an Application record that is not answered with an Alert record.
- The system observes all packets in a session, **Server side data is trusted** is enabled, and the session includes a Finished message from the client and at least one packet from the client with an Application record and without an Alert record.
- The system misses some of the traffic, **Server side data is trusted** is enabled, and the session includes at least one packet from the client with an Application record that is not answered with an Alert record.

If you choose to stop processing on encrypted traffic, the system ignores future packets in a session after it marks the session as encrypted.

In addition, during the SSL handshake, the preprocessor monitors heartbeat requests and responses. The preprocessor generates an event if it detects:

- a heartbeat request containing a payload length value greater than the payload itself

- a heartbeat response that is larger than the value stored in the Max Heartbeat Length field

Note You can add the `ssl_state` and `ssl_version` keywords to a rule to use SSL state or version information within the rule.

Related Topics

SSL Keywords, on page 1609
TLS/SSL Inspection Requirements

SSL Preprocessor Options

Note The system-provided network analysis policies enable the SSL preprocessor by default. Cisco recommends that you do not disable the SSL preprocessor in custom deployments if you expect encrypted traffic to cross your network.

Without SSL inspection configured, the system attempts to inspect encrypted traffic for malware and intrusions without decrypting it. When you enable the SSL preprocessor, it detects when a session becomes encrypted. After the SSL preprocessor is enabled, the rules engine can invoke the preprocessor to obtain SSL state and version information. If you enable rules using the `ssl_state` and `ssl_version` keywords in an intrusion policy, you should also enable the SSL preprocessor in that policy.

Ports

Specifies the ports, separated by commas, where the SSL preprocessor should monitor traffic for encrypted sessions. Only ports specified in this field will be checked for encrypted traffic.

Note If the SSL preprocessor detects non-SSL traffic over the ports specified for SSL monitoring, it tries to decode the traffic as SSL traffic, and then flags it as corrupt.

Stop inspecting encrypted traffic

Enables or disables inspection of traffic in a session after the session is marked as encrypted.

Enable this option to disable inspection and reassembly for encrypted sessions. The SSL preprocessor maintains state for the session so it can disable inspection of all traffic in the session. When this option is enabled a few packets of a session are verified to ensure the flow is encrypted after which deep inspection is bypassed. Every bypassed session increases the fast-forwarded flows count shown in the response of the **show snort statistics** command. Moreover, since deep inspection is bypassed, the initiator and responder bytes in the connection event are not accurate. They are less than the value of the actual session, since it only includes the packets inspected by Snort and it does not include any packets after the deep inspection is bypassed. This behavior holds good for connection summary events and all traffic values shown in the widgets.

The system only stops inspecting traffic in encrypted sessions if both:

- SSL preprocessing is enabled

- this option is selected

If you clear this option, you cannot modify the **Server side data is trusted** option.

Server side data is trusted

When Stop inspecting encrypted traffic is enabled, enables identification of encrypted traffic based only on the client-side traffic,

Max Heartbeat Length

By specifying a number of bytes, enables inspection of heartbeat requests and responses within the SSL handshake for Heartbleed bug exploit attempts. You can specify an integer from `1` to `65535`, or `0` to disable the option.

If the preprocessor detects a heartbeat request whose payload length is greater than the actual payload length and rule 137:3 is enabled, or a heartbeat response greater in size than the value configured for this option when rule 137:4 is enabled, the preprocessor generates an event and, in an inline deployment, drops offending packets.

Configuring the SSL Preprocessor

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation panel.

Step 4 If **SSL Configuration** under **Application Layer Preprocessors** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **SSL Configuration**.

Step 6 Modify any of the settings described in SSL Preprocessor Options, on page 1747.

- Enter a value in the **Ports** field. Separate multiple values with commas.
- Check or clear the **Stop inspecting encrypted traffic** check box.
- If you checked **Stop inspecting encrypted traffic**, check or clear **Server side data is trusted**.
- Enter a value in the **Max Heartbeat Length** field.

Tip A value of `0` disables this option.

Step 7 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, cached changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want to enable intrusion events, enable SSL preprocessor rules (GID 137). For more information, see Setting Intrusion Rule States, on page 1504.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Managing Layers, on page 1470
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461
TLS/SSL Inspection Requirements

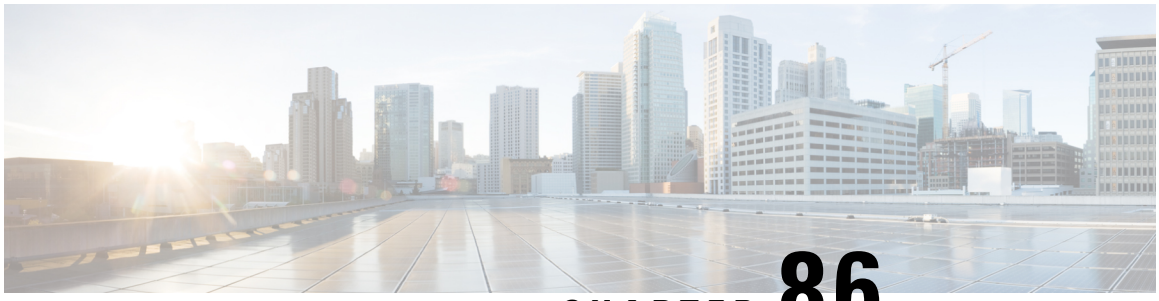
SSL Preprocessor Rules

If you want to generate events and, in an inline deployment, drop offending packets, enable SSL preprocessor rules (GID 137).

The following table describes the SSL preprocessor rules you can enable.

Table 205: SSL Preprocessor Rules

Preprocessor Rule GID:SID	Description
137:1	Detects a ClientHello message after a ServerHello message, which is invalid and considered to be anomalous behavior.
137:2	Detects a ServerHello message without a ClientHello message when the SSL preprocessor option **Server side data is trusted** is disabled, which is invalid and considered to be anomalous behavior.
137:3	Detects a heartbeat request with a payload length greater than the payload itself when the SSL preprocessor option **Max Heartbeat Length** contains a non-zero value, which indicates an attempt to exploit the Heartbleed bug.
137:4	Detects a heartbeat response larger than a non-zero value specified in the SSL preprocessor option **Max Heartbeat Length**, which indicates an attempt to exploit the Heartbleed bug.



CHAPTER 86

SCADA Preprocessors

The following topics explain preprocessors for Supervisory Control and Data Acquisition (SCADA) protocols, and how to configure them:

- Introduction to SCADA Preprocessors, on page 1751
- License Requirements for SCADA Preprocessors, on page 1751
- Requirements and Prerequisites for SCADA Preprocessors, on page 1752
- The Modbus Preprocessor, on page 1752
- The DNP3 Preprocessor, on page 1754

Introduction to SCADA Preprocessors

Supervisory Control and Data Acquisition (SCADA) protocols monitor, control, and acquire data from industrial, infrastructure, and facility processes such as manufacturing, production, water treatment, electric power distribution, airport and shipping systems, and so on. The Firepower System provides preprocessors for the Modbus and Distributed Network Protocol (DNP3) SCADA protocols that you can configure as part of your network analysis policy.

If the Modbus or DNP3 preprocessor is disabled, and you enable and deploy an intrusion rule that requires one of these preprocessors, the system automatically uses the required preprocessor, with its current settings, although the preprocessor remains disabled in the web interface for the corresponding network analysis policy.

License Requirements for SCADA Preprocessors

FTD License

Threat

Classic License

Protection

Requirements and Prerequisites for SCADA Preprocessors

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Intrusion Admin

The Modbus Preprocessor

The Modbus protocol, which was first published in 1979 by Modicon, is a widely used SCADA protocol. The Modbus preprocessor detects anomalies in Modbus traffic and decodes the Modbus protocol for processing by the rules engine, which uses Modbus keywords to access certain protocol fields.

A single configuration option allows you to modify the default setting for the port that the preprocessor inspects for Modbus traffic.

Related Topics

SCADA Keywords, on page 1631

Modbus Preprocessor Ports Option

Ports

Specifies the ports that the preprocessor inspects for Modbus traffic. Separate multiple ports with commas.

Configuring the Modbus Preprocessor

You should not enable this preprocessor in a network analysis policy that you apply to traffic if your network does not contain any Modbus-enabled devices.

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation panel.

Step 4 If **Modbus Configuration** under **SCADA Preprocessors** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **Modbus Configuration**.

Step 6 Enter a value in the **Ports** field.

Separate multiple values with commas.

Step 7 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, cached changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want to generate events and, in an inline deployment, drop offending packets, enable Modbus preprocessor rules (GID 144). For more information, see Setting Intrusion Rule States, on page 1504 and Modbus Preprocessor Rules, on page 1753.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Managing Layers, on page 1470
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

Modbus Preprocessor Rules

You must enable the Modbus preprocessor rules in the following table if you want these rules to generate events and, in an inline deployment, drop offending packets.

Table 206: Modbus Preprocessor Rules

Preprocessor Rule GID:SID	Description
144:1	Generates an event when the length in the Modbus header does not match the length required by the Modbus function code. Each Modbus function has an expected format for requests and responses. If the length of the message does not match the expected format, this event is generated.

Preprocessor Rule GID:SID	Description
144:2	Generates an event when the Modbus protocol ID is non-zero. The protocol ID field is used for multiplexing other protocols with Modbus. Because the preprocessor does not process these other protocols, this event is generated instead.
144:3	Generates an event when the preprocessor detects a reserved Modbus function code.

The DNP3 Preprocessor

The Distributed Network Protocol (DNP3) is a SCADA protocol that was originally developed to provide consistent communication between electrical stations. DNP3 has also become widely used in the water, waste, transportation, and many other industries.

The DNP3 preprocessor detects anomalies in DNP3 traffic and decodes the DNP3 protocol for processing by the rules engine, which uses DNP3 keywords to access certain protocol fields.

Related Topics

DNP3 Keywords, on page 1632

DNP3 Preprocessor Options

Ports

Enables inspection of DNP3 traffic on each specified port. You can specify a single port or a comma-separated list of ports.

Log bad CRCs

Validates the checksums contained in DNP3 link layer frames. Frames with invalid checksums are ignored.

You can enable rule 145:1 to generate events and, in an inline deployment, drop offending packets when invalid checksums are detected.

Configuring the DNP3 Preprocessor

You should not enable this preprocessor in a network analysis policy that you apply to traffic if your network does not contain any DNP3-enabled devices.

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation panel.

Step 4 If **DNP3 Configuration** under **SCADA Preprocessors** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **DNP3 Configuration**.

Step 6 Enter a value for **Ports**.

Separate multiple values with commas.

Step 7 Check or clear the **Log bad CRCs** check box.

Step 8 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, cached changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want to generate events and, in an inline deployment, drop offending packets, enable DNP3 preprocessor rules (GID 145). For more information, see Setting Intrusion Rule States, on page 1504, DNP3 Preprocessor Options, on page 1754, and DNP3 Preprocessor Rules, on page 1755.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Managing Layers, on page 1470
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

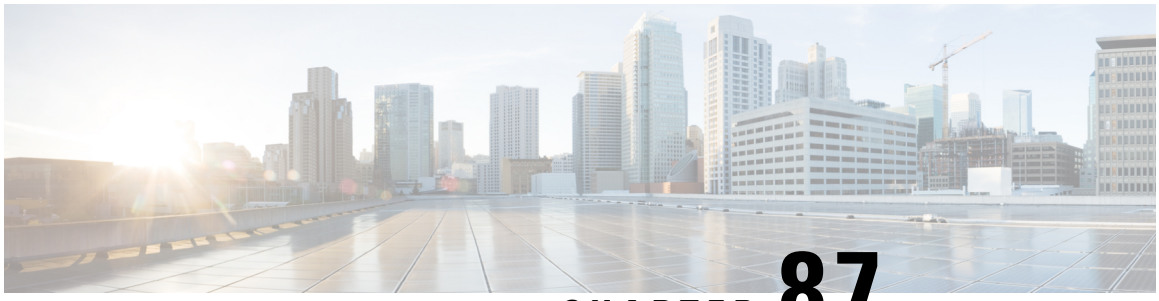
DNP3 Preprocessor Rules

You must enable the DNP3 preprocessor rules in the following table if you want these rules to generate events and, in an inline deployment, drop offending packets.

Table 207: DNP3 Preprocessor Rules

Preprocessor Rule GID:SID	Description
145:1	When **Log bad CRC** is enabled, generates an event when the preprocessor detects a link layer frame with an invalid checksum.
145:2	Generates an event and blocks the packet when the preprocessor detects a DNP3 link layer frame with an invalid length.

Preprocessor Rule GID:SID	Description
145:3	Generates an event and blocks the packet during reassembly when the preprocessor detects a transport layer segment with an invalid sequence number.
145:4	Generates an event when the DNP3 reassembly buffer is cleared before a complete fragment can be reassembled. This happens when a segment carrying the FIR flag appears after other segments have been queued.
145:5	Generates an event when the preprocessor detects a DNP3 link layer frame that uses a reserved address.
145:6	Generates an event when the preprocessor detects a DNP3 request or response that uses a reserved function code.



CHAPTER 87

Transport & Network Layer Preprocessors

The following topics explain transport and network layer preprocessors and how to configure them:

- Introduction to Transport and Network Layer Preprocessors, on page 1757
- License Requirements for Transport and Network Layer Preprocessors, on page 1757
- Requirements and Prerequisites for Transport and Network Layer Preprocessors, on page 1758
- Advanced Transport/Network Preprocessor Settings, on page 1758
- Checksum Verification, on page 1761
- The Inline Normalization Preprocessor, on page 1762
- The IP Defragmentation Preprocessor, on page 1769
- The Packet Decoder, on page 1774
- TCP Stream Preprocessing, on page 1778
- UDP Stream Preprocessing, on page 1788

Introduction to Transport and Network Layer Preprocessors

Transport and network layer preprocessors detect attacks that exploit IP fragmentation, checksum validation, and TCP and UDP session preprocessing. Before packets are sent to preprocessors, the packet decoder converts packet headers and payloads into a format that can be easily used by the preprocessors and the intrusion rules engine and detects various anomalous behaviors in packet headers. After packet decoding and before sending packets to other preprocessors, the inline normalization preprocessor normalizes traffic for inline deployments.

When an intrusion rule or rule argument requires a disabled preprocessor, the system automatically uses it with its current configuration even though it remains disabled in the network analysis policy's web interface.

License Requirements for Transport and Network Layer Preprocessors

FTD License

Threat

Classic License

Protection

Requirements and Prerequisites for Transport and Network Layer Preprocessors

Model Support

Any.

Supported Domains

Any

User Roles

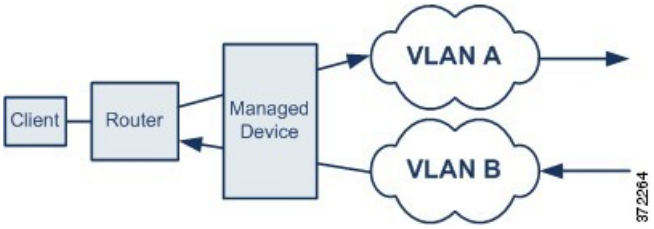
- Admin
- Intrusion Admin

Advanced Transport/Network Preprocessor Settings

Advanced transport and network preprocessor settings apply globally to all networks, zones, and VLANs where you deploy your access control policy. You configure these advanced settings in an access control policy rather than in a network analysis policy.

Ignored VLAN Headers

Different VLAN tags in traffic traveling in different directions for the same connection can affect traffic reassembly and rule processing. For example, in the following graphic traffic for the same connection could be transmitted over VLAN A and received over VLAN B.



You can configure the system to ignore the VLAN header so packets can be correctly processed for your deployment.

Note This option is not supported on ASA FirePOWER.

Active Responses in Intrusion Drop Rules

A drop rule is an intrusion or preprocessor rule whose rule state is set to Drop and Generate Events. In an inline deployment, the system responds to TCP or UDP drop rules by dropping the triggering packet and blocking the session where the packet originated.

Tip Because UDP data streams are not typically thought of in terms of *sessions*, the stream preprocessor uses the source and destination IP address fields in the encapsulating IP datagram header and the port fields in the UDP header to determine the direction of flow and identify a UDP session.

You can configure the system to initiate one or more *active responses* to more precisely and specifically close a TCP connection or UDP session when an offending packet triggers a TCP or UDP drop rule. You can use active responses in inline, including routed and transparent, deployments. Active responses are not suited or supported for passive deployments.

To configure active responses:

- Create or modify a TCP or UDP (**resp** keyword only) intrusion rule. See Intrusion Rule Header Protocol, on page 1546.
- Add the **react** or **resp** keyword to the intrusion rule; see xActive Response Keywords, on page 1636.
- Optionally, for a TCP connection, specify the maximum number of additional active responses to send and the number of seconds to wait between active respones; see **Maximum Active Responses** and **Minimum Response Seconds** in Advanced Transport/Network Preprocessor Options, on page 1759.

Active responses close the sesson when matching traffic triggers a drop rule, as follows:

- **TCP**—drops the triggering packet and inserts a TCP Reset (RST) packet in both the client and server traffic.
- **UDP**—sends an ICMP unreachable packet to each end of the session.

Advanced Transport/Network Preprocessor Options

Ignore the VLAN header when tracking connections

Specifies whether to ignore or include VLAN headers when identifying traffic, as follows:

- When this option is selected, the system ignores VLAN headers. Use this setting for deployed devices that might detect different VLAN tags for the same connection in traffic traveling in different directions
- When this option is disabled, the system includes VLAN headers. Use this setting for deployed devices that will not detect different VLAN tags for the same connection traffic traveling in different directions.

Note This option is not supported on ASA FirePOWER.

Maximum Active Responses

Specifies a maximum number of active responses per TCP connection. When additional traffic occurs on a connection where an active response has been initiated, and the traffic occurs more than **Minimum Response Seconds** after a previous active response, the system sends another active response unless the specified maximum has been reached. A setting of 0 disables additional active responses triggered by **resp** or **react** rules. See Active Responses in Intrusion Drop Rules, on page 1759 and Active Response Keywords, on page 1636.

Note that a triggered **resp** or **react** rule initiates an active response regardless of the configuration of this option.

Minimum Response Seconds

Until **Maximum Active Responses** occur, specifies the number of seconds to wait before any additional traffic on a connection where the system has initiated an active response results in a subsequent active response.

Troubleshooting Options: Session Termination Logging Threshold

Caution Do not modify Session Termination Logging Threshold unless instructed to do so by Support.

Support might ask you during a troubleshooting call to configure your system to log a message when an individual connection exceeds the specified threshold. Changing the setting for this option will affect performance and should be done only with Support guidance.

This option specifies for the number of bytes that result in a logged message when the session terminates and the specified number was exceeded.

Note The upper limit of 1GB is also restricted by the amount of memory on the managed device allocated for stream processing.

Related Topics

Active Response Keywords, on page 1636

Configuring Advanced Transport/Network Preprocessor Settings

You must be an Admin, Access Admin, or Network Admin to perform this task.

Procedure

Step 1 In the access control policy editor, click **Advanced**.

Step 2 Click **Edit** ( ) next to the Transport/Network Layer Settings section.

Step 3 Except for the troubleshooting option **Session Termination Logging Threshold**, modify the options described in Advanced Transport/Network Preprocessor Options, on page 1759.

Note The **Ignore the VLAN header when tracking connectons** option is not available on the ASA FirePOWER module.

Caution Do not modify **Session Termination Logging Threshold** unless instructed to do so by Support.

Step 4 Click **OK**.

What to do next

- Optionally, further configure the policy as described in Editing an Access Control Policy, on page 1205.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Checksum Verification

The system can verify all protocol-level checksums to ensure that complete IP, TCP, UDP, and ICMP transmissions are received and that, at a basic level, packets have not been tampered with or accidentally altered in transit. A checksum uses an algorithm to verify the integrity of a protocol in the packet. The packet is considered to be unchanged if the system computes the same value that is written in the packet by the end host.

Disabling checksum verification may leave your network susceptible to insertion attacks. Note that the system does not generate checksum verification events. In an inline deployment, you can configure the system to drop packets with invalid checksums.

Checksum Verification Options

You can set any of the following options to **Enabled** or **Disabled** in a passive or inline deployment, or to **Drop** in an inline deployment:

- **ICMP Checksums**
- **IP Checksums**
- **TCP Checksums**
- **UDP Checksums**

To drop offending packets, in addition to setting an option to **Drop** you must also enable **Inline Mode** in the associated network analysis policy and ensure that the device is deployed inline.

Setting these options to **Drop** in a passive deployment, or in an inline deployment in tap mode, is the same as setting them to **Enabled**.

The default for all checksum verification options is **Enabled**. However, Firepower Threat Defense routed and transparent interfaces always drop packets that fail IP checksum verification. Note that the Firepower Threat Defense routed and transparent interfaces fix UDP packets with a bad checksum before passing the packets to the Snort process.

Related Topics

Preprocessor Traffic Modification in Inline Deployments, on page 1683

Verifying Checksums

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation panel.

Step 4 If **Checksum Verification** under **Transport/Network Layer Preprocessors** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **Checksum Verification**.

Step 6 Modify the options described in Checksum Verification, on page 1761.

Step 7 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, cached changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Layer Management, on page 1468
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

The Inline Normalization Preprocessor

The inline normalization preprocessor normalizes traffic to minimize the chances of attackers evading detection in inline deployments.

Note For the system to affect traffic, you must deploy relevant configurations to managed devices using routed, switched, or transparent interfaces, or inline interface pairs.

You can specify normalization of any combination of IPv4, IPv6, ICMPv4, ICMPv6, and TCP traffic. Most normalizations are on a per-packet basis and are conducted by the inline normalization preprocessor. However,

the TCP stream preprocessor handles most state-related packet and stream normalizations, including TCP payload normalization.

Inline normalization takes place immediately after decoding by the packet decoder and before processing by other preprocessors. Normalization proceeds from the inner to outer packet layers.

The inline normalization preprocessor does not generate events; it prepares packets for use by other preprocessors and the rules engine in inline deployments. The preprocessor also helps ensure that the packets the system processes are the same as the packets received by the hosts on your network.

Note In an inline deployment, Cisco recommends that you enable inline mode and configure the inline normalization preprocessor with the **Normalize TCP Payload** option enabled. In a passive deployment, Cisco recommends that you use adaptive profile updates.

Related Topics

Preprocessor Traffic Modification in Inline Deployments, on page 1683
About Adaptive Profiles, on page 1811

Inline Normalization Options

Minimum TTL

When **Reset TTL** is greater than or equal to the value set for this option, specifies the following:

- the minimum value the system will permit in the IPv4 Time to Live (TTL) field when **Normalize IPv4** is enabled; a lower value results in normalizing the packet value for TTL to the value set for **Reset TTL**
- the minimum value the system will permit in the IPv6 Hop Limit field when **Normalize IPv6** is enabled; a lower value results in normalizing the packet value for Hop Limit to the value set for **Reset TTL**

The system assumes a value of 1 when the field is empty.

Note For Firepower Threat Defense routed and transparent interfaces, the **Minimum TTL** and **Reset TTL** options are ignored. The maximum TTL for a connection is determined by the TTL in the initial packet. The TTL for subsequent packets can decrease, but it cannot increase. The system will reset the TTL to the lowest previously-seen TTL for that connection. This prevents TTL evasion attacks.

When the packet decoding **Detect Protocol Header Anomalies** option is enabled, you can enable the following rules in the decoder rule category to generate events and, in an inline deployment, drop offending packets for this option:

- You can enable rule 116:428 to trigger when the system detects an IPv4 packet with a TTL less than the specified minimum.
- You can enable rule 116:270 to trigger when the system detects an IPv6 packet with a hop limit that is less than the specified minimum.

Reset TTL

When set to a value greater than or equal to **Minimum TTL**, normalizes the following:

- the IPv4 TTL field when **Normalize IPv4** is enabled
- the IPv6 Hop Limit field when **Normalize IPv6** is enabled

The system normalizes the packet by changing its TTL or Hop Limit value to the value set for this option when the packet value is less than **Minimum TTL**. Leaving this field blank, or setting it to `0`, or to any value less than **Minimum TTL**, disables the option.

Normalize IPv4

Enables normalization of IPv4 traffic. The system also normalizes the TTL field as needed when:

- this option is enabled, and
- the value set for **Reset TTL** enables TTL normalization.

You can also enable additional IPv4 options when this option is enabled.

When you enable this option, the system performs the following base IPv4 normalizations:

- truncates packets with excess payload to the datagram length specified in the IP header
- clears the Differentiated Services (DS) field, formerly known as the Type of Service (TOS) field
- sets all option octets to 1 (No Operation)

This option is ignored for Firepower Threat Defense routed and transparent interfaces. Firepower Threat Defense devices will drop any RSVP packet that contains IP options other than the router alert, end of options list (EOOL), and no operation (NOP) options on any routed or transparent interface.

Normalize Don't Fragment Bit

Clears the single-bit Don’t Fragment subfield of the IPv4 Flags header field. Enabling this option allows a downstream router to fragment packets if necessary instead of dropping them; enabling this option can also prevent evasions based on crafting packets to be dropped. You must enable **Normalize IPv4** to select this option.

Normalize Reserved Bit

Clears the single-bit Reserved subfield of the IPv4 Flags header field. You would typically enable this option. You must enable **Normalize IPv4** to select this option.

Normalize TOS Bit

Clears the one byte Differentiated Services field, formerly known as Type of Service. You must enable **Normalize IPv4** to select this option.

Normalize Excess Payload

Truncates packets with excess payload to the datagram length specified in the IP header plus the Layer 2 (for example, Ethernet) header, but does not truncate below the minimum frame length. You must enable **Normalize IPv4** to select this option.

This option is ignored for Firepower Threat Defense routed and transparent interfaces. Packets with excess payload are always dropped on these interfaces.

Normalize IPv6

Sets all Option Type fields in the Hop-by-Hop Options and Destination Options extension headers to 00 (Skip and continue processing). The system also normalizes the Hop Limit field as needed when this option is enabled and the value set for **Reset TTL** enables hop limit normalization.

Normalize ICMPv4

Clears the 8-bit Code field in Echo (Request) and Echo Reply messages in ICMPv4 traffic.

Normalize ICMPv6

Clears the 8-bit Code field in Echo (Request) and Echo Reply messages in ICMPv6 traffic.

Normalize/Clear Reserved Bits

Clears the Reserved bits in the TCP header.

Normalize/Clear Option Padding Bytes

Clears any TCP option padding bytes.

Clear Urgent Pointer if URG=0

Clears the 16-bit TCP header Urgent Pointer field if the urgent (URG) control bit is not set.

Clear Urgent Pointer/URG on Empty Payload

Clears the TCP header Urgent Pointer field and the URG control bit if there is no payload.

Clear URG if Urgent Pointer is Not Set

Clears the TCP header URG control bit if the urgent pointer is not set.

Normalize Urgent Pointer

Sets the two-byte TCP header Urgent Pointer field to the payload length if the pointer is greater than the payload length.

Normalize TCP Payload

Enables normalization of the TCP Data field to ensure consistency in retransmitted data. Any segment that cannot be properly reassembled is dropped.

Remove Data on SYN

Removes data in synchronization (SYN) packets if your TCP operating system policy is **not** Mac OS.

This option also disables rule 129:2, which can otherwise trigger when the TCP stream preprocessor **Policy** option is not set to **Mac OS**.

Remove Data on RST

Removes any data from a TCP reset (RST) packet.

Trim Data to Window

Trims the TCP Data field to the size specified in the Window field.

Trim Data to MSS

Trims the TCP Data field to the Maximum Segment Size (MSS) if the payload is longer than MSS.

Block Unresolvable TCP Header Anomalies

When you enable this option, the system blocks anomalous TCP packets that, if normalized, would be invalid and likely would be blocked by the receiving host. For example, the system blocks any SYN packet transmitted subsequent to an established session.

The system also drops any packet that matches any of the following TCP stream preprocessor rules, regardless of whether the rules are enabled:

- 129:1
- 129:3
- 129:4
- 129:6
- 129:8
- 129:11
- 129:14 through 129:19

The Total Blocked Packets performance graph tracks the number of packets blocked in inline deployments and, in passive deployments and inline deployments in tap mode, the number that would have been blocked in an inline deployment.

Explicit Congestion Notification

Enables per-packet or per-stream normalization of Explicit Congestion Notification (ECN) flags as follows:

- select **Packet** to clear ECN flags on a per-packet basis regardless of negotiation
- select **Stream** to clear ECN flags on a per-stream basis if ECN use was not negotiated

If you select **Stream**, you must also ensure that the TCP stream preprocessor **Require TCP 3-Way Handshake** option is enabled for this normalization to take place.

Clear Existing TCP Options

Enables **Allow These TCP Options**.

Allow These TCP Options

Disables normalization of specific TCP options you allow in traffic.

The system does not normalize options that you explicitly allow. It normalizes options that you do not explicitly allow by setting the options to No Operation (TCP Option 1).

The system always allows the following options regardless of the configuration of **Allow These TCP Options** because they are commonly used for optimal TCP performance:

- Maximum Segment Size (MSS)
- Window Scale
- Time Stamp TCP

The system does not automatically allow other less commonly used options.

You can allow specific options by configuring a comma-separated list of option keywords, option numbers, or both as shown in the following example:

```
sack, echo, 19
```

Specifying an option keyword is the same as specifying the number for one or more TCP options associated with the keyword. For example, specifying `sack` is the same as specifying TCP options 4 (Selective Acknowledgment Permitted) and 5 (Selective Acknowledgment). Option keywords are not case sensitive.

You can also specify `any`, which allows all TCP options and effectively disables normalization of all TCP options.

The following table summarizes how you can specify TCP options to allow. If you leave the field empty, the system allows only the MSS, Window Scale, and Time Stamp options.

Specify...	To allow...
sack	TCP options 4 (Selective Acknowledgment Permitted) and 5 (Selective Acknowledgment)
echo	TCP options 6 (Echo Request) and 7 (Echo Reply)
partial_order	TCP options 9 (Partial Order Connection Permitted) and 10 (Partial Order Service Profile)
conn_count	TCP Connection Count options 11 (CC), 12 (CC.New), and 13 (CC.Echo)
alt_checksum	TCP options 14 (Alternate Checksum Request) and 15 (Alternate Checksum)
md5	TCP option 19 (MD5 Signature)
the option number, 2 to 255	a specific option, including options for which there is no keyword
any	all TCP options; this setting effectively disables TCP option normalization

When you do not specify `any` for this option, normalizations include the following:

- except MSS, Window Scale, Time Stamp, and any explicitly allowed options, sets all option bytes to No Operation (TCP Option 1)

- sets the Time Stamp octets to No Operation if Time Stamp is present but invalid, or valid but not negotiated
- blocks the packet if Time Stamp is negotiated but not present
- clears the Time Stamp Echo Reply (TSecr) option field if the Acknowledgment (ACK) control bit is not set
- sets the MSS and Window Scale options to No Operation (TCP Option 1) if the SYN control bit is not set

Related Topics

Intrusion Event Performance Statistics Graph Types, on page 2290

Configuring Inline Normalization

Before you begin

- If you want to normalize or drop offending packets, enable **Inline Mode** as described in Preprocessor Traffic Modification in Inline Deployments, on page 1683. The managed device must also be deployed inline.

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation panel (NOT the caret; click the word).

Step 4 If **Inline Normalization** under **Transport/Network Layer Preprocessors** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **Inline Normalization**.

Step 6 Set the options described in The Inline Normalization Preprocessor, on page 1762.

Step 7 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, cached changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want the inline normalizaition Minimum TTL option to generate intrusion events, enable either or both packet decoder rules 116:429 (IPv4) and 116:270 (IPv6). For more information, see Setting Intrusion Rule States, on page 1504, and Inline Normalization Options, on page 1763.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Layer Management, on page 1468
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

The IP Defragmentation Preprocessor

When an IP datagram is broken into two or more smaller IP datagrams because it is larger than the maximum transmission unit (MTU), it is *fragmented*. A single IP datagram fragment may not contain enough information to identify a hidden attack. Attackers may attempt to evade detection by transmitting attack data in fragmented packets. The IP defragmentation preprocessor reassembles fragmented IP datagrams before the rules engine executes rules against them so the rules can more appropriately identify attacks in those packets. If fragmented datagrams cannot be reassembled, rules do not execute against them.

IP Fragmentation Exploits

Enabling IP defragmentation helps you detect attacks against hosts on your network, like the teardrop attack, and resource consumption attacks against the system itself, like the Jolt2 attack.

The Teardrop attack exploits a bug in certain operating systems that causes them to crash when trying to reassemble overlapping IP fragments. When enabled and configured to do so, the IP defragmentation preprocessor identifies the overlapping fragments. The IP defragmentation preprocessor detects the first packets in an overlapping fragment attack such as Teardrop, but does not detect subsequent packets for the same attack.

The Jolt2 attack sends a large number of copies of the same fragmented IP packet in an attempt to overuse IP defragmentors and cause a denial of service attack. A memory usage cap disrupts this and similar attacks in the IP defragmentation preprocessor, and places the system self-preservation above exhaustive inspection. The system is not overwhelmed by the attack, remains operational, and continues to inspect network traffic.

Different operating systems reassemble fragmented packets in different ways. Attackers who can determine which operating systems your hosts are running can also fragment malicious packets so that a target host reassembles them in a specific manner. Because the system does not know which operating systems the hosts on your monitored network are running, the preprocessor may reassemble and inspect the packets incorrectly, thus allowing an exploit to pass through undetected. To mitigate this kind of attack, you can configure the defragmentation preprocessor to use the appropriate method of defragmenting packets for each host on your network.

Note that you can also use adaptive profile updates in a passive deployment to dynamically select target-based policies for the IP defragmentation preprocessor using host operating system information for the target host in a packet.

Target-Based Defragmentation Policies

A host's operating system uses three criteria to determine which packet fragments to favor when reassembling the packet:

- the order in which the fragment was received by the operating system
- its offset (the fragment's distance, in bytes, from the beginning of the packet)
- its beginning and ending position compared to overlap fragments.

Although every operating system uses these criteria, different operating systems favor different fragments when reassembling fragmented packets. Therefore, two hosts with different operating systems on your network could reassemble the same overlapping fragments in entirely different ways.

An attacker, aware of the operating system of one of your hosts, could attempt to evade detection and exploit that host by sending malicious content hidden in overlapping packet fragments. This packet, when reassembled and inspected, seems innocuous, but when reassembled by the target host, contains a malicious exploit. However, if you configure the IP defragmentation preprocessor to be aware of the operating systems running on your monitored network segment, it will reassemble the fragments the same way that the target host does, allowing it to identify the attack.

IP Defragmentation Options

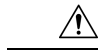
You can choose to simply enable or disable IP defragmentation; however, Cisco recommends that you specify the behavior of the enabled IP defragmentation preprocessor at a more granular level.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

You can configure the following global option:

Preallocated Fragments

The maximum number of individual fragments that the preprocessor can process at once. Specifying the number of fragment nodes to preallocate enables static memory allocation.



Caution Processing an individual fragment uses approximately 1550 bytes of memory. If the preprocessor requires more memory to process the individual fragments than the predetermined allowable memory limit for the managed device, the memory limit for the device takes precedence.

You can configure the following options for each IP defragmentation policy:

Networks

The IP address of the host or hosts to which you want to apply the defragmentation policy.

You can specify a single IP address or address block, or a comma-separated list of either or both. You can specify up to 255 total profiles, including the default policy.

Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Note that the `default` setting in the default policy specifies all IP addresses on your monitored network segment that are not covered by another target-based policy. Therefore, you cannot and do not need to specify an IP address or CIDR block/prefix length for the default policy, and you cannot leave this setting blank in another policy or use address notation to represent `any` (for example, 0.0.0.0/0 or ::/0).

Policy

The defragmentation policy you want to use for a set of hosts on your monitored network segment.

You can select one of seven defragmentation policies, depending on the operating system of the target host. The following table lists the seven policies and the operating systems that use each one. The First and Last policy names reflect whether those policies favor original or subsequent overlapping packets.

This option is ignored for Firepower Threat Defense routed and transparent interfaces.

Table 208: Target-Based Defragmentation Policies

Policy	Operating Systems
BSD	AIX FreeBSD IRIX VAX/VMS
BSD-right	HP JetDirect
First	Mac OS HP-UX
Linux	Linux OpenBSD
Last	Cisco IOS
Solaris	SunOS
Windows	Windows

Timeout

Specifies the maximum amount of time, in seconds, that the preprocessor engine can use when reassembling a fragmented packet. If the packet cannot be reassembled within the specified time period, the preprocessor engine stops attempting to reassemble the packet and discards received fragments.

Min TTL

Specifies the minimum acceptable TTL value a packet may have. This option detects TTL-based insertion attacks.

You can enable rule 123:11 to generate events and, in an inline deployment, drop offending packets for this option.

Detect Anomalies

Identifies fragmentation problems such as overlapping fragments.

This option is ignored for Firepower Threat Defense routed and transparent interfaces.

You can enable the following rules to generate events and, in an inline deployment, drop offending packets for this option:

- 123:1 through 123:4
- 123:5 (BSD policy)
- 123:6 through 123:8

Overlap Limit

Specifies that when the configured number of overlapping segments in a session has been detected, defragmentation stops for that session.

You must enable **Detect Anomalies** to configure this option. A blank value disables this option. A value of 0 specifies an unlimited number overlapping segments.

This option is ignored for Firepower Threat Defense routed and transparent interfaces. Overlapping fragments are always dropped on those interfaces.

You can enable rule 123:12 to generate events and, in an inline deployment, drop offending packets for this option.

Minimum Fragment Size

Specifies that when a non-last fragment smaller than the configured number of bytes has been detected, the packet is considered malicious.

You must enable **Detect Anomalies** to configure this option. A blank value disables this option. A value of 0 specifies an unlimited number of bytes.

You can enable rule 123:13 to generate events and, in an inline deployment, drop offending packets for this option.

Related Topics

Firepower System IP Address Conventions, on page 15

Configuring IP Defragmentation

The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Before you begin

- Confirm that any networks you want to identify in a custom target-based policy match or are a subset of the networks, zones, and VLANs handled by its parent network analysis policy. See Advanced Settings for Network Analysis Policies, on page 1672 for more information.

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation panel.

Step 4 If **IP Defragmentation** under **Transport/Network Layer Preprocessors** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **IP Defragmentation**.

Step 6 Optionally, enter a value in the **Preallocated Fragments** field.

Step 7 You have the following choices:

- Add a server profile — Click **Add** ( ) next to **Servers** on the left side of the page, then enter a value in the **Host Address** field and click **OK**. You can specify a single IP address or address block, or a comma-separated list of either or both. You can create a total of 255 target-based policies including the default policy.
- Edit a server profile — Click the configured address for under **Servers** on the left side of the page, or click **default**.
- Delete a profile — Click **Delete** ( ) next to the policy.

Step 8 Modify the options described in IP Defragmentation Options, on page 1770.

Step 9 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, cached changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want to generate events and, in an inline deployment, drop offending packets, enable IP defragmentation rules (GID 123). For more information, see Setting Intrusion Rule States, on page 1504 and IP Defragmentation Options, on page 1770.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Firepower System IP Address Conventions, on page 15
Layer Basics, on page 1463
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

The Packet Decoder

Before sending captured packets to a preprocessor, the system first sends the packets to the packet decoder. The packet decoder converts packet headers and payloads into a format that preprocessors and the rules engine can easily use. Each stack layer is decoded in turn, beginning with the data link layer and continuing through the network and transport layers.

Packet Decoder Options

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Decode GTP Data Channel

Decodes the encapsulated GTP (General Packet Radio Service [GPRS] Tunneling Protocol) data channel. By default, the decoder decodes version 0 data on port 3386 and version 1 data on port 2152. You can use the `GTP_PORTS` default variable to modify the ports that identify encapsulated GTP traffic.

You can enable rules 116:297 and 116:298 to generate events and, in an inline deployment, drop offending packets for this option.

Detect Teredo on Non-Standard Ports

Inspects Teredo tunneling of IPv6 traffic that is identified on a UDP port other than port 3544.

The system always inspects IPv6 traffic when it is present. By default, IPv6 inspection includes the 4in6, 6in4, 6to4, and 6in6 tunneling schemes, and also includes Teredo tunneling when the UDP header specifies port 3544.

In an IPv4 network, IPv4 hosts can use the Teredo protocol to tunnel IPv6 traffic through an IPv4 Network Address Translation (NAT) device. Teredo encapsulates IPv6 packets within IPv4 UDP datagrams to permit IPv6 connectivity behind an IPv4 NAT device. The system normally uses UDP port 3544 to identify Teredo traffic. However, an attacker could use a non-standard port in an attempt to avoid detection. You can enable **Detect Teredo on Non-Standard Ports** to cause the system to inspect all UDP payloads for Teredo tunneling.

Teredo decoding occurs only on the first UDP header, and only when IPv4 is used for the outer network layer. When a second UDP layer is present after the Teredo IPv6 layer because of UDP data encapsulated in the IPv6 data, the rules engine uses UDP intrusion rules to analyze both the inner and outer UDP layers.

Note that intrusion rules 12065, 12066, 12067, and 12068 in the **policy-other** rule category detect, but do not decode, Teredo traffic. Optionally, you can use these rules to drop Teredo traffic in an inline deployment; however, you should ensure that these rules are disabled or set to generate events without dropping traffic when you enable **Detect Teredo on Non-Standard Ports**.

Detect Excessive Length Value

Detects when the packet header specifies a packet length that is greater than the actual packet length.

This option is ignored for Firepower Threat Defense routed, transparent, and inline interfaces. Packets that have excessive header length are always dropped. However, this option does apply to Firepower Threat Defense inline tap and passive interfaces.

You can enable rules 116:6, 116:47, 116:97, and 116:275 to generate events and, in an inline deployment, drop offending packets for this option.

Detect Invalid IP Options

Detects invalid IP header options to identify exploits that use invalid IP options. For example, there is a denial of service attack against a firewall which causes the system to freeze. The firewall attempts to parse invalid Timestamp and Security IP options and fails to check for a zero length, which causes an irrecoverable infinite loop. The rules engine identifies the zero length option, and provides information you can use to mitigate the attack at the firewall.

Firepower Threat Defense devices will drop any RSVP packet that contains IP options other than the router alert, end of options list (EOOL), and no operation (NOP) options on any routed or transparent interface. For inline, inline tap, or passive interfaces, IP options will be handled as described above.

You can enable rules 116:4 and 116:5 to generate events and, in an inline deployment, drop offending packets for this option.

Detect Experimental TCP Options

Detects TCP headers with experimental TCP options. The following table describes these options.

TCP Option	Description
9	Partial Order Connection Permitted
10	Partial Order Service Profile
14	Alternate Checksum Request
15	Alternate Checksum Data
18	Trailer Checksum
20	Space Communications Protocol Standards (SCPS)
21	Selective Negative Acknowledgements (SCPS)
22	Record Boundaries (SCPS)
23	Corruption (SPCS)
24	SNAP
26	TCP Compression Filter

Because these are experimental options, some systems do not account for them and may be open to exploits.



Note In addition to the experimental options listed in the above table, the system considers any TCP option with an option number greater than 26 to be experimental.

You can enable rule 116:58 to generate events and, in an inline deployment, drop offending packets for this option.

Detect Obsolete TCP Options

Detects TCP headers with obsolete TCP options. Because these are obsolete options, some systems do not account for them and may be open to exploits. The following table describes these options.

TCP Option	Description
6	Echo
7	Echo Reply
16	Skeeter
17	Bubba
19	MD5 Signature
25	Unassigned

You can enable rule 116:57 to generate events and, in an inline deployment, drop offending packets for this option.

Detect T/TCP

Detects TCP headers with the CC.ECHO option. The CC.ECHO option confirms that TCP for Transactions (T/TCP) is being used. Because T/TCP header options are not in widespread use, some systems do not account for them and may be open to exploits.

You can enable rule 116:56 to generate events and, in an inline deployment, drop offending packets for this option.

Detect Other TCP Options

Detects TCP headers with invalid TCP options not detected by other TCP decoding event options. For example, this option detects TCP options with the incorrect length or with a length that places the option data outside the TCP header.

This option is ignored for Firepower Threat Defense routed and transparent interfaces. Packets that have invalid TCP options are always dropped.

You can enable rules 116:54, 116:55, and 116:59 to generate events and, in an inline deployment, drop offending packets for this option.

Detect Protocol Header Anomalies

Detects other decoding errors not detected by the more specific IP and TCP decoder options. For example, the decoder might detect a malformed data-link protocol header.

This option is ignored for Firepower Threat Defense routed, transparent, and inline interfaces. Packets that have header anomalies are always dropped. However, this option does apply to Threat Defense inline tap and passive interfaces.

To generate events and, in an inline deployment, drop offending packets for this option, you can enable any of the following rules:

GID:SID	Generates an event if:
116:467	The packet is smaller than the minimum size of a packet encapsulated with a Cisco FabricPath header.
116:468	The Cisco Meta Data (CMD) field in the header contains a header length smaller than the minimum size of a valid CMD header. The CMD field is associated with the Cisco Trustsec protocol.
116:469	The CMD field in the header contains an invalid field length.
116:470	The CMD field in the header contains an invalid Security Group Tag (SGT) option type.
116:471	The CMD field in the header contains an SGT with a reserved value.

You can also enable any packet decoder rule not associated with other packet decoder options.

Related Topics

Predefined Default Variables, on page 414

Configuring Packet Decoding

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation panel.

Step 4 If **Packet Decoding** under **Transport/Network Layer Preprocessors** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **Packet Decoding**.

Step 6 Enable or disable the options described in Packet Decoder Options, on page 1774.

Step 7 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, cached changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want to generate events and, in an inline deployment, drop offending packets, enable packet decoder rules (GID 116). For more information, see Setting Intrusion Rule States, on page 1504 and Packet Decoder Options, on page 1774.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Layer Basics, on page 1463
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

TCP Stream Preprocessing

The TCP protocol defines various states in which connections can exist. Each TCP connection is identified by the source and destination IP addresses and source and destination ports. TCP permits only one connection with the same connection parameter values to exist at a time.

State-Related TCP Exploits

If you add the `flow` keyword with the `established` argument to an intrusion rule, the intrusion rules engine inspects packets matching the rule and the flow directive in stateful mode. Stateful mode evaluates only the traffic that is part of a TCP session established with a legitimate three-way handshake between a client and server.

You can configure the system so that the preprocessor detects any TCP traffic that cannot be identified as part of an established TCP session, although this is not recommended for typical use because the events would quickly overload the system and not provide meaningful data.

Attacks like stick and snot use the system's extensive rule sets and packet inspection against itself. These tools generate packets based on the patterns in Snort-based intrusion rules, and send them across the network. If your rules do not include the `flow` or `flowbits` keyword to configure them for stateful inspection, each packet will trigger the rule, overwhelming the system. Stateful inspection allows you to ignore these packets because they are not part of an established TCP session and do not provide meaningful information. When performing stateful inspection, the rules engine detects only those attacks that are part of an established TCP session, allowing analysts to focus on these rather than the volume of events caused by stick or snot.

Target-Based TCP Policies

Different operating systems implement TCP in different ways. For example, Windows and some other operating systems require a TCP reset segment to have a precise TCP sequence number to reset a session, while Linux and other operating systems permit a range of sequence numbers. In this example, the stream preprocessor must understand exactly how the destination host will respond to the reset based on the sequence number. The stream preprocessor stops tracking the session only when the destination host considers the reset to be valid, so an attack cannot evade detection by sending packets after the preprocessor stops inspecting the

stream. Other variations in TCP implementations include such things as whether an operating system employs a TCP timestamp option and, if so, how it handles the timestamp, and whether an operating system accepts or ignores data in a SYN packet.

Different operating systems also reassemble overlapping TCP segments in different ways. Overlapping TCP segments could reflect normal retransmissions of unacknowledged TCP traffic. They could also represent an attempt by an attacker, aware of the operating system of one of your hosts, to evade detection and exploit that host by sending malicious content hidden in overlapping segments. However, you can configure the stream preprocessor to be aware of the operating systems running on your monitored network segment so it reassembles segments the same way the target host does, allowing it to identify the attack.

You can create one or more TCP policies to tailor TCP stream inspection and reassembly to the different operating systems on your monitored network segment. For each policy, you identify one of thirteen operating system policies. You bind each TCP policy to a specific IP address or address block using as many TCP policies as you need to identify any or all of the hosts using a different operating system. The default TCP policy applies to any hosts on the monitored network that you do not identify in any other TCP policy, so there is no need to specify an IP address or address block for the default TCP policy.

Note that you can also use adaptive profile updates in a passive deployment to dynamically select target-based policies for the TCP stream preprocessor using host operating system information for the target host in a packet.

TCP Stream Reassembly

The stream preprocessor collects and reassembles all the packets that are part of a TCP session's server-to-client communication stream, client-to-server communication stream, or both. This allows the rules engine to inspect the stream as a single, reassembled entity rather than inspecting only the individual packets that are part of a given stream.

Stream reassembly allows the rules engine to identify stream-based attacks, which it may not detect when inspecting individual packets. You can specify which communication streams the rules engine reassembles based on your network needs. For example, when monitoring traffic on your web servers, you may only want to inspect client traffic because you are much less likely to receive malicious traffic from your own web server.

In each TCP policy, you can specify a comma-separated list of ports to identify the traffic for the stream preprocessor to reassemble. If adaptive profile updates are enabled, you can also list services that identify traffic to reassemble, either as an alternative to ports or in combination with ports.

You can specify ports, services, or both. You can specify separate lists of ports for any combination of client ports, server ports, and both. You can also specify separate lists of services for any combination of client services, server services, and both. For example, assume that you wanted to reassemble the following:

- SMTP (port 25) traffic from the client
- FTP server responses (port 21)
- telnet (port 23) traffic in both directions

You could configure the following:

- For client ports, specify `23, 25`
- For server ports, specify `21, 23`

Or, instead, you could configure the following:

- For client ports, specify `25`
- For server ports, specify `21`
- For both ports, specify `23`

Additionally, consider the following example which combines ports and services and would be valid when adaptive profile updates are enabled:

- For client ports, specify `23`
- For client services, specify `smtp`
- For server ports, specify `21`
- For server services, specify `telnet`

Negating a port (for example, `!80`) can improve performance by preventing the TCP stream preprocessor from processing traffic for that port.

Although you can also specify `all` as the argument to provide reassembly for all ports, Cisco does **not** recommend setting ports to `all` because it may increase the amount of traffic inspected by this preprocessor and slow performance unnecessarily.

TCP reassembly automatically and transparently includes ports that you add to other preprocessors. However, if you do explicitly add ports to TCP reassembly lists that you have added to other preprocessor configurations, these additional ports are handled normally. This includes port lists for the following preprocessors:

- FTP/Telnet (server-level FTP)
- DCE/RPC
- HTTP Inspect
- SMTP
- Session Initiation Protocol
- POP
- IMAP
- SSL

Note that reassembling additional traffic types (client, server, both) increases resource demands.

TCP Stream Preprocessing Options

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

You can configure the following global TCP option:

Packet Type Performance Boost

Enables ignoring TCP traffic for all ports and application protocols that are not specified in enabled intrusion rules, except when a TCP rule with both the source and destination ports set to `any` has a `flow` or `flowbits` option. This performance improvement could result in missed attacks.

You can configure the following options for each TCP policy.

Network

Specifies the host IP addresses to which you want to apply the TCP stream reassembly policy.

You can specify a single IP address or address block. You can specify up to 255 total profiles including the default policy.

Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Note that the `default` setting in the default policy specifies all IP addresses on your monitored network segment that are not covered by another target-based policy. Therefore, you cannot and do not need to specify an IP address or CIDR block/prefix length for the default policy, and you cannot leave this setting blank in another policy or use address notation to represent `any` (for example, 0.0.0.0/0 or ::/0).

Policy

Identifies the TCP policy operating system of the target host or hosts. If you select a policy other than **Mac OS**, the system removes the data from the synchronization (SYN) packets and disables event generation for rule 129:2. Note that enabling the inline normalization preprocessor **Remove Data on SYN** option also disables rule 129:2.

The following table identifies the operating system policies and the host operating systems that use each.

Table 209: TCP Operating System Policies

Policy	Operating Systems
First	unknown OS
Last	Cisco IOS
BSD	AIX FreeBSD OpenBSD
Linux	Linux 2.4 kernel Linux 2.6 kernel
Old Linux	Linux 2.2 and earlier kernel
Windows	Windows 98 Windows NT Windows 2000 Windows XP
Windows 2003	Windows 2003

Policy	Operating Systems
Windows Vista	Windows Vista
Solaris	Solaris OS SunOS
IRIX	SGI Irix
HPUX	HP-UX 11.0 and later
HPUX 10	HP-UX 10.2 and earlier
Mac OS	Mac OS 10 (Mac OS X)



Tip The First operating system policy could offer some protection when you do not know the host operating system. However, it may result in missed attacks. You should edit the policy to specify the correct operating system if you know it.

Timeout

The number of seconds between 1 and 86400 the intrusion rules engine keeps an inactive stream in the state table. If the stream is not reassembled in the specified time, the intrusion rules engine deletes it from the state table.



Note If your managed device is deployed on a segment where the network traffic is likely to reach the device's bandwidth limits, you should consider setting this value higher (for example, to 600 seconds) to lower the amount of processing overhead.

Firepower Threat Defense devices use this option only for connections that are inspected by Snort. For other connections, you need to configure a global TCP timeout in your platform settings policy.

Maximum TCP Window

Specifies the maximum TCP window size between 1 and 1073725440 bytes allowed as specified by a receiving host. Setting the value to 0 disables checking for the TCP window size.



Caution The upper limit is the maximum window size permitted by RFC, and is intended to prevent an attacker from evading detection, but setting a significantly large maximum window size could result in a self-imposed denial of service.

When **Stateful Inspection Anomalies** is enabled, you can enable rule 129:6 to generate events and, in an inline deployment, drop offending packets for this option.

Overlap Limit

Specifies that when the configured number between 0 (unlimited) and 255 of overlapping segments in a session has been detected, segment reassembly stops for that session and, if **Stateful Inspection Anomalies** is enabled and the accompanying preprocessor rule is enabled, an event is generated.

You can enable rule 129:7 to generate events and, in an inline deployment, drop offending packets for this option.

Flush Factor

In an inline deployment, specifies that when a segment of decreased size has been detected subsequent to the configured number between 1 and 2048 of segments of non-decreasing size, the system flushes segment data accumulated for detection. Setting the value to 0 disables detection of this segment pattern, which can indicate the end of a request or response. Note that the Inline Normalization **Normalize TCP Payload** option must be enabled for this option the be effective.

Stateful Inspection Anomalies

Detects anomalous behavior in the TCP stack. When accompanying preprocessor rules are enabled, this may generate many events if TCP/IP stacks are poorly written.

This option is ignored for Firepower Threat Defense routed and transparent interfaces.

You can enable the following rules to generate events and, in an inline deployment, drop offending packets for this option:

- 129:1 through 129:5
- 129:6 (Mac OS only)
- 129:8 through 129:11
- 129:13 through 129:19

Note the following:

- for rule 129:6 to trigger you must also configure a value greater than 0 for **Maximum TCP Window**.
- for rules 129:9 and 129:10 to trigger you must also enable **TCP Session Hijacking**.

TCP Session Hijacking

Detects TCP session hijacking by validating the hardware (MAC) addresses detected from both sides of a TCP connection during the 3-way handshake against subsequent packets received on the session. When the MAC address for one side or the other does not match, if **Stateful Inspection Anomalies** is enabled and one of the two corresponding preprocessor rules are enabled, the system generates events.

This option is ignored for Firepower Threat Defense routed and transparent interfaces.

You can enable rules 129:9 and 129:10 to generate events and, in an inline deployment, drop offending packets for this option. Note that for either of these rules to generate events you must also enable **Stateful Inspection Anomalies**.

Consecutive Small Segments

When **Stateful Inspection Anomalies** is enabled, specifies a maximum number of 1 to 2048 consecutive small TCP segments allowed. Setting the value to 0 disables checking for consecutive small segments.

You must set this option together with the **Small Segment Size** option, either disabling both or setting a non-zero value for both. Note that receiving as many as 2000 consecutive segments, even if each segment was 1 byte in length, without an intervening ACK would be far more consecutive segments than you would normally expect.

This option is ignored for Firepower Threat Defense routed and transparent interfaces.

You can enable rule 129:12 to generate events and, in an inline deployment, drop offending packets for this option.

Small Segment Size

When **Stateful Inspection Anomalies** is enabled, specifies the 1 to 2048 byte TCP segment size that is considered small. Setting the value to 0 disables specifying the size of a small segment.

This option is ignored for Firepower Threat Defense routed and transparent interfaces.

You must set this option together with the **Consecutive Small Segments** option, either disabling both or setting a non-zero value for both. Note that a 2048 byte TCP segment is larger than a normal 1500 byte Ethernet frame.

Ports Ignoring Small Segments

When **Stateful Inspection Anomalies**, **Consecutive Small Segments**, and **Small Segment Size** are enabled, specifies a comma-separated list of one or more ports that ignore small TCP segment detection. Leaving this option blank specifies that no ports are ignored.

This option is ignored for Firepower Threat Defense routed and transparent interfaces.

You can add any port to the list, but the list only affects ports specified in one of the **Perform Stream Reassembly on** port lists in the TCP policy.

Require TCP 3-Way Handshake

Specifies that sessions are treated as established only upon completion of a TCP three-way handshake. Disable this option to increase performance, protect from SYN flood attacks, and permit operation in a partially asynchronous environment. Enable it to avoid attacks that attempt to generate false positives by sending information that is not part of an established TCP session.

You can enable rule 129:20 to generate events and, in an inline deployment, drop offending packets for this option.

3-Way Handshake Timeout

Specifies the number of seconds between 0 (unlimited) and 86400 (twenty-four hours) by which a handshake must be completed when **Require TCP 3-Way Handshake** is enabled. You must enable **Require TCP 3-Way Handshake** to modify the value for this option.

For Firepower Software devices and Firepower Threat Defense inline, inline tap, and passive interfaces, the default is 0. For Firepower Threat Defense routed and transparent interfaces, the timeout is always 30 seconds; the value configured here is ignored.

Packet Size Performance Boost

Sets the preprocessor to not queue large packets in the reassembly buffer. This performance improvement could result in missed attacks. Disable this option to protect against evasion attempts using small packets of one to twenty bytes. Enable it when you are assured of no such attacks because all traffic is comprised of very large packets.

Legacy Reassembly

Sets the stream preprocessor to emulate the deprecated Stream 4 preprocessor when reassembling packets, which lets you compare events reassembled by the stream preprocessor to events based on the same data stream reassembled by the Stream 4 preprocessor.

Asynchronous Network

Specifies whether the monitored network is an asynchronous network, that is, a network where the system sees only half the traffic. When this option is enabled, the system does not reassemble TCP streams to increase performance.

This option is ignored for Firepower Threat Defense routed and transparent interfaces.

Perform Stream Reassembly on Client Ports

Enables stream reassembly based on ports for the client side of the connection. In other words, it reassembles streams destined for web servers, mail servers, or other IP addresses typically defined by the IP addresses specified in $HOME_NET. Use this option when you expect malicious traffic to originate from clients.

This option is ignored for Firepower Threat Defense routed and transparent interfaces.

Perform Stream Reassembly on Client Services

Enables stream reassembly based on services for the client side of the connection. Use this option when you expect malicious traffic to originate from clients.

At least one client detector must be enabled for each client service you select. By default, all Cisco-provided detectors are activated. If no detector is enabled for an associated client application, the system automatically enables all Cisco-provided detectors for the application; if none exist, the system enables the most recently modified user-defined detector for the application.

This feature requires Protection and Control licenses.

This option is ignored for Firepower Threat Defense routed and transparent interfaces.

Perform Stream Reassembly on Server Ports

Enables stream reassembly based on ports for the server side of the connection only. In other words, it reassembles streams originating from web servers, mail servers, or other IP addresses typically defined by the IP addresses specified in $EXTERNAL_NET. Use this option when you want to watch for server side attacks. You can disable this option by not specifying ports.

This option is ignored for Firepower Threat Defense routed and transparent interfaces.

Note For a thorough inspection of a service, add the service name in the Perform Stream Reassembly on Server Services field in addition to adding the port number in the Perform Stream Reassembly on Server Ports field. For example, add `'HTTP'` service in the Perform Stream Reassembly on Server Services field to inspect HTTP service in addition to adding port number 80 in the Perform Stream Reassembly on Server Ports field.

Perform Stream Reassembly on Server Services

Enables stream reassembly based on services for the server side of the connection only. Use this option when you want to watch for server side attacks. You can disable this option by not specifying services.

At least one detector must be enabled. By default, all Cisco-provided detectors are activated. If no detector is enabled for a service, the system automatically enables all Cisco-provided detectors for the associated application protocol; if none exist, the system enables the most recently modified user-defined detector for the application protocol.

This feature requires Protection and Control licenses.

This option is ignored for Firepower Threat Defense routed and transparent interfaces.

Perform Stream Reassembly on Both Ports

Enables stream reassembly based on ports for both the client and server side of the connection. Use this option when you expect that malicious traffic for the same ports may travel in either direction between clients and servers. You can disable this option by not specifying ports.

This option is ignored for Firepower Threat Defense routed and transparent interfaces.

Perform Stream Reassembly on Both Services

Enables stream reassembly based on services for both the client and server side of the connection. Use this option when you expect that malicious traffic for the same services may travel in either direction between clients and servers.You can disable this option by not specifying services.

At least one detector must be enabled. By default, all Cisco-provided detectors are activated. If no detector is enabled for an associated client application or application protocol, the system automatically enables all Cisco-provided detectors for the application or application protocol; if none exist, the system enables the most recently modified user-defined detector for the application or application protocol.

This feature requires Protection and Control licenses.

This option is ignored for Firepower Threat Defense routed and transparent interfaces.

Troubleshooting Options: Maximum Queued Bytes

Support might ask you during a troubleshooting call to specify the amount of data that can be queued on one side of a TCP connection. A value of 0 specifies an unlimited number of bytes.

Caution Changing the setting for this troubleshooting option will affect performance and should be done only with Support guidance.

Troubleshooting Options: Maximum Queued Segments

Support might ask you during a troubleshooting call to specify the maximum number of bytes of data segments that can be queued on one side of a TCP connection. A value of 0 specifies an unlimited number of data segment bytes.

Caution Changing the setting for this troubleshooting option will affect performance and should be done only with Support guidance.

Related Topics

Firepower System IP Address Conventions, on page 15
Activating and Deactivating Detectors, on page 1889
Layer Management, on page 1468
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461

Configuring TCP Stream Preprocessing

The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Before you begin

- Confirm that networks you want to identify in a custom target-based policy match or are a subset of the networks, zones, and VLANs handled by its parent network analysis policy. See Advanced Settings for Network Analysis Policies, on page 1672 for more information.

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to modify.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation panel on the left.

Step 4 If the **TCP Stream Configuration** setting is disabled under Transport/Network Layer Preprocessors, enable it by clicking **Enabled**.

Step 5 Click **Edit** ( ) next to **TCP Stream Configuration**.

Step 6 Check or clear the **Packet Type Performance Boost** check box in the **Global Settings** section.

Step 7 You can:

- Add a target-based policy — Click **Add** () next to **Hosts** in the Targets section. Specify one or more IP addresses in the **Host Address** field. You can specify a single IP address or address block. You can create a total of 255 target-based policies including the default policy. When done, click **OK**.
- Edit an exist target-based policy — Under **Hosts**, click on the address for the policy you want to edit, or click default to edit the **default** configuration values.
- Modify the TCP Stream Preprocessing options — See TCP Stream Preprocessing Options, on page 1780.

Caution Do not modify **Maximum Queued Bytes** or **Maximum Queued Segments** unless instructed to do so by Support.

Tip To modify stream reassembly settings based on client, server, or both services, click inside the field you want to modify or click **Edit** next to the field. Use arrow to move services between the **Available** and **Enabled** lists in the pop-up window, then click **OK**.

- Delete an existing target-based policy — Click **Delete** () next to the policy you want to remove.

Step 8 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, cached changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want to generate events and, in an inline deployment, drop offending packets, enable TCP Stream preprocessor rules (GID 129). For more information, see Setting Intrusion Rule States, on page 1504 and TCP Stream Preprocessing Options, on page 1780.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Layer Management, on page 1468
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461
Firepower System IP Address Conventions, on page 15

UDP Stream Preprocessing

UDP stream preprocessing occurs when the rules engine processes packets against a UDP rule that includes the `flow` keyword using any of the following arguments:

- `Established`
- `To Client`
- `From Client`
- `To Server`
- `From Server`

UDP data streams are not typically thought of in terms of *sessions*. UDP is a connectionless protocol that does not provide a means for two endpoints to establish a communication channel, exchange data, and close the channel. However, the stream preprocessor uses the source and destination IP address fields in the encapsulating IP datagram header and the port fields in the UDP header to determine the direction of flow and identify a session. A session ends when a configurable timer is exceeded, or when either endpoint receives an ICMP message that the other endpoint is unreachable or the requested service is unavailable.

Note that the system does not generate events related to UDP stream preprocessing; however, you can enable related packet decoder rules to detect UDP protocol header anomalies.

Related Topics

TCP Header Values and Stream Size, on page 1605

UDP Stream Preprocessing Options

Timeout

Specifies the number of seconds the preprocessor keeps an inactive stream in the state table. If additional datagrams are not seen in the specified time, the preprocessor deletes the stream from the state table.

Firepower Threat Defense devices use this option only for connections that are inspected by Snort. For other connections, you need to configure a global UDP timeout in your platform settings policy.

Packet Type Performance Boost

Sets to preprocessor to ignore UDP traffic for all ports and application protocols that are not specified in enabled rules, except when a UDP rule with both the source and destination ports set to `any` has a `flow` or `flowbits` option. This performance improvement could result in missed attacks.

Configuring UDP Stream Preprocessing

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation panel.

Step 4 If **UDP Stream Configuration** under **Transport/Network Layer Preprocessors** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **UDP Stream Configuration**.

Step 6 Set the options described in UDP Stream Preprocessing Options, on page 1789.

Step 7 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

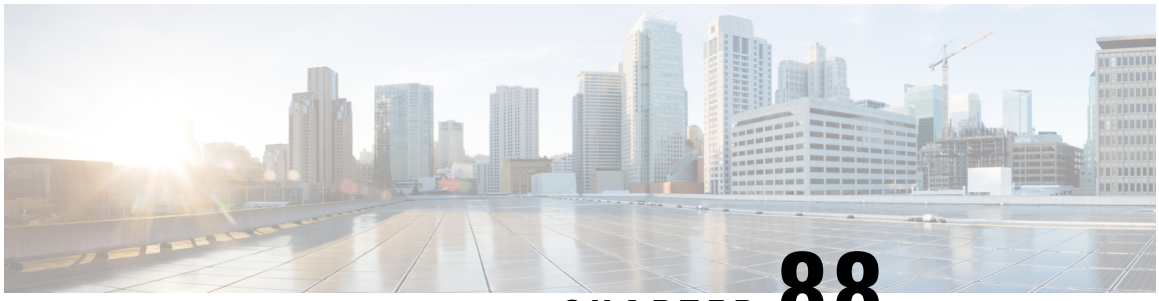
If you leave the policy without committing changes, cached changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want to generate events and, in an inline deployment, drop offending packets, enable related packet decoder rules (GID 116). For more information, see Setting Intrusion Rule States, on page 1504 and The Packet Decoder, on page 1774.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Layer Management, on page 1468
Conflicts and Changes: Network Analysis and Intrusion Policies, on page 1461



CHAPTER 88

Detecting Specific Threats

The following topics explain how to use preprocessors in a network analysis policy to detect specific threats:

- Introduction to Specific Threat Detection, on page 1791
- License Requirements for Specific Threat Detection, on page 1791
- Requirements and Prerequisites for Specific Threat Detection, on page 1792
- Back Orifice Detection, on page 1792
- Portscan Detection, on page 1793
- Rate-Based Attack Prevention, on page 1800

Introduction to Specific Threat Detection

You can use several preprocessors in a network analysis policy to detect specific threats to your monitored network, such as Back Orifice attacks, several portscan types, and rate-based attacks that attempt to overwhelm your network with excessive traffic. When the GID Signatures specific to pre-processor is enabled, the Network Analysis Policy on Web will show disabled. However, the pre-processors will be turned on device using the available default settings.

You can also use sensitive data detection, which you configure in an intrusion policy, to detect unsecured transmission of sensitive numerical data.

License Requirements for Specific Threat Detection

FTD License

Threat

Classic License

Protection

Requirements and Prerequisites for Specific Threat Detection

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Intrusion Admin

Back Orifice Detection

The Firepower System provides a preprocessor that detects the existence of the Back Orifice program. This program can be used to gain admin access to your Windows hosts.

Back Orifice Detection Preprocessor

The Back Orifice preprocessor analyzes UDP traffic for the Back Orifice magic cookie, "`*!*QWTY?`", which is located in the first eight bytes of the packet and is XOR-encrypted.

The Back Orifice preprocessor has a configuration page, but no configuration options. When it is enabled, you must also enable preprocessor rules for the preprocessor to generate events and, in an inline deployment, drop offending packets.

Table 210: Back Orifice GID:SIDs

Preprocessor rule GID:SID	Description
105:1	Back Orifice traffic detected
105:2	Back Orifice client traffic detected
105:3	Back Orifice server traffic detected
105:4	Back Orifice Snort buffer attack detected

Detecting Back Orifice

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings** in the navigation panel.

Step 4 If **Back Orifice Detection** under **Specific Threat Detection** is disabled, click **Enabled**.

Note There are no user-configurable options for Back Orifice.

Step 5 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want to generate events and, in an inline deployment, drop offending packets, enable Back Orifice Detection rules 105:1, 105:2, 105:3, or 105:4. For more information, see Intrusion Rule States, on page 1503 and Back Orifice Detection Preprocessor, on page 1792.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Portscan Detection

A portscan is a form of network reconnaissance that is often used by attackers as a prelude to an attack. In a portscan, an attacker sends specially crafted packets to a targeted host. By examining the packets that the host responds with, the attacker can often determine which ports are open on the host and, either directly or by inference, which application protocols are running on these ports.

By itself, a portscan is not evidence of an attack. In fact, some of the portscanning techniques used by attackers can also be employed by legitimate users on your network. Cisco's portscan detector is designed to help you determine which portscans might be malicious by detecting patterns of activity.

Portscan Types, Protocols, and Filtered Sensitivity Levels

Attackers are likely to use several methods to probe your network. Often they use different protocols to draw out different responses from a target host, hoping that if one type of protocol is blocked, another may be available.

Table 211: Protocol Types

Protocol	Description
TCP	Detects TCP probes such as SYN scans, ACK scans, TCP connect() scans, and scans with unusual flag combinations such as Xmas tree, FIN, and NULL
UDP	Detects UDP probes such as zero-byte UDP packets
ICMP	Detects ICMP echo requests (pings)
IP	Detects IP protocol scans. These scans differ from TCP and UDP scans because the attacker, instead of looking for open ports, is trying to discover which IP protocols are supported on a target host.

Portscans are generally divided into four types based on the number of targeted hosts, the number of scanning hosts, and the number of ports that are scanned.

Table 212: Portscan Types

Type	Description
Portscan Detection	A one-to-one portscan in which an attacker uses one or a few hosts to scan multiple ports on a single target host. One-to-one portscans are characterized by: • a low number of scanning hosts • a single host that is scanned • a high number of ports scanned This option detects TCP, UDP, and IP portscans.
Port Sweep	A one-to-many portsweep in which an attacker uses one or a few hosts to scan a single port on multiple target hosts. Portsweeps are characterized by: • a low number of scanning hosts • a high number of scanned hosts • a low number of unique ports scanned This option detects TCP, UDP, ICMP, and IP portsweeps.

Type	Description
Decoy Portscan	A one-to-one portscan in which the attacker mixes spoofed source IP addresses with the actual scanning IP address. Decoy portscans are characterized by: • a high number of scanning hosts • a low number of ports that are scanned only once • a single (or a low number of) scanned hosts The decoy portscan option detects TCP, UDP, and IP protocol portscans.
Distributed Portscan	A many-to-one portscan in which multiple hosts query a single host for open ports. Distributed portscans are characterized by: • a high number of scanning hosts • a high number of ports that are scanned only once • a single (or a low number of) scanned hosts The distributed portscan option detects TCP, UDP, and IP protocol portscans.

The information that the portscan detector learns about a probe is largely based on seeing negative responses from the probed hosts. For example, when a web client tries to connect to a web server, the client uses port 80/tcp and the server can be counted on to have that port open. However, when an attacker probes a server, the attacker does not know in advance if it offers web services. When the portscan detector sees a negative response (that is, an ICMP unreachable or TCP RST packet), it records the response as a potential portscan. The process is more difficult when the targeted host is on the other side of a device such as a firewall or router that filters negative responses. In this case, the portscan detector can generate *filtered* portscan events based on the sensitivity level that you select.

Table 213: Sensitivity Levels

Level	Description
Low	Detects only negative responses from targeted hosts. Select this sensitivity level to suppress false positives, but keep in mind that some types of portscans (slow scans, filtered scans) might be missed. This level uses the shortest time window for portscan detection.
Medium	Detects portscans based on the number of connections to a host, which means that you can detect filtered portscans. However, very active hosts such as network address translators and proxies may generate false positives. Note that you can add the IP addresses of these active hosts to the **Ignore Scanned** field to mitigate this type of false positive. This level uses a longer time window for portscan detection.

Level	Description
High	Detects portscans based on a time window, which means that you can detect time-based portscans. However, if you use this option, you should be careful to tune the detector over time by specifying IP addresses in the **Ignore Scanned** and **Ignore Scanner** fields. This level uses a much longer time window for portscan detection.

Portscan Event Generation

When portscan detection is enabled, you must enable rules with Generator ID (GID) 122 and a Snort ID (SID) from among SIDs 1 through 27 to detect the various portscans and portsweeps.



Note For events generated by the portscan connection detector, the protocol number is set to 255. Because portscan does not have a specific protocol associated with it by default, the Internet Assigned Numbers Authority (IANA) does not have a protocol number assigned to it. IANA designates 255 as a reserved number, so that number is used in portscan events to indicate that there is not an associated protocol for the event.

Table 214: Portscan Detection SIDs (GID 122)

Portscan Type	Protocol	Sensitivity Level	Preprocessor Rule SID
Portscan Detection	TCP	Low	1
	UDP	Medium or High	5
	ICMP	Low	17
	IP	Medium or High	21
		Low	Does not generate events.
		Medium or High	Does not generate events.
		Low	9
		Medium or High	13
Port Sweep	TCP	Low	3, 27
	UDP	Medium or High	7
	ICMP	Low	19
	IP	Medium or High	23
		Low	25
		Medium or High	26
		Low	11
		Medium or High	15

Portscan Type	Protocol	Sensitivity Level	Preprocessor Rule SID
Decoy Portscan	TCP	Low	2
	UDP	Medium or High	6
	ICMP	Low	18
	IP	Medium or High	22
		Low	Does not generate events.
		Medium or High	Does not generate events.
		Low	10
		Medium or High	14
Distributed Portscan	TCP	Low	4
	UDP	Medium or High	8
	ICMP	Low	20
	IP	Medium or High	24
		Low	Does not generate events.
		Medium or High	Does not generate events.
		Low	12
		Medium or High	16

Portscan Event Packet View

When you enable the accompanying preprocessor rules, the portscan detector generates intrusion events that you can view just as you would any other intrusion event. However, the information presented on the packet view is different from the other types of intrusion events.

Begin by using the intrusion event views to drill down to the packet view for a portscan event. Note that you cannot download a portscan packet because single portscan events are based on multiple packets; however, the portscan packet view provides all usable packet information.

For any IP address, you can click the address to view the context menu and select **whois** to perform a lookup on the IP address or **View Host Profile** to view the host profile for that host.

Table 215: Portscan Packet View

Information	Description
Device	The device that detected the event.
Time	The time when the event occurred.
Message	The event message generated by the preprocessor.
Source IP	The IP address of the scanning host.

Information	Description
Destination IP	The IP address of the scanned host.
Priority Count	The number of negative responses (for example, TCP RSTs and ICMP unreachables) from the scanned host. The higher the number of negative responses, the higher the priority count.
Connection Count	The number of active connections on the hosts. This value is more accurate for connection-based scans such as TCP and IP.
IP Count	The number of times that the IP addresses that contact the scanned host changes. For example, if the first IP address is 10.1.1.1, the second IP is 10.1.1.2, and the third IP is 10.1.1.1, then the IP count is 3. This number is less accurate for active hosts such as proxies and DNS servers.
Scanner/Scanned IP Range	The range of IP addresses for the scanned hosts or the scanning hosts, depending on the type of scan. For portsweeps, this field shows the IP range of scanned hosts. For portscans, this shows the IP range of the scanning hosts.
Port/Proto Count	For TCP and UDP portscans, the number of times that the port being scanned changes. For example, if the first port scanned is 80, the second port scanned is 8080, and the third port scanned is again 80, then the port count is 3. For IP protocol portscans, the number of times that the protocol being used to connect to the scanned host changes.
Port/Proto Range	For TCP and UDP portscans, the range of the ports that were scanned. For IP protocol portscans, the range of IP protocol numbers that were used to attempt to connect to the scanned host.
Open Ports	The TCP ports that were open on the scanned host. This field appears only when the portscan detects one or more open ports.

Related Topics

About Intrusion Events, on page 2251

Configuring Portscan Detection

The portscan detection configuration options allow you to finely tune how the portscan detector reports scan activity.

The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings**.

Step 4 If **Portscan Detection** under **Specific Threat Detection** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **Portscan Detection**.

Step 6 In the **Protocol** field, specify protocols to enable.

Note You must ensure TCP stream processing is enabled to detect scans over TCP, and that UDP stream processing is enabled to detect scans over UDP.

Step 7 In the **Scan Type** field, specify portscan types you want to detect.

Step 8 Choose a level from the **Sensitivity Level** list; see Portscan Types, Protocols, and Filtered Sensitivity Levels, on page 1794.

Step 9 If you want to monitor specific hosts for signs of portscan activity, enter the host IP address in the **Watch IP** field.

You can specify a single IP address or address block, or a comma-separated lists of either or both. Leave the field blank to watch all network traffic.

Step 10 If you want to ignore hosts as scanners, enter the host IP address in the **Ignore Scanners** field.

You can specify a single IP address or address block, or a comma-separated lists of either or both.

Step 11 If you want to ignore hosts as targets of a scan, enter the host IP address in the **Ignore Scanned** field.

You can specify a single IP address or address block, or a comma-separated lists of either or both.

Tip Use the **Ignore Scanners** and **Ignore Scanned** fields to indicate hosts on your network that are especially active. You may need to modify this list of hosts over time.

Step 12 If you want to discontinue monitoring of sessions picked up in mid-stream, clear the **Detect Ack Scans** check box.

Note Detection of mid-stream sessions helps to identify ACK scans, but may cause false events, particularly on networks with heavy traffic and dropped packets.

Step 13 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- If you want portscan detection to detect various portscans and portsweeps, enable rules 122:1 through 122:27. For more information, see Intrusion Rule States, on page 1503 and Portscan Event Generation, on page 1796.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Firepower System IP Address Conventions, on page 15

Rate-Based Attack Prevention

Rate-based attacks are attacks that depend on frequency of connection or repeated attempts to perpetrate the attack. You can use rate-based detection criteria to detect a rate-based attack as it occurs and respond to it when it happens, then return to normal detection settings after it stops.

You can configure your network analysis policy to include rate-based filters that detect excessive activity directed at hosts on your network. You can use this feature on managed devices deployed in inline mode to block rate-based attacks for a specified time, then revert to only generating events and not drop traffic.

The SYN attack prevention option helps you protect your network hosts against SYN floods. You can protect individual hosts or whole networks based on the number of packets seen over a period of time. If your device is deployed passively, you can generate events. If your device is placed inline, you can also drop the malicious packets. After the timeout period elapses, if the rate condition has stopped, the event generation and packet dropping stops.

For example, you could configure a setting to allow a maximum number of SYN packets from any one IP address, and block further connections from that IP address for 60 seconds.

You can also limit TCP/IP connections to or from hosts on your network to prevent denial of service (DoS) attacks or excessive activity by users. When the system detects the configured number of successful connections to or from a specified IP address or range of addresses, it generates events on additional connections. The rate-based event generation continues until the timeout period elapses without the rate condition occurring. In an inline deployment you can choose to drop packets until the rate condition times out.

For example, you could configure a setting to allow a maximum of 10 successful simultaneous connections from any one IP address, and block further connections from that IP address for 60 seconds.

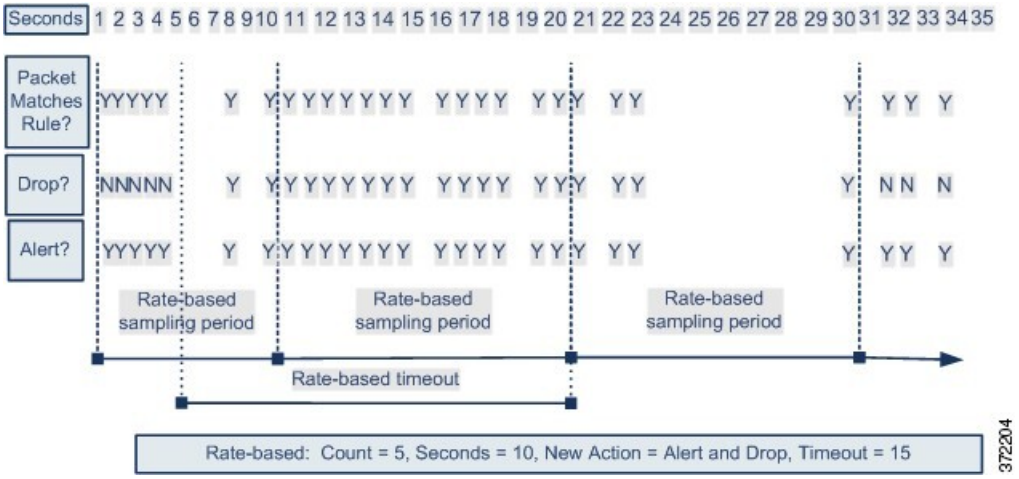


Note Devices load-balance inspection across internal resources. When you configure rate-based attack prevention, you configure the triggering rate per resource, not per device. If rate-based attack prevention is not working as expected, you may need to lower the triggering rate. It triggers alert, if users send too many connection attempts within prescribed time intervals. Hence it is recommended to rate limit the rule. For help determining the correct rate, contact Support.

The following diagram shows an example where an attacker is attempting to access a host. Repeated attempts to find a password trigger a rule which has rate-based attack prevention configured. The rate-based settings change the rule attribute to Drop and Generate Events after rule matches occur five times in a 10-second span. The new rule attribute times out after 15 seconds.

After the timeout, note that packets are still dropped in the rate-based sampling period that follows. If the sampled rate is above the threshold in the current or previous sampling period, the new action continues. The

new action reverts to generating events only after a sampling period completes where the sampled rate is below the threshold rate.



Related Topics

Dynamic Intrusion Rule States, on page 1511

Rate-Based Attack Prevention Examples

The `detection_filter` keyword and the thresholding and suppression features provide other ways to filter either the traffic itself or the events that the system generates. You can use rate-based attack prevention alone or in any combination with thresholding, suppression, or the `detection_filter` keyword.

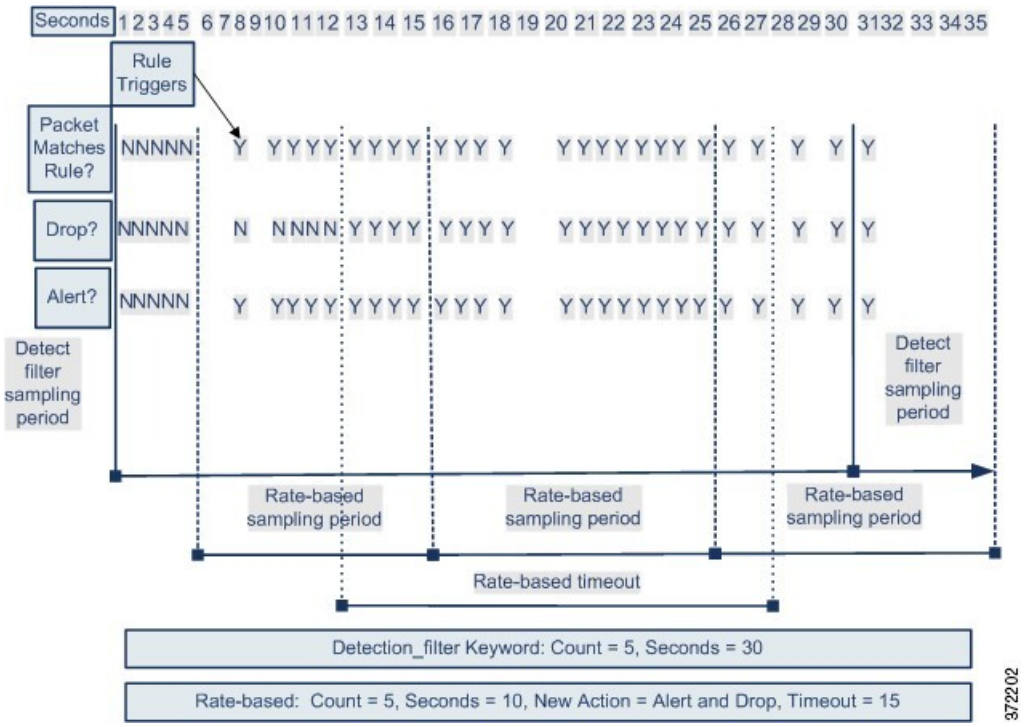
The `detection_filter` keyword, thresholding or suppression, and rate-based criteria may all apply to the same traffic. When you enable suppression for a rule, events are suppressed for the specified IP addresses even if a rate-based change occurs.

detection_filter Keyword Example

The following example shows an attacker attempting a brute-force login. Repeated attempts to find a password trigger a rule that also includes the `detection_filter` keyword, with a count set to 5. This rule has rate-based attack prevention configured. The rate-based settings change the rule attribute to Drop and Generate Events for 20 seconds when there are five hits on the rule in a 10-second span.

As shown in the diagram, the first five packets matching the rule do not generate events because the rule does not trigger until the rate exceeds the rate indicated by the `detection_filter` keyword. After the rule triggers, event notification begins, but the rate-based criteria do not trigger the new action of Drop and Generate Events until five more packets pass.

After the rate-based criteria are met, events are generated and the packets are dropped until the rate-based timeout period expires and the rate falls below the threshold. After twenty seconds elapse, the rate-based action times out. After the timeout, note that packets are still dropped in the rate-based sampling period that follows. Because the sampled rate is above the threshold rate in the previous sampling period when the timeout happens, the rate-based action continues.



Note that although the example does not depict this, you can use the Drop and Generate Events rule state in combination with the `detection_filter` keyword to start dropping traffic when hits for the rule reach the specified rate. When deciding whether to configure rate-based settings for a rule, consider whether setting the rule to Drop and Generate Events and including the `detection_filter` keyword would achieve the same result, or whether you want to manage the rate and timeout settings in the intrusion policy.

Related Topics

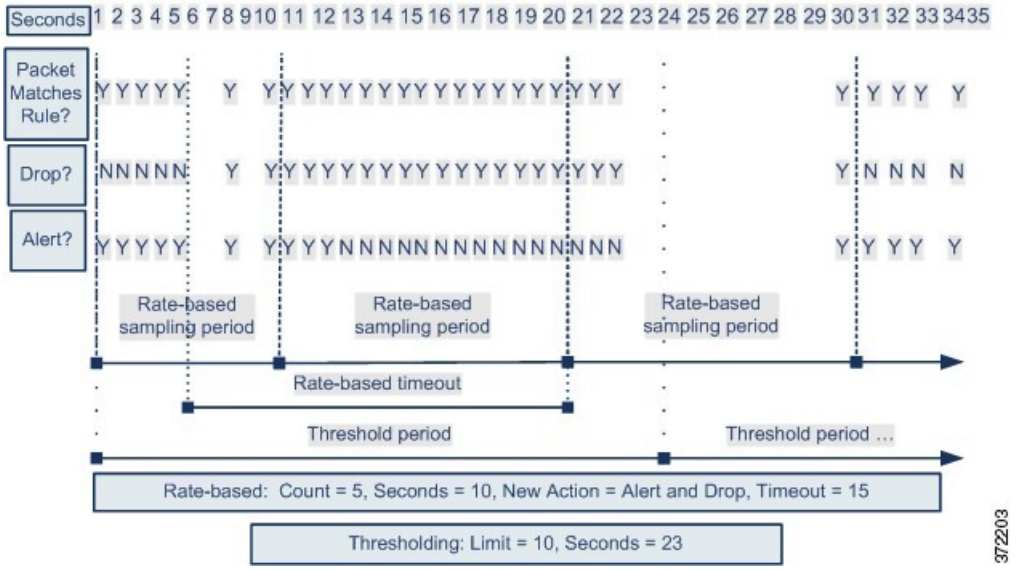
Intrusion Rule States, on page 1503

Dynamic Rule State Thresholding or Suppression Example

The following example shows an attacker attempting a brute-force login. Repeated attempts to find a password trigger a rule that has rate-based attack prevention configured. The rate-based settings change the rule attribute to Drop and Generate Events for 15 seconds when there are five hits on the rule in 10 seconds. In addition, a limit threshold limits the number of events the rule can generate to 10 events in 23 seconds.

As shown in the diagram, the rule generates events for the first five matching packets. After five packets, the rate-based criteria trigger the new action of Drop and Generate Events, and for the next five packets the rule generates events and the system drops the packet. After the tenth packet, the limit threshold has been reached, so for the remaining packets the system does not generate events but does drop the packets.

After the timeout, note that packets are still dropped in the rate-based sampling period that follows. If the sampled rate is above the threshold rate in the current or previous sampling period, the new action continues. The new action reverts to Generate Events only after a sampling period completes where the sampled rate is below the threshold rate.



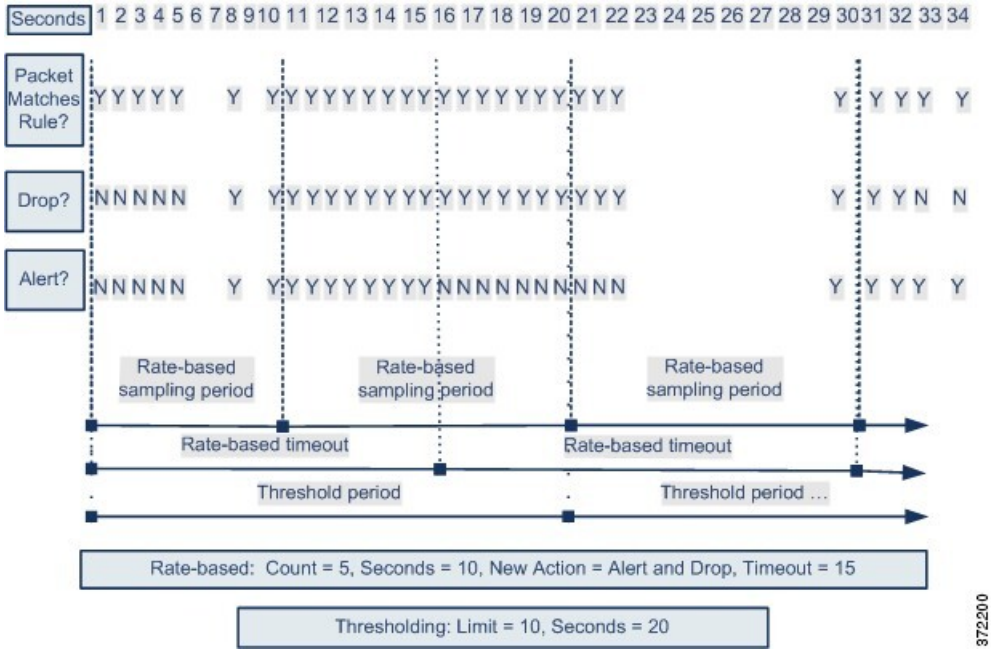
Note that although it is not shown in this example, if a new action triggers because of rate-based criteria *after* a threshold has been reached, the system generates a single event to indicate the change in action. So, for example, when the limit threshold of 10 is reached and the system stops generating events and the action changes from Generate Events to Drop and Generate Events on the 14th packet, the system generates an eleventh event to indicate the change in action.

Policy-Wide Rate-Based Detection and Thresholding or Suppression Example

The following example shows an attacker attempting denial of service (DoS) attacks on hosts in your network. Many simultaneous connections to hosts from the same sources trigger a policy-wide Control Simultaneous Connections setting. The setting generates events and drops malicious traffic when there are five connections from one source in 10 seconds. In addition, a global limit threshold limits the number of events any rule or setting can generate to 10 events in 20 seconds.

As shown in the diagram, the policy-wide setting generates events for the first ten matching packets and drops the traffic. After the tenth packet, the limit threshold is reached, so for the remaining packets no events are generated but the packets are dropped.

After the timeout, note that packets are still dropped in the rate-based sampling period that follows. If the sampled rate is above the threshold rate in the current or previous sampling period, the rate-based action of generating events and dropping traffic continues. The rate-based action stops only after a sampling period completes where the sampled rate is below the threshold rate.



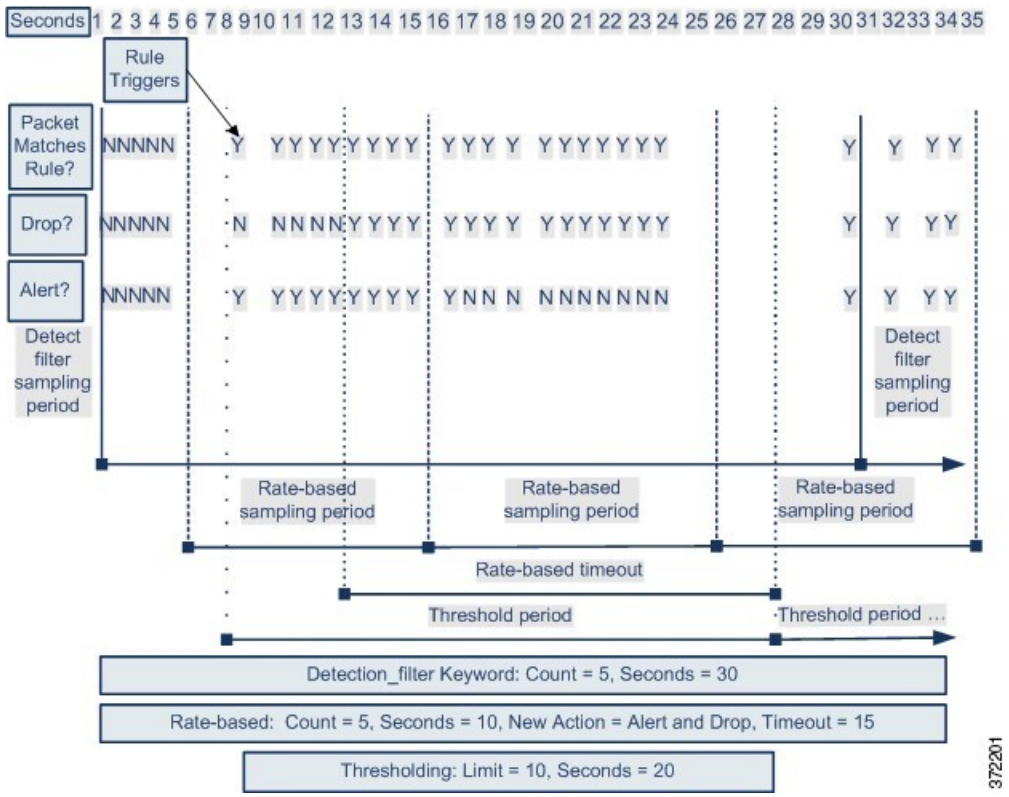
Note that although it is not shown in this example, if a new action triggers because of rate-based criteria *after* a threshold has been reached, the system generates a single event to indicate the change in action. So, for example, if the limit threshold of 10 has been reached and the system stops generating events and the action changes to Drop and Generate events on the 14th packet, the system generates an eleventh event to indicate the change in action.

Rate-Based Detection with Multiple Filtering Methods Example

The following example shows an attacker attempting a brute force login, and describes a case where a `detection_filter` keyword, rate-based filtering, and thresholding interact. Repeated attempts to find a password trigger a rule which includes the `detection_filter` keyword, with a count set to 5. This rule also has rate-based attack prevention settings that change the rule attribute to Drop and Generate Events for 30 seconds when there are five rule hits in 15 seconds. In addition, a limit threshold limits the rule to 10 events in 30 seconds.

As shown in the diagram, the first five packets matching the rule do not cause event notification because the rule does not trigger until the rate indicated in the `detection_filter` keyword is exceeded. After the rule triggers, event notification begins, but the rate-based criteria do not trigger the new action of Drop and Generate Events until five more packets pass. After the rate-based criteria are met, the system generates events for packets 11-15 and drops the packets. After the fifteenth packet, the limit threshold has been reached, so for the remaining packets the system does not generate events but does drop the packets.

After the rate-based timeout, note that packets are still dropped in the rate-based sampling period that follows. Because the sampled rate is above the threshold rate in the previous sampling period, the new action continues.



Rate-Based Attack Prevention Options and Configuration

Rate-based attack prevention identifies abnormal traffic patterns and attempts to minimize the impact of that traffic on legitimate requests. Rate-based attacks usually have one of the following characteristics:

- Any traffic containing excessive incomplete connections to hosts on the network, indicating a SYN flood attack
- Any traffic containing excessive complete connections to hosts on the network, indicating a TCP/IP connection flood attack
- Excessive rule matches in traffic going to a particular destination IP address or addresses or coming from a particular source IP address or addresses
- Excessive matches for a particular rule across all traffic

In a network analysis policy, you can either configure SYN flood or TCP/IP connection flood detection for the entire policy; in an intrusion policy, you can set rate-based filters for individual intrusion or preprocessor rules. Note that you cannot manually add a rate-based filter to GID 135 rules or modify their rule state. Rules with GID 135 use the client as the source value and the server as the destination value.

When **SYN Attack Prevention** is enabled, rule 135:1 triggers if a defined rate condition is exceeded.

When **Control Simultaneous Connections** is enabled, rule 135:2 tiggers if a defined rate condition is exceeded, and rule 135:3 triggeres if a session closes or times out.

Important The precedence of rate based preprocessors is as follows:

TCP SYN Rate based filtering > SI (IP reputation) > TCP Connection Rate based filtering

TCP SYN based rate filtering has the highest priority, but rate-based filtering depends on configuration like sample time and timeout. If there is a drop action, there is no further inspection.

Note Devices load-balance inspection across internal resources. When you configure rate-based attack prevention, you configure the triggering rate per resource, not per device. If rate-based attack prevention is not working as expected, you may need to lower the triggering rate. It triggers alert, if users send too many connection attempts within prescribed time intervals. Hence it is recommended to rate limit the rule. For help determining the correct rate, contact Support.

Each rate-based filter contains several components:

- For policy-wide or rule-based source or destination settings, the network address designation
- The rule matching rate, which you configure as a count of rule matches within a specific number of seconds
- A new action to be taken when the rate is exceeded

  When you set a rate-based setting for the entire policy, the system generates events when it detects a rate-based attack, and can drop the traffic in an inline deployment. When setting rate-based actions for individual rules, you have three available actions: Generate Events, Drop and Generate Events, and Disable.
- The duration of the action, which you configure as a timeout value

Note that when started, the new action occurs until the timeout is reached, even if the rate falls below the configured rate during that time period. When the timeout period expires, if the rate has fallen below the threshold, the action for the rule reverts to the action initially configured for the rule. For policy-wide settings, the action reverts to the action of each rule the traffic matches or stops if it does not match any rules.

You can configure rate-based attack prevention in an inline deployment to block attacks, either temporarily or permanently. Without rate-based configuration, rules set to Generate Events create events, but the system does not drop packets for those rules. However, if the attack traffic matches rules that have rate-based criteria configured, the rate action may cause packet dropping to occur for the period of time that the rate action is active, even if those rules are not initially set to Drop and Generate Events.

Note Rate-based actions cannot enable disabled rules or drop traffic that matches disabled rules. However, if you set a rate-based filter at the policy level, you can generate events on or generate events on and drop traffic that contains an excessive number of SYN packets or SYN/ACK interactions within a designated time period.

You can define multiple rate-based filters on the same rule. The first filter listed in the intrusion policy has the highest priority. Note that when two rate-based filter actions conflict, the system implements the action of the first rate-based filter. Similarly, policy-wide rate-based filters override rate-based filters set on individual rules if the filters conflict.

Related Topics

Setting a Dynamic Rule State from the Rules Page, on page 1512

Rate-Based Attack Prevention, Detection Filtering, and Thresholding or Suppression

The `detection_filter` keyword prevents a rule from triggering until a threshold number of rule matches occur within a specified time. When a rule includes the `detection_filter` keyword, the system tracks the number of incoming packets matching the pattern in the rule per timeout period. The system can count hits for that rule from particular source or destination IP addresses. After the rate exceeds the rate in the rule, event notification for that rule begins.

You can use thresholding and suppression to reduce excessive events by limiting the number of event notifications for a rule, a source, or destination, or by suppressing notifications altogether for that rule. You can also configure a global rule threshold that applies to each rule that does not have an overriding specific threshold.

If you apply suppression to a rule, the system suppresses event notifications for that rule for all applicable IP addresses even if a rate-based action change occurs because of a policy-wide or rule-specific rate-based setting.

Related Topics

Intrusion Event Thresholds, on page 1505
Intrusion Policy Suppression Configuration, on page 1509
Global Rule Thresholding Basics, on page 1537

Configuring Rate-Based Attack Prevention

You can configure rate-based attack prevention at the policy level to stop SYN flood attacks. You can also stop excessive connections from a specific source or to a specific destination.

Procedure

Step 1 Choose **Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**.

Note If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Click **Edit** ( ) next to the policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Settings**.

Step 4 If **Rate-Based Attack Prevention** under **Specific Threat Detection** is disabled, click **Enabled**.

Step 5 Click **Edit** ( ) next to **Rate-Based Attack Prevention**.

Step 6 You have two choices:

- To prevent incomplete connections intended to flood a host, click **Add** under **SYN Attack Prevention**.
- To prevent excessive numbers of connections, click **Add** under **Control Simultaneous Connections**.

Step 7 Specify how you want to track traffic:

- To track all traffic from a specific source or range of sources, choose **Source** from the **Track By** drop-down list, and enter a single IP address or address block in the **Network** field.
- To track all traffic to a specific destination or range of destinations, choose **Destination** from the **Track By** drop-down list, and enter an IP address or address block in the **Network** field.

Note

- Do not enter the IP address 0.0.0.0/0 in the Network field to monitor all subnets or IPs. The system does not support this IP address (which is usually used to identify all subnets or IPs) for Rate Based Attack Prevention.
- The system tracks traffic separately for each IP address included in the **Network** field. Traffic from an IP address that exceeds the configured rate results in generated events only for that IP address. As an example, you might set a source CIDR block of `10.1.0.0/16` for the network setting and configure the system to generate events when there are ten simultaneous connections open. If eight connections are open from 10.1.4.21 and six from 10.1.5.10, the system does not generate events, because neither source has the triggering number of connections open. However, if eleven simultaneous connections are open from 10.1.4.21, the system generates events only for the connections from 10.1.4.21.

  The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Step 8 Specify the triggering rate for the rate tracking setting:

- For SYN attack configuration, enter the number of SYN packets per number of seconds in the **Rate** fields.
- For simultaneous connection configuration, enter the number of connections in the **Count** field.

Devices load-balance inspection across internal resources. When you configure rate-based attack prevention, you configure the triggering rate per resource, not per device. If rate-based attack prevention is not working as expected, you may need to lower the triggering rate. It triggers alert, if users send too many connection attempts within prescribed time intervals. Hence it is recommended to rate limit the rule. For help determining the correct rate, contact Support.

Step 9 To drop packets matching the rate-based attack prevention settings, check the **Drop** check box.

Step 10 In the **Timeout** field, enter the time period after which to stop generating events (and if applicable, dropping) for traffic with the matching pattern of SYNs or simultaneous connections.

Caution Setting a high timeout value may entirely block connection to a host in an inline deployment.

Step 11 Click **OK**.

Step 12 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

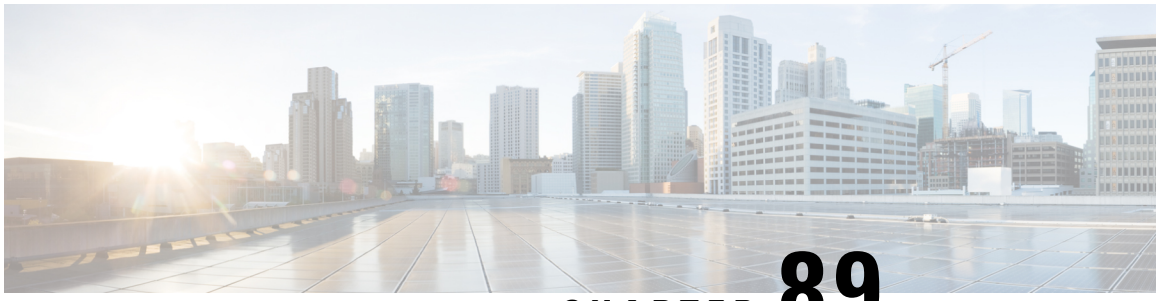
If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Firepower System IP Address Conventions, on page 15



CHAPTER 89

Adaptive Profiles

The following topics describe how to configure adaptive profiles:

- About Adaptive Profiles, on page 1811
- License Requirements for Adaptive Profiles, on page 1812
- Requirements and Prerequisites for Adaptive Profiles, on page 1812
- Adaptive Profile Updates, on page 1812
- Adaptive Profile Updates and Firepower Recommended Rules, on page 1813
- Adaptive Profile Options, on page 1813
- Configuring Adaptive Profiles, on page 1814

About Adaptive Profiles

Adaptive profiles must be enabled in order to:

- Perform application and file control, including malware protection (AMP), and to allow intrusion rules to use service metadata.

Caution Adaptive profiling **must** be enabled (its default state) as described in Configuring Adaptive Profiles, on page 1814 for access control rules to perform application and file control, including malware protection (AMP), and for intrusion rules to use service metadata.

- For passive deployments, enable adaptive profile updates to defragment and reassemble IP traffic according to the destination hosts' operating systems.

Note For inline deployments Cisco recommends that, instead of enabling adaptive profile updates, you configure the inline normalization preprocessor with the **Normalize TCP Payload** option enabled.

License Requirements for Adaptive Profiles

FTD License

Threat

Classic License

Protection

Requirements and Prerequisites for Adaptive Profiles

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Adaptive Profile Updates

Typically, the system uses the static settings in your network analysis policy to preprocess and analyze traffic. With adaptive profile updates, the system can adapt processing behavior using host information either detected by network discovery or imported from a third party.

Profile updates, like the target-based profiles you can configure manually in a network analysis policy, help to defragment IP packets and reassemble streams in the same way as the operating system on the target host. The intrusion rules engine then analyzes the data in the same format as that used by the destination host.

Manually configured target-based profiles apply either the default operating system profile you select, or profiles you bind to specific hosts. Profile updates, however, switch to the appropriate operating system profile based on the operating system in the host profile for the target host.

Consider a scenario where you configure profile updates for the 10.6.0.0/16 subnet and set the default IP Defragmentation target-based policy to Linux. The Firepower Management Center where you configure the settings has a network map that includes the 10.6.0.0/16 subnet.

- When the system detects traffic from Host A, which is not in the 10.6.0.0/16 subnet, it uses the Linux target-based policy to reassemble IP fragments.

- When the system detects traffic from Host B, which is in the 10.6.0.0/16 subnet, it retrieves Host B's operating system data from the network map. The system uses a profile based on that operating system to defragment the traffic destined for Host B.

Adaptive Profile Updates and Firepower Recommended Rules

The adaptive profile updates feature is an advanced setting in an access control policy that applies globally to all intrusion policies invoked by that access control policy. The Firepower recommended rules feature applies to the individual intrusion policy where you configure it.

Like Firepower recommended rules, profile updates compare metadata in a rule to host information to determine whether a rule should apply for a particular host. However, while Firepower recommended rules provide recommendations for enabling or disabling rules using that information, profile updates use the information to apply specific rules to specific traffic.

Firepower recommended rules require your interaction to implement suggested changes to rule states. Profile updates, on the other hand, do not modify intrusion policies. Treatment of rules based on profile updates happens on a packet-by-packet basis.

Additionally, Firepower recommended rules can result in enabling disabled rules. Profile updates, in contrast, only affect the application of rules that are already enabled in intrusion policies. Profile updates never change the rule state.

You can use profile updates and Firepower recommended rules in combination. Profile updates use the rule state for a rule when your intrusion policy is deployed to determine whether to include it as a candidate for applying, and your choices to accept or decline recommendations are reflected in that rule state. You can use both features to ensure that you have enabled or disabled the most appropriate rules for each network you monitor, and then to apply enabled rules most efficiently for specific traffic.

Related Topics

About Firepower Recommended Rules, on page 1517

Adaptive Profile Options

Enable

Enabling this option is required for:

- access control rules to perform application and file control, including malware protection (AMP)
- intrusion rules to use service metadata

This option is enabled by default.

Enable Profile Updates

In passive deployments, enable profile updates to defragment and reassemble IP traffic according to a profile of the operating system used by the hosts in your network map

Adaptive Profiles - Attribute Update Interval

When profile updates are enabled, you can control how frequently in minutes network map data is synced from the Firepower Management Center to its managed devices. The system uses the data to determine what profiles should be used when processing traffic. Increasing the value for this option can improve performance in a large network.

Adaptive Profiles - Networks

Optionally, when profile updates are enabled, you can improve performance by constraining profile updates to a comma-separated list of IP addresses, address blocks, and network variables. If you use a network variable, the system uses the variable's value in the variable set linked to the default intrusion policy for your access control policy. For example, you could enter: **`192.168.1.101, 192.168.4.0/24, $HOME_NET`**. IPv4 and IPv6 are supported.

The default value (**`0.0.0.0/0`**) applies adaptive profile updates to all networks.

Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. If you enable and enforce profile updates in an ancestor policy, Cisco recommends you keep the default network constraint of **`0.0.0.0/0`**, or use a network variable with a value of `any` . This setting applies profile updates to all monitored hosts in all subdomains.

Related Topics

Inspection of Packets That Pass Before Traffic Is Identified, on page 1670
Firepower System IP Address Conventions, on page 15
Variable Sets, on page 412

Configuring Adaptive Profiles

In a passive deployment, Cisco recommends that you configure adaptive profile updates. In an inline deployment, configure the inline normalization preprocessor with the **Normalize TCP Payload** option enabled.

Caution Adaptive profiling **must** be enabled (its default state) as described in this procedure for access control rules to perform application or file control, including malware protection (AMP), and for intrusion rules to use service metadata. Enabling or disabling adaptive profiles restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Before you begin

The access control policy must have a network discovery policy that is enabled to do host/service discovery, or host data must be imported from a third-party source.

Procedure

Step 1 In the access control policy editor, click **Advanced**, then click **Edit** () next to the Detection Enhancement Settings section.

If **View** () appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings.If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 Set adaptive profile options as described in Adaptive Profile Options, on page 1813.

Step 3 Click **OK**.

Step 4 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

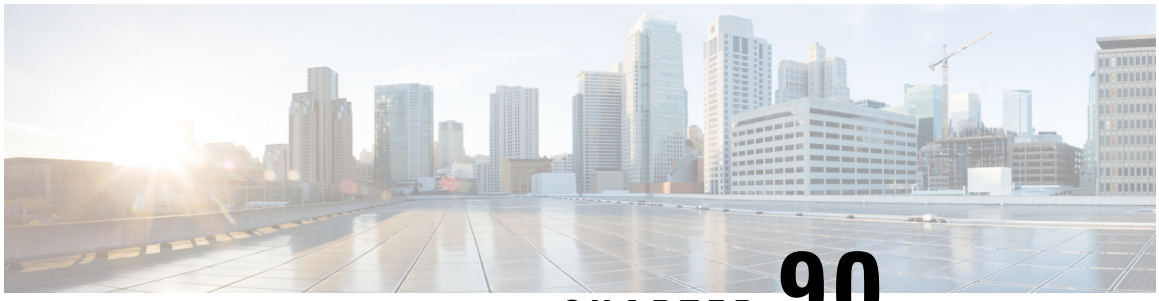
The Inline Normalization Preprocessor, on page 1762
Snort® Restart Scenarios, on page 348



PART XX

Discovery and Identity

- Introduction to Network Discovery and Identity, on page 1819
- Host Identity Sources, on page 1835
- Application Detection, on page 1873
- User Identity Sources, on page 1891
- Network Discovery Policies, on page 1917
- Realms and Identity Policies, on page 1941



CHAPTER 90

Introduction to Network Discovery and Identity

The following topics provide an introduction to network discovery and identity policies and data:

- Uses for Host, Application, and User Discovery and Identity Data, on page 1819
- Host and Application Detection Fundamentals, on page 1820
- About User Identity, on page 1827
- Firepower System Host and User Limits, on page 1831

Uses for Host, Application, and User Discovery and Identity Data

Logging discovery and identity data allows you to take advantage of many features in the Firepower System, including:

- Viewing the network map, which is a detailed representation of your network assets and topology that you can view by grouping hosts and network devices, host attributes, application protocols, or vulnerabilities.
- Performing application and user control; that is, writing access control rules using application, realm, user, user group, and ISE attribute conditions.
- Viewing host profiles, which are complete views of all the information available for your detected hosts.
- Viewing dashboards, which (among other capabilities) can provide you with an at-a-glance view of your network assets and user activity.
- Viewing detailed information on the discovery events and user activity logged by the system.
- Associating hosts and any servers or clients they are running with the exploits to which they are susceptible.

  This enables you to identify and mitigate vulnerabilities, evaluate the impact that intrusion events have on your network, and tune intrusion rule states so that they provide maximum protection for your network assets
- Alerting you by email, SNMP trap, or syslog when the system generates either an intrusion event with a specific impact flag, or a specific type of discovery event
- Monitoring your organization's compliance with a white list of allowed operating systems, clients, application protocols, and protocols

- Creating correlation policies with rules that trigger and generate correlation events when the system generates discovery events or detects user activity
- Logging and using NetFlow connections, if applicable.

Host and Application Detection Fundamentals

You can configure your network discovery policy to perform host and application detection.

For more information, see Overview: Host Data Collection, on page 1835 and Overview: Application Detection, on page 1873.

Passive Detection of Operating System and Host Data

Passive detection is the system's default method of populating the network map by analyzing network traffic (and any exported NetFlow data). Passive detection provides contextual information about your network assets, such as operating systems and running applications.

If traffic from a monitored host does not offer conclusive evidence of the host's operating system, the network map displays the most likely operating system. For example, a NAT device may appear to be running several operating systems because of the hosts "behind" the NAT device. To make this most-likely determination, the system uses a confidence value it assigns to each detected operating system, and the amount of corroborating data among detected operating systems.



Note The system does not consider reported "`unknown`" applications and operating systems in its determination.

If passive detection inaccurately identifies your network assets, consider the placement of your managed devices. You can also augment the system's passive detection capabilities with custom operating-system fingerprints and custom application detectors. Or, you can use *active detection*, which is not based on traffic analysis, but instead allows you to directly update the network map using scan results or other information sources.

Active Detection of Operating System and Host Data

Active detection adds host information collected by active sources to network maps. For example, you can use the Nmap scanner to actively scan the hosts that you target on your network. Nmap discovers operating systems and applications on hosts.

In addition, the host input feature allows you to actively add *host input data* to network maps. There are two different categories of host input data:

- *user input data*—Data added through the Firepower System user interface. You can modify a host's operating system or application identity through this interface.
- *host import input data*—Data imported using a command line utility.

The system retains one identity for each active source. When you run an Nmap scan instance, for example, the results of the previous scan are replaced with the new scan results. However, if you run an Nmap scan and then replace those results with data from a client whose results are imported through the command line,

the system retains both the identities from the Nmap results and the identities from the import client. The system then uses the priorities set in the network discovery policy to determine which active identity to use as the current identity.

Note that user input is considered one source, even if it comes from different users. As an example, if UserA sets the operating system through the host profile, and then UserB changes that definition through the host profile, the definition set by UserB is retained, and the definition set by UserA is discarded. In addition, note that user input overrides all other active sources and is used as the current identity if it exists.

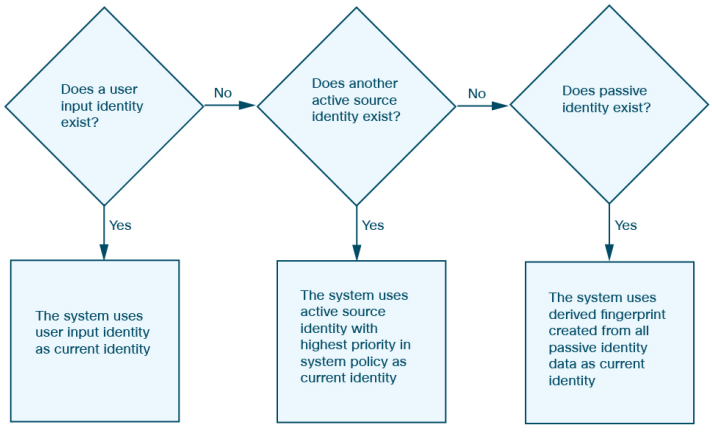
Current Identities for Applications and Operating Systems

The *current identity* for an application or an operating system on a host is the identity that the system finds most likely to be correct.

The system uses the current identity for an operating system or application for the following purposes:

- to assign vulnerabilities to a host
- for impact assessment
- when evaluating correlation rules written against operating system identifications, host profile qualifications, and compliance white lists
- for display in the Hosts and Servers table views in workflows
- for display in the host profile
- to calculate the operating system and application statistics on the Discovery Statistics page

The system uses source priorities to determine which active identity should be used as the current identity for an application or operating system.



For example, if a user sets the operating system to Windows 2003 Server on a host, Windows 2003 Server is the current identity. Attacks which target Windows 2003 Server vulnerabilities on that host are given a higher impact, and the vulnerabilities listed for that host in the host profile include Windows 2003 Server vulnerabilities.

The database may retain information from several sources for the operating system or for a particular application on a host.

The system treats an operating system or application identity as the current identity when the source for the data has the highest source priority. Possible sources have the following priority order:

1. user
2. scanner and application (set in the network discovery policy)
3. managed devices
4. NetFlow records

A new higher priority application identity will not override a current application identity if it has less detail than the current identity.

In addition, when an identity conflict occurs, the resolution of the conflict depends on settings in the network discovery policy or on your manual resolution.

Current User Identities

When the system detects multiple logins to the same host by different users, the system assumes that only one user is logged into any given host at a time, and that the current user of a host is the last authoritative user login. If only non-authoritative user logins have been logged into the host, the last non-authoritative user login is considered the current user. If multiple users are logged in through remote sessions, the last user reported by the server is the user reported to the Firepower Management Center.

When the system detects multiple logins to the same host by the same user, the system records the first time that a user logs into a specific host and disregards subsequent logins. If an individual user is the only person who logs into a specific host, the only login that the system records is the original login.

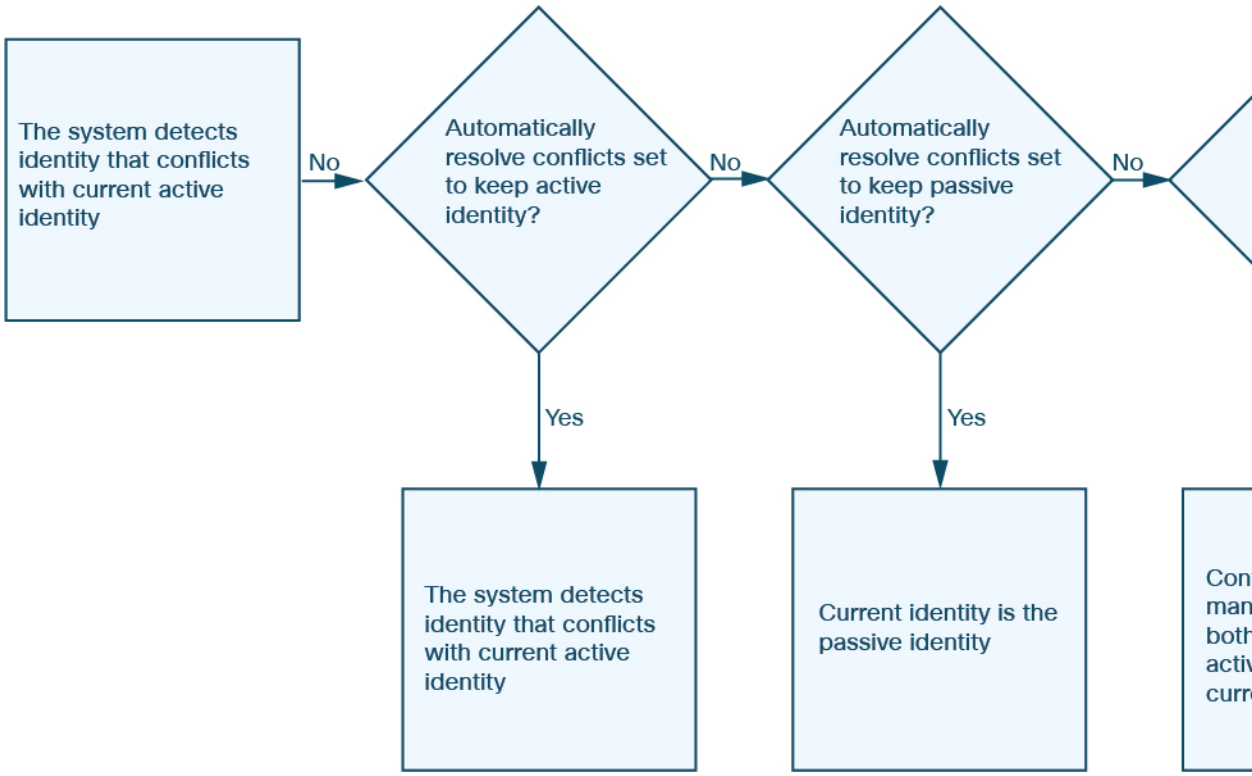
If another user logs into that host, however, the system records the new login. Then, if the original user logs in again, his or her new login is recorded.

Application and Operating System Identity Conflicts

An *identity conflict* occurs when the system reports a new passive identity that conflicts with the current active identity and previously reported passive identities. For example, the previous passive identity for an operating system is reported as Windows 2000, then an active identity of Windows XP becomes current. Next, the system detects a new passive identity of Ubuntu Linux 8.04.1. The Windows XP and the Ubuntu Linux identities are in conflict.

When an identity conflict exists for the identity of the host's operating system or one of the applications on the host, the system lists both conflicting identities as current and uses both for impact assessment until the conflict is resolved.

A user with Administrator privileges can resolve identity conflicts automatically by choosing to always use the passive identity or always use the active identity. Unless you disable automatic resolution of identity conflicts, identity conflicts are always automatically resolved.



A user with Administrator privileges can also configure the system to generate an event when an identity conflict occurs. That user can then set up a correlation policy with a correlation rule that uses an Nmap scan as a correlation response. When an event occurs, Nmap scans the host to obtain updated host operating system and application data.

Netflow Data in the Firepower System

NetFlow is a Cisco IOS application that provides statistics on packets flowing through a router. It is available on Cisco networking devices and can also be embedded in Juniper, FreeBSD, and OpenBSD devices.

When NetFlow is enabled on a network device, a database on the device (the NetFlow cache) stores records of the flows that pass through the router. A flow, called a *connection* in the Firepower System, is a sequence of packets that represents a session between a source and destination host, using specific ports, protocol, and application protocol. The network device can be configured to export this NetFlow data. In this documentation, network devices configured in this way are called *NetFlow exporters*.

Firepower System managed devices can be configured to collect records from NetFlow exporters, generate unidirectional end-of-connection events based on the data in those records, and finally send those events to the Firepower Management Center to be logged in the connection event database. You can also configure the network discovery policy to add host and application protocol information to the database based on the information in NetFlow connections.

You can use this discovery and connection data to supplement the data gathered directly by your managed devices. This is especially useful if you have NetFlow exporters monitoring networks that your managed devices cannot monitor.

Requirements for Using NetFlow Data

Before you configure the Firepower System to analyze NetFlow data, you must enable the NetFlow feature on the routers or other NetFlow-enabled network devices you plan to use, and configure the devices to broadcast NetFlow data to a destination network where the sensing interface of a managed device is connected.

The Firepower System can parse both NetFlow version 5 and NetFlow version 9 records. NetFlow exporters **must** use one of those versions if you want to export the data to the Firepower System. In addition, the system requires that specific fields be present in the exported NetFlow templates and records. If your NetFlow exporters are using version 9, which you can customize, you **must** make sure that the exported templates and records contain the following fields, in any order:

- IN_BYTES (1)
- IN_PKTS (2)
- PROTOCOL (4)
- TCP_FLAGS (6)
- L4_SRC_PORT (7)
- IPV4_SRC_ADDR (8)
- L4_DST_PORT (11)
- IPV4_DST_ADDR (12)
- LAST_SWITCHED (21)
- FIRST_SWITCHED (22)
- IPV6_SRC_ADDR (27)
- IPV6_DST_ADDR (28)

Because the Firepower System uses managed devices to analyze NetFlow data, your deployment must include at least one managed device that can monitor NetFlow exporters. At least one sensing interface on that managed device must be connected to a network where it can collect the exported NetFlow data. Because the sensing interfaces on managed devices do not usually have IP addresses, the system does not support the direct collection of NetFlow records.

Note that the Sampled NetFlow feature available on some network devices collects NetFlow statistics on only a subset of packets that pass through the devices. Although enabling this feature can improve CPU utilization on the network device, it may affect the NetFlow data you are collecting for analysis by the Firepower System.

Differences between NetFlow and Managed Device Data

The traffic represented by NetFlow data is not directly analysed. Instead, the exported NetFlow records are converted into connection logs and host and application protocol data.

As a result, there are several differences between converted NetFlow data and the discovery and connection data gathered directly by your managed devices. You should keep these differences in mind when performing analysis that requires:

- Statistics on the number of detected connections
- Operating system and other host-related information (including vulnerabilities)

- Application data, including client information, web application information, and vendor and version server information
- Knowing which host in a connection is the initiator and which is the responder

Network Discovery Policy versus Access Control Policy

You configure NetFlow data collection, including connection logging, using rules in the network discovery policy. Contrast this with connection logging for connections detected by managed devices, which you configure per access control rule.

Types of Connection Events

Because NetFlow data collection is linked to networks rather than access control rules, you do not have granular control over which NetFlow connections the system logs.

NetFlow data cannot generate Security Intelligence events.

NetFlow-based connection events can be stored in the connection event database only; you cannot send them to the system log or an SNMP trap server.

Number of Connection Events Generated Per Monitored Session

For connections detected directly by managed devices, you can configure the access control rule to log a bidirectional connection event at the beginning or end of a connection, or both.

In contrast, because exported NetFlow records contain unidirectional connection data, the system generates at least two connection events for each NetFlow record it processes. This also means that a summary's connection count is incremented by two for every connection based on NetFlow data, providing an inflated count of the number of connections that are actually occurring on your network.

Because the NetFlow exporter outputs records at a fixed interval even if a connection is still ongoing, long-running sessions can result in multiple exported records, each of which generates a connection event. For example, if the NetFlow exporter exports every five minutes, and a particular connection lasts twelve minutes, the system generates six connection events for that session:

- One pair of events for the first five minutes
- One pair for the second five minutes
- A final pair when the connection is terminated

Host and Operating System Data

Hosts added to the network map from NetFlow data do not have operating system, NetBIOS, or host type (host vs network device) information. You can, however, manually set a host’s operating system identity using the host input feature.

Application Data

For connections detected directly by managed devices, the system can identify application protocols, clients, and web applications by examining the packets in the connection.

When the NetFlow records are processesed, the system uses a port correlation in `/etc/sf/services` to extrapolate application protocol identity. However, there is no vendor or version information for those

application protocols, nor do connection logs contain information on client or web applications used in the session. You can, however, manually provide this information using the host input feature.

Note that a simple port correlation means that application protocols running on non-standard ports may be unidentified or misidentified. Additionally, if no correlation exists, the system marks the application protocol as `unknown` in connection logs.

Vulnerability Mappings

The system cannot map vulnerabilities to hosts monitored by NetFlow exporters, unless you use the host input feature to manually set either a host's operating system identity or an application protocol identity. Note that because there is no client information in NetFlow connections, you cannot associate client vulnerabilities with hosts created from NetFlow data.

Initiator and Responder Information in Connections

For connections detected directly by managed devices, the system can identify which host is the initiator, or source, and which is the responder, or destination. However, NetFlow data does not contain initiator or responder information.

When the system processes NetFlow records, it uses an algorithm to determine this information based on the ports each host is using, and whether those ports are well-known:

- If both or neither port being used is a well-known port, the system considers the host using the lower-number port to be the responder.
- If only one of the hosts is using a well-known port, the system considers that host to be the responder.

For this purpose, a well-known port is any port that is either numbered from 1 to 1023, or that contains application protocol information in `/etc/sf/services` on the managed device.

In addition, for connections detected directly by managed devices, the system records two byte counts in the corresponding connection event:

- The **Initiator Bytes** field records bytes sent.
- The **Responder Bytes** field records bytes received.

Connection events based on unidirectional NetFlow records contain only one byte count, which the system assigns to either **Initiator Bytes** or **Responder Bytes**, depending on the port-based algorithm. The system sets the other field to 0. Note that if you are viewing connection summaries (aggregated connection data) of NetFlow records, both fields may be populated.

NetFlow-only Connection Event Fields

A small number of fields are present only in connection events generated from NetFlow records; see Information Available in Connection Event Fields, on page 2242.

Related Topics

Information Available in Connection Event Fields, on page 2242

About User Identity

User identity information can help you to identify the source of policy breaches, attacks, or network vulnerabilities, and trace them to specific users. For example, you could determine:

- Who owns the host targeted by an intrusion event that has a Vulnerable (level 1: red) impact level.
- Who initiated an internal attack or portscan.
- Who is attempting unauthorized access to a specified host.
- Who is consuming an unreasonable amount of bandwidth.
- Who has not applied critical operating system updates.
- Who is using instant messaging software or peer-to-peer file-sharing applications in violation of company policy.
- Who is associated with each indication of compromise on your network.

Armed with this information, you can use other features of the Firepower System to mitigate risk, perform access control, and take action to protect others from disruption. These capabilities also significantly improve audit controls and enhance regulatory compliance.

After you configure user identity sources to gather user data, you can perform user awareness and user control.

Video YouTube videos for configuring identity.

Related Topics

Identity Terminology, on page 1827
Identity Deployments, on page 1830

Identity Terminology

This topic discusses common terminology for user identity and user control.

User awareness

Identifying users on your network using *identity sources* (such as user agent). User awareness enables you to identify users from both *authoritative* (such as Active Directory) and *non-authoritative* (application-based) sources. To use Active Directory as an identity source, you must configure a realm and directory. For more information, see About User Identity Sources, on page 1891.

User control

Configuring an *identity policy* that you associate with an *access control policy*. (The identity policy is then referred to as an access control *subpolicy*.) The identity policy specifies the identity source and, optionally, users and groups belonging to that source.

By associating the identity policy with an access control policy, you determine whether to monitor, trust, block, or allow users or user activity in traffic on your network. For more information, see Access Control Policies, on page 1201.

Authoritative identity sources

A trusted server validated the user login (for example, Active Directory). You can use the data obtained from authoritative logins to perform user awareness and user control. Authoritative user logins are obtained from passive and active authentications:

- *Passive authentications* occur when a user authenticates through an external server. ISE and the user agent, are the passive authentication methods supported by the Firepower System.
- *Active authentications* occur when a user authenticates through preconfigured managed devices. Captive portal is the only active authentication method supported by the Firepower System.

Non-authoritative identity sources

An unknown or untrusted server validated the user login. Traffic-based detection is the only non-authoritative identity source supported by the Firepower System. You can use the data obtained from non-authoritative logins to perform user awareness.

Best Practices for User Identity

We recommend you review the following information before you set up identity policies.

- Know user limits
- Create one realm per AD domain, something about trust
- Health monitor
- Use latest version of ISE/ISE-PIC, two types of remediation
- User agent support drops in 6.7
- Captive portal requires routed interface, several individual tasks
- See TS Agent troubleshooting

Active Directory, LDAP, and realms

The Firepower System supports either Active Directory or LDAP for user awareness and control. The association between an Active Directory or LDAP repository and the FMC is referred to as a *realm*. You should create one realm per LDAP server or Active Directory domain. For details about which versions are supported, see Supported Servers for Realms, on page 1943.

The only user identity source supported by LDAP is captive portal. To use other identity sources (with the exception of ISE), you must use Active Directory.

For Active Directory only:

- Create one *directory* per domain controller.

  For details, see Configure a Realm Directory, on page 1959.

Health monitor

The FMC health monitor provides valuable information about the status of various FMC functions, including:

- User/realm mismatches

- Snort memory usage
- ISE connection status

For more information about health modules, see Health Modules, on page 277.

To set up policies to monitor health modules, see Creating Health Policies, on page 283.

Device-specific user limits

Every physical or virtual FMC device has limits to the number of users that can be downloaded. When the user limit is reached, the FMC can run out of memory and can function unreliably as a result.

User limits are discussed in Firepower System User Limit, on page 1833.

If you use the ISE/ISE-PIC identity source, you can optionally limit the subnets the FMC monitors to reduce memory usage using identity mapping filters as discussed in Create an Identity Policy, on page 1961.

Use the latest version of ISE

If you expect to use the ISE/ISE-PIC identity source, we strongly recommend you always use the latest version to make sure you get the latest features and bug fixes.

pxGrid 2.0 (which is used by version 2.6 patch 6 or later; or 2.7 patch 2 or later) also changes the remediation used by ISE from Endpoint Protection Service (EPS) to Adaptive Network Control (ANC). If you upgrade ISE, you must migrate your mediation policies from EPS to ANC.

More information about using ISE can be found in ISE Guidelines and Limitations, on page 1896.

To set up the ISE identity source, see How to Configure ISE for User Control, on page 1895.

Captive portal information

Captive portal is the only user identity source for which you can use either LDAP or Active Directory. In addition, your managed devices must be configured to use a routed interface.

Additional guidelines can be found in Captive Portal Guidelines and Limitations, on page 1903.

Setting up captive portal requires performing several independent tasks. For more information, see How to Configure the Captive Portal for User Control, on page 1905.

TS Agent information

The TS Agent user identity source is required to identify user sessions on a Windows Terminal Server. The TS Agent software must be installed on the Terminal Server machine as discussed in the *Cisco Terminal Services (TS) Agent Guide*. In addition, you must synchronize the time on your TS Agent server with the time on the Firepower Management Center.

TS Agent data is visible in the Users, User Activity, and Connection Event tables and can be used for user awareness and user control.

Associate the identity policy with an access control policy

After you configure your realm, directory, and user identity source, you must set up identity rules in an identity policy. To make the policy effective, you must associate the identity policy with an access control policy.

For more information about creating an identity policy, see Create an Identity Policy, on page 1961.

For more information about creating identity rules, see Create an Identity Rule, on page 1962.

To associate an identity policy with an access control policy, see Associating Other Policies with Access Control, on page 1213.

User agent deprecation and end of support by FMC

End of support is planned for FMC integration with the Cisco Firepower User Agent (hereafter referred to as *user agent*) in a future release.

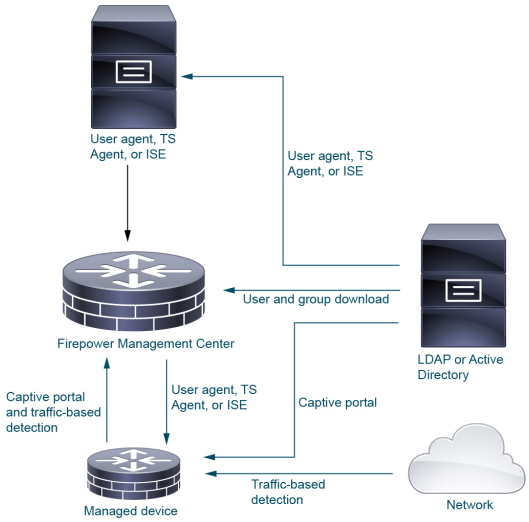
For more information, see End-of-Life and End-of-Support for the Cisco Firepower User Agent.

Identity Deployments

When the system detects user data from a user login, from any identity source, the user from the login is checked against the list of users in the Firepower Management Center user database. If the login user matches an existing user, the data from the login is assigned to the user. Logins that do not match existing users cause a new user to be created, unless the login is in SMTP traffic. Non-matching logins in SMTP traffic are discarded.

The group to which the user belongs is associated with the user when the user passes traffic on the network.

The following diagram illustrates how the Firepower System collects and stores user data:



The User Activity Database

The user activity database on the Firepower Management Center contains records of user activity on your network detected or reported by all of your configured identity sources. The system logs events in the following circumstances:

- When it detects individual logins or logoffs.
- When it detects a new user.
- When a system administrator manually delete a user.

- When the system detects a user that is not in the database, but cannot add the user because you have reached your user limit.
- When you resolve an indication of compromise associated with a user, or enable or disable indication of compromise rules for a user.

You can view user activity detected by the system using the Firepower Management Center web interface. (**Analysis** > **Users** > **User Activity**).

The Users Database

The users database on the Firepower Management Center contains a record for each user detected or reported by all of your configured identity sources. You can use data obtained from an authoritative source for user control.

See About User Identity Sources, on page 1891 for more information about the supported non-authoritative and authoritative identity sources.

The total number of users the Firepower Management Center can store depends on the Firepower Management Center model, as described in Firepower System User Limit, on page 1833. After the user limit is reached, the system prioritizes previously-undetected user data based on its identity source, as follows:

- If the new user is from a non-authoritative identity source, the system does not add the user to the database. To allow new users to be added, you must delete users manually or with a database purge.
- If the new user is from an authoritative identity source, the system deletes the non-authoritative user who has remained inactive for the longest period and adds the new user to the database.

If an identity source is configured to exclude specific user names, user activity data for those user names are not reported to the Firepower Management Center. These excluded user names remain in the database, but are not associated with IP addresses. For more information about the type of data stored by the system, see User Data, on page 2395.

If you have Firepower Management Center high availability configured and the primary fails, no logins reported by a captive portal, user agent or ISE device can be identified during failover downtime, even if the users were previously seen and downloaded to the Firepower Management Center. The unidentified users are logged as Unknown users on the Firepower Management Center. After the downtime, the Unknown users are reidentified and processed according to the rules in your identity policy.

When the system detects a new user session, the user session data remains in the users database until one of the following occurs:

- A user on the Firepower Management Center manually deletes the user session.
- An identity source reports the logoff of that user session.
- A realm ends the user session as specified by the realm's **User Session Timeout: Authenticated Users**, **User Session Timeout: Failed Authentication Users**, or **User Session Timeout: Guest Users** setting.

Firepower System Host and User Limits

Your Firepower Management Center model determines how many individual hosts you can monitor with your deployment, as well as how many users you can monitor and use to perform user control.

Related Topics

Purging Data from the FMC Database, on page 198

Firepower System Host Limit

The system adds a host to the network map when it detects activity associated with an IP address in your monitored network (as defined in your network discovery policy). The number of hosts a Firepower Management Center can monitor, and therefore store in the network map, depends on its model.

Table 216: Host Limits by Firepower Management Center Model

FMC Model	Hosts
MC750	2,000
MC1000	50,000
MC1500	50,000
MC2000	150,000
MC2500	150,000
MC3500	300,000
MC4000	600,000
MC4500	600,000
virtual	50,000

You cannot view contextual data for hosts not in the network map. However, you can perform access control. For example, you can perform application control on traffic to and from a host not in the network map, even though you cannot use a compliance white list to monitor the host's network compliance.

Note The system counts MAC-only hosts separately from hosts identified by both IP addresses and MAC addresses. All IP addresses associated with a host are counted together as one host.

Reaching the Host Limit and Deleting Hosts

The network discovery policy controls what happens when you detect a new host after you reach the host limit; you can drop the new host, or replace the host that has been inactive for the longest time. You can also set the period after which the system removes a host from the network map due to inactivity. Although you can manually delete a host, an entire subnet, or all of your hosts from the network map, if the system detects activity associated with a deleted host, it re-adds the host.

In a multidomain deployment, each leaf domain has its own network discovery policy. Therefore, each leaf domain governs its own behavior when the system discovers a new host.

Related Topics

Domain Properties, on page 337

Network Discovery Data Storage Settings, on page 1934

Firepower System User Limit

Your Firepower Management Center model determines how many individual users you can monitor. The user is added to the Firepower Management Center user database when:

- The user is downloaded from a realm.
- A captive portal or RA-VPN user logs in.
- A user is detected from any identity source (for example, TS Agent).

Only authoritative users are available for user control with access control policies.

Note the following:

The *concurrent user limit* is the number of users logged in to the system at the same time. These users are all *authoritative users*, which means they were reported to the Firepower Management Center by an authoritative user source: User Agent, ISE, the TS Agent, and captive portal.

Table 217: Maximum Downloaded Users by Firepower Management Center Model[1]

FMC Model	Maximum Downloaded Users
FMC750	2,000
FMC1000	50,000
FMC1500	50,000
FMC2000	150,000
FMC2500	150,000
FMC3500	300,000
FMC4000	600,000
FMC4500	600,000
FMCv (any supported hypervisor)	50,000
FMCv 300 (any supported hypervisor)	150,000

[1]—FMC models are subject to end of life and end of sale. For more information, see End-Of-Life and End-Of-Sale Notices.

When the system detects a new, previously-undetected user after the limit has been reached, it prioritizes user data based on their identity source:

- If the new user is from a non-authoritative source, the system does not add the non-authoritative user to the database. To allow new users to be added, you must delete users manually or purge the database.
- If the new user is from an authoritative identity source, the system deletes the non-authoritative user who has remained inactive for the longest period of and adds the new authoritative user to the database.

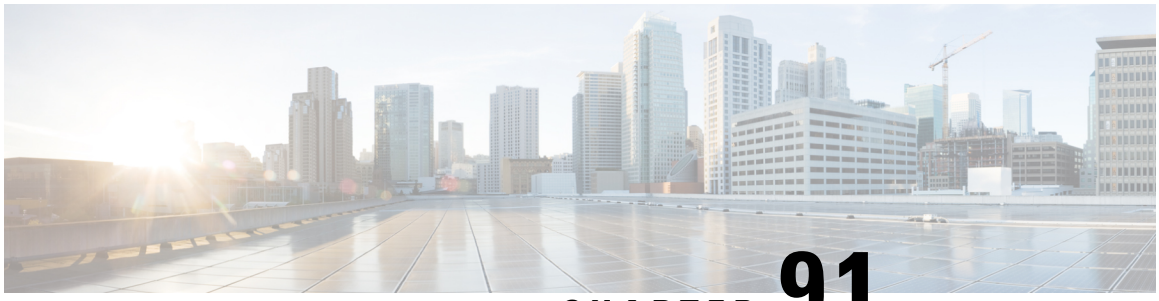
Troubleshooting information can be found in Troubleshoot User Control, on page 387.



Note If your deployment includes an ASA FirePOWER module managed via ASDM, you can store a maximum of 2,000 authoritative users, regardless of your Firepower Management Center model.



Tip Note that if you are using traffic-based detection, you can restrict user logging by protocol to help minimize username clutter and preserve space in the database. For example, you could prevent the system from adding users discovered in AIM, POP3, and IMAP traffic because you know it is traffic from specific contractors or visitors you do not want to monitor.



CHAPTER 91

Host Identity Sources

The following topics provide information on host identity sources:

- Overview: Host Data Collection, on page 1835
- Requirements and Prerequisites for Host Identity Sources, on page 1836
- Determining Which Host Operating Systems the System Can Detect, on page 1836
- Identifying Host Operating Systems, on page 1836
- Custom Fingerprinting, on page 1837
- Host Input Data, on page 1845
- Nmap Scanning, on page 1855

Overview: Host Data Collection

As the Firepower System passively monitors the traffic that travels through your network, it compares specific packet header values and other unique data from network traffic against established definitions (called *fingerprints*) to determine information about the hosts on your network, including:

- the number and types of hosts (including network devices such as bridges, routers, load balancers, and NAT devices)
- basic network topology data, including the number of hops from the discovery point on the network to the hosts
- the operating systems running on the hosts
- applications on the hosts and users associated with these applications

If the system cannot identify a host's operating system, you can create custom client or server fingerprints. The system uses these fingerprints to identify new hosts. You can map fingerprints to systems in the vulnerability database (VDB) to allow the appropriate vulnerability information to be displayed whenever a host is identified using the custom fingerprint.

Note In addition to collecting host data from monitored network traffic, the system can collect host data from exported NetFlow records, and you can actively add host data using Nmap scans and the host input feature.

Requirements and Prerequisites for Host Identity Sources

Model Support

Any.

Supported Domains

Any, with the exception of custom fingerprinting, which is Leaf only.

User Roles

- Admin
- Discovery Admin, except for third-party data and custom mappings.

Determining Which Host Operating Systems the System Can Detect

To learn which exact operating systems the system can fingerprint, view the list of available fingerprints that is shown during the process of creating a custom OS fingerprint.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Custom Operating Systems**.

Step 3 Click **Create Custom Fingerprint**.

Step 4 View the lists of options in the drop-down lists in the **OS Vulnerability Mappings** section. These options are the operating systems that the system can fingerprint.

What to do next

As needed, see Identifying Host Operating Systems, on page 1836.

Identifying Host Operating Systems

If the system does not correctly identify a host's operating system (for example, it shows in the Host Profile as Unknown or is incorrectly identified), try the strategies below.

Procedure

Try one of the following strategies:

- Check the Network Discovery Identity Conflict Settings.
- Create a custom fingerprint for the host.
- Run an Nmap scan against the host.
- Import data into the network map, using the host input feature.
- Manually enter operating system information.

Custom Fingerprinting

The Firepower System includes operating system *fingerprints* that the system uses to identify the operating system on each host it detects. However, sometimes the system cannot identify a host operating system or misidentifies it because no fingerprints exist that match the operating system. To correct this problem, you can create a *custom fingerprint,* which provides a pattern of operating system characteristics unique to the unknown or misidentified operating system, to supply the name of the operating system for identification purposes.

If the system cannot match a host's operating system, it cannot identify the vulnerabilities for the host, because the system derives the list of vulnerabilities for each host from its operating system fingerprint. For example, if the system detects a host running Microsoft Windows, the system has a stored Microsoft Windows vulnerability list that it adds to the host profile for that host based on the detected Windows operating system.

As an example, if you have several devices on your network running a new beta version of Microsoft Windows, the system cannot identify that operating system or map vulnerabilities to the hosts. However, knowing that the system has a list of vulnerabilities for Microsoft Windows, you may want to create a custom fingerprint for one of the hosts to help identify the other hosts running the same operating system. You can include a mapping of the vulnerability list for Microsoft Windows in the fingerprint to associate that list with each host that matches the fingerprint.

When you create a custom fingerprint, the Firepower Management Center lists the set of vulnerabilities associated with that fingerprint for any hosts running the same operating system. If the custom fingerprint you create does not have any vulnerabilities mappings in it, the system uses the fingerprint to assign the custom operating system information you provide in the fingerprint. When the system sees new traffic from a previously detected host, the system updates the host with the new fingerprint information. the system also uses the new fingerprint to identify any new hosts with that operating system the first time they are detected.

Before creating a custom fingerprint, you should determine why the host is not being identified correctly to decide whether custom fingerprinting is a viable solution.

You can create two types of fingerprints with the system:

- Client fingerprints, which identify operating systems based on the SYN packet that the host sends when it connects to a TCP application running on another host on the network.
- Server fingerprints, which identify operating systems based on the SYN-ACK packet that the host uses to respond to an incoming connection to a running TCP application.

Note If both a client and server fingerprint match the same host, the client fingerprint is used.

After creating fingerprints, you must activate them before the system can associate them with hosts.

Related Topics

Creating a Custom Fingerprint for Clients, on page 1840
Creating a Custom Fingerprint for Servers, on page 1842

Managing Fingerprints

After a fingerprint is created and activated, you can edit a fingerprint to make changes or add vulnerability mappings.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Custom Operating Systems**. If the system is awaiting data to create a fingerprint, it automatically refreshes the page every 10 seconds until the fingerprint is created.

Step 3 Manage your custom fingerprints:

- Activate/Deactivate — Activate or deactivate a fingerprint as described in Activating and Deactivating Fingerprints, on page 1838.
- Create — Create fingerprints as described in Creating a Custom Fingerprint for Clients, on page 1840 and Creating a Custom Fingerprint for Servers, on page 1842.
- Edit — Edit a fingerprint as described in Editing an Active Fingerprint, on page 1839 and Editing an Inactive Fingerprint, on page 1839.
- Delete — Click **Delete** ( ) next to the fingerprint you want to delete, and click **OK** to confirm. You can only delete deactivated fingerprints.

Activating and Deactivating Fingerprints

You must activate a custom fingerprint before the system can use it to identify hosts. After the new fingerprint is activated, the system uses it to re-identify previously discovered hosts and discover new hosts.

If you want to stop using a fingerprint, you can deactivate it. Deactivating a fingerprint causes a fingerprint to no longer be used, but allows it to remain on the system. When you deactivate a fingerprint, the operating system is marked as unknown for hosts that use the fingerprint. If the hosts are detected again and match a different active fingerprint, they are then identified by that active fingerprint.

Deleting a fingerprint removes it from the system completely. After deactivating a fingerprint, you can delete it.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Custom Operating Systems**.

Step 3 Click the slider next to the fingerprint you want to activate or deactivate.

Note The activate option is only available if the fingerprint you created is valid. If the slider is not available, try creating the fingerprint again.

Editing an Active Fingerprint

If a fingerprint is *active*, you can modify the fingerprint name, description, custom operating system display, and map additional vulnerabilities to it.

You can modify the fingerprint name, description, custom operating system display, and map additional vulnerabilities to it.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Custom Operating Systems**

Step 3 Click **Edit** () next to the fingerprint you want to edit.

Step 4 Modify the fingerprint name, description, and custom OS display, if necessary.

Step 5 If you want to delete a vulnerability mapping, click **Delete** next to the mapping in the **Pre-Defined OS Product Maps** section of the page.

Step 6 If you want to add additional operating systems for vulnerability mapping, choose the **Product** and, if applicable, **Major Version**, **Minor Version**, **Revision Version**, **Build**, **Patch**, and **Extension** and then click **Add OS Definition**.

The vulnerability mapping is added to the **Pre-Defined OS Product Maps** list.

Step 7 Click **Save**.

Editing an Inactive Fingerprint

If a fingerprint is *inactive*, you can modify all elements of the fingerprint and resubmit it to the Firepower Management Center. This includes all properties you specified when creating the fingerprint, such as fingerprint type, target IP addresses and ports, vulnerability mappings, and so on. When you edit an inactive fingerprint and submit it, it is resubmitted to the system and, if it is a client fingerprint, you must resend traffic to the appliance before activating it. Note that you can choose only a single vulnerability mapping for an inactive

fingerprint. After you activate the fingerprint, you can map additional operating systems and versions to its vulnerabilities list.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Custom Operating Systems**.

Step 3 Click **Edit** ( ) next to the fingerprint you want to edit.

Step 4 Make changes to the fingerprint as necessary:

- If you are modifying a client fingerprint, see Creating a Custom Fingerprint for Clients, on page 1840.
- If you are modifying a server fingerprint, see Creating a Custom Fingerprint for Servers, on page 1842.

Step 5 Click **Save**.

What to do next

- If you modified a client fingerprint, remember to send traffic from the host to the appliance gathering the fingerprint.

Creating a Custom Fingerprint for Clients

Client fingerprints identify operating systems based on the SYN packet a host sends when it connects to a TCP application running on another host on the network.

If the Firepower Management Center does not have direct contact with monitored hosts, you can specify a device that is managed by the FMC and is closest to the host you intend to fingerprint when specifying client fingerprint properties.

Before you begin the fingerprinting process, obtain the following information about the host you want to fingerprint:

- The number of network hops between the host and the Firepower Management Center or the device you use to obtain the fingerprint. (Cisco strongly recommends that you directly connect the Firepower Management Center or the device to the same subnet that the host is connected to.)
- The network interface (on the Firepower Management Center or the device) that is connected to the network where the host resides.
- The actual operating system vendor, product, and version of the host.
- Access to the host in order to generate client traffic.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Custom Operating Systems**.

Step 3 Click **Create Custom Fingerprint**.

Step 4 From the **Device** drop-down list, choose the Firepower Management Center or the device that you want to use to collect the fingerprint.

Step 5 Enter a **Fingerprint Name**.

Step 6 Enter a **Fingerprint Description**.

Step 7 From the **Fingerprint Type** list, choose **Client**.

Step 8 In the **Target IP Address** field, enter an IP address of the host you want to fingerprint.

Note that the fingerprint will only be based on traffic to and from the host IP address you specify, not any of the host's other IP addresses (if it has any).

Step 9 In the **Target Distance** field, enter the number of network hops between the host and the device that you chose earlier to collect the fingerprint.

Caution This must be the actual number of physical network hops to the host, which may or may not be the same as the number of hops detected by the system.

Step 10 From the **Interface** list, choose the network interface that is connected to the network segment where the host resides.

Caution Cisco recommends that you do **not** use the sensing interface on a managed device for fingerprinting for several reasons. First, fingerprinting does not work if the sensing interface is on a span port. Also, if you use the sensing interface on a device, the device stops monitoring the network for the amount of time it takes to collect the fingerprint. You can, however, use the management interface or any other available network interfaces to perform fingerprint collection. If you do not know which interface is the sensing interface on your device, refer to the *Installation Guide* for the specific model you are using to fingerprint.

Step 11 If you want to display custom information in the host profile for fingerprinted hosts (or if the host you want to fingerprint does not reside in the **OS Vulnerability Mappings** section), choose **Use Custom OS Display** and provide the values you want to display for the following:

- In the **Vendor String** field, enter the operating system's vendor name. For example, the vendor for Microsoft Windows would be Microsoft.
- In the **Product String** field, enter the operating system's product name. For example, the product name for Microsoft Windows 2000 would be Windows.
- In the **Version String** field, enter the operating system's version number. For example, the version number for Microsoft Windows 2000 would be 2000.

Step 12 In the OS Vulnerability Mappings section, choose the operating system, product, and versions you want to use for vulnerability mapping.

You must specify **Vendor** and **Product** values in this section if you want to use the fingerprint to identify vulnerabilities for matching hosts or if you do not assign custom operating system display information.

To map vulnerabilities for all versions of an operating system, specify only the **Vendor** and **Product** values.

Note Not all options in the **Major Version**, **Minor Version**, **Revision Version**, **Build**, **Patch**, and **Extension** drop-down lists may apply to the operating system you choose. In addition, if no definition appears in a list that matches the operating system you want to fingerprint, you can leave these values empty. Be aware that if you do not create any OS vulnerability mappings in a fingerprint, the system cannot use the fingerprint to assign a vulnerabilities list with hosts identified by the fingerprint.

Example:

If you want your custom fingerprint to assign the list of vulnerabilities from Redhat Linux 9 to matching hosts, choose **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the major version.

Example:

To add all versions of the Palm OS, you would choose **PalmSource, Inc.** from the **Vendor** list, **Palm OS** from the **Product** list, and leave all other lists at their default settings.

Step 13 Click **Create**.

The status briefly shows `New`, then switches to `Pending`, where it remains until traffic is seen for the fingerprint. Once traffic is seen, it switches to `Ready`.

The Custom Fingerprint status page refreshes every ten seconds until it receives data from the host in question.

Step 14 Using the IP address you specified as the target IP address, access the host you are trying to fingerprint and initiate a TCP connection to the appliance.

To create an accurate fingerprint, traffic **must** be seen by the appliance collecting the fingerprint. If you are connected through a switch, traffic to a system other than the appliance may not be seen by the system.

Example:

Access the web interface of the Firepower Management Center from the host you want to fingerprint or SSH into the FMC from the host. If you are using SSH, use the command below, where localIPv6address is the IPv6 address specified in step 7 that is currently assigned to the host and DCmanagementIPv6address is the management IPv6 address of the FMC. The Custom Fingerprint page should then reload with a “Ready” status.

```
ssh -b localIPv6address DCmanagementIPv6address
```

What to do next

- Activate the fingerprint as described in Activating and Deactivating Fingerprints, on page 1838.

Creating a Custom Fingerprint for Servers

Server fingerprints identify operating systems based on the SYN-ACK packet that the host uses to respond to an incoming connection to a running TCP application. Before you begin, you should obtain the following information about the host you want to fingerprint:

- The number of network hops between the host and the appliance you use to obtain the fingerprint. Cisco strongly recommends that you directly connect an unused interface on the appliance to the same subnet that the host is connected to.
- The network interface (on the appliance) that is connected to the network where the host resides.

- The actual operating system vendor, product, and version of the host.
- An IP address that is not currently in use and is authorized on the network where the host is located.



Tip If the Firepower Management Center does not have direct contact with monitored hosts, you can specify a managed device that is closest to the host you intend to fingerprint when specifying server fingerprint properties.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Custom Operating Systems**.

Step 3 Click **Create Custom Fingerprint**.

Step 4 From the **Device** list, choose the Firepower Management Center or the managed device that you want to use to collect the fingerprint.

Step 5 Enter a **Fingerprint Name**.

Step 6 Enter a **Fingerprint Description**.

Step 7 From the **Fingerprint Type** list, choose **Server** to display the server fingerprinting options.

Step 8 In the **Target IP Address** field, enter an IP address of the host you want to fingerprint.

Note that the fingerprint will only be based on traffic to and from the host IP address you specify, not any of the host's other IP addresses (if it has any).

Caution You can capture IPv6 fingerprints only with appliances running Version 5.2 and later of the Firepower System.

Step 9 In the **Target Distance** field, enter the number of network hops between the host and the device that you chose earlier to collect the fingerprint.

Caution This must be the actual number of physical network hops to the host, which may or may not be the same as the number of hops detected by the system.

Step 10 From the **Interface** list, choose the network interface that is connected to the network segment where the host resides.

Caution Cisco recommends that you do **not** use the sensing interface on a managed device for fingerprinting for several reasons. First, fingerprinting does not work if the sensing interface is on a span port. Also, if you use the sensing interface on a device, the device stops monitoring the network for the amount of time it takes to collect the fingerprint. You can, however, use the management interface or any other available network interfaces to perform fingerprint collection. If you do not know which interface is the sensing interface on your device, refer to the *Installation Guide* for the specific model you are using to fingerprint.

Step 11 Click **Get Active Ports**.

Step 12 In the **Server Port** field, enter the port that you want the device chose to collect the fingerprint to initiate contact with, or choose a port from the **Get Active Ports** drop-down list.

You can use any server port that you know is open on the host (for instance, 80 if the host is running a web server).

Step 13 In the **Source IP Address** field, enter an IP address that should be used to attempt to communicate with the host.

You should use a source IP address that is authorized for use on the network but is not currently being used, for example, a DHCP pool address that is currently not in use. This prevents you from temporarily knocking another host offline while you create the fingerprint.

You should exclude that IP address from monitoring in your network discovery policy while you create the fingerprint. Otherwise, the network map and discovery event views will be cluttered with inaccurate information about the host represented by that IP address.

Step 14 In the **Source Subnet Mask** field, enter the subnet mask for the IP address you are using.

Step 15 If the **Source Gateway** field appears, enter the default gateway IP address that should be used to establish a route to the host.

Step 16 If you want to display custom information in the host profile for fingerprinted hosts or if the fingerprint name you want to use does not exist in the OS Definition section, choose **Use Custom OS Display** in the Custom OS Display section.

Provide the values you want to appear in host profiles for the following:

- In the **Vendor String** field, enter the operating system's vendor name. For example, the vendor for Microsoft Windows would be Microsoft.
- In the **Product String** field, enter the operating system's product name. For example, the product name for Microsoft Windows 2000 would be Windows.
- In the **Version String** field, enter the operating system's version number. For example, the version number for Microsoft Windows 2000 would be 2000.

Step 17 In the OS Vulnerability Mappings section, choose the operating system, product, and versions you want to use for vulnerability mapping.

You must specify a Vendor and Product name in this section if you want to use the fingerprint to identify vulnerabilities for matching hosts or if you do not assign custom operating system display information.

To map vulnerabilities for all versions of an operating system, specify only the vendor and product name.

Note Not all options in the **Major Version**, **Minor Version**, **Revision Version**, **Build**, **Patch**, and **Extension** drop-down lists may apply to the operating system you choose. In addition, if no definition appears in a list that matches the operating system you want to fingerprint, you can leave these values empty. Be aware that if you do not create any OS vulnerability mappings in a fingerprint, the system cannot use the fingerprint to assign a vulnerabilities list with hosts identified by the fingerprint.

Example:

If you want your custom fingerprint to assign the list of vulnerabilities from Redhat Linux 9 to matching hosts, choose **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the version.

Example:

To add all versions of the Palm OS, you would choose **PalmSource, Inc.** from the **Vendor** list, **Palm OS** from the **Product** list, and leave all other lists at their default settings.

Step 18 Click **Create**.

The Custom Fingerprint status page refreshes every ten seconds and should reload with a “Ready” status.

Note If the target system stops responding during the fingerprinting process, the status shows an `ERROR: No Response` message. If you see this message, submit the fingerprint again. Wait three to five minutes (the time period may vary depending on the target system), click **Edit** () to access the Custom Fingerprint page, and then click **Create**.

What to do next

- Activate the fingerprint as described in Activating and Deactivating Fingerprints, on page 1838.

Host Input Data

You can augment the network map by importing network map data from third parties. You can also use the host input feature by modifying operating system or application identities or deleting application protocols, protocols, host attributes, or clients using the web interface.

The system may reconcile data from multiple sources to determine the current identity of an operating system or application.

All data except third-party vulnerabilities is discarded when the affected host is removed from the network map. For more information on setting up scripts or import files, see the *Firepower System Host Input API Guide*.

To include imported data in impact correlations, you must map the data to the operating system and application definitions in the database.

Requirements for Using Third-Party Data

You can import discovery data from third-party systems on your network. However, to enable features where intrusion and discovery data are used together, such as Firepower recommendations, adaptive profile updates, or impact assessment, you should map as many elements of it as possible to corresponding definitions. Consider the following requirements for using third-party data:

- If you have a third-party system that has specific data on your network assets, you can import that data using the host input feature. However, because third parties may name the products differently, you must map the third-party vendor, product, and versions to the corresponding Cisco product definition. After you map the products, you must enable vulnerability mappings for impact assessment in the Firepower Management Center configuration to allow impact correlation. For versionless or vendorless application protocols, you need to map vulnerabilities for the application protocols in the Firepower Management Center configuration.
- If you import patch information from a third party and you want to mark all vulnerabilities fixed by that patch as invalid, you must map the third-party fix name to a fix definition in the database. All vulnerabilities addressed by the fix will then be removed from hosts where you add that fix.
- If you import operating system and application protocol vulnerabilities from a third party and you want to use them for impact correlation, you must map the third-party vulnerability identification string to vulnerabilities in the database. Note that although many clients have associated vulnerabilities, and clients

are used for impact assessment, you cannot import and map third-party client vulnerabilities. After the vulnerabilities are mapped, you must enable third-party vulnerability mappings for impact assessment in the Firepower Management Center configuration. To cause application protocols without vendor or version information to map to vulnerabilities, an administrative user must also map vulnerabilities for the applications in the Firepower Management Center configuration.

- If you import application data and you want to use that data for impact correlation, you must map the vendor string for each application protocol to the corresponding Cisco application protocol definition.

Related Topics

Mapping Third-Party Products, on page 1846
Mapping Third-Party Product Fixes, on page 1848
Mapping Third-Party Vulnerabilities, on page 1849
Mapping Vulnerabilities for Servers, on page 918
Creating Custom Product Mappings, on page 1850

Third-Party Product Mappings

When you add data from third parties to the network map through the user input feature, you must map the vendor, product, and version names used by the third party to the Cisco product definitions. Mapping the products to Cisco definitions assigns vulnerabilities based on those definitions.

Similarly, if you are importing patch information from a third party, such as a patch management product, you must map the name for the fix to the appropriate vendor and product and the corresponding fix in the database.

Mapping Third-Party Products

If you import data from a third party, you must map the Cisco product to the third-party name to assign vulnerabilities and perform impact correlation using that data. Mapping the product associates Cisco vulnerability information with the third-party product name, which allows the system to perform impact correlation using that data.

If you import data using the host input import feature, you can also use the AddScanResult function to map third-party products to operating system and application vulnerabilities during the import.

For example, if you import data from a third party that lists Apache Tomcat as an application and you know it is version 6 of that product, you could add a third-party map where:

- **Vendor Name** is set to `Apache`.
- **Product Name** is set to `Tomcat`.
- **Apache** is chosen from the **Vendor** drop-down list.
- **Tomcat** is chosen from the **Product** drop-down list.
- **6** is chosen from the **Version** drop-down list

This mapping would cause any vulnerabilities for Apache Tomcat 6 to be assigned to hosts with an application listing for Apache Tomcat.

Note that for versionless or vendorless applications, you must map vulnerabilities for the application types in the Firepower Management Center configuration. Although many clients have associated vulnerabilities, and clients are used for impact assessment, you cannot import and map third-party client vulnerabilities.

Tip If you have already created a third-party mapping on another Firepower Management Center, you can export it and then import it onto this FMC. You can then edit the imported mapping to suit your needs.

Procedure

Step 1 Choose **Policies** > **Application Detectors**.

Step 2 Click **User Third-Party Mappings**.

Step 3 You have two choices:

- Create — To create a new map set, click **Create Product Map Set**.
- Edit — To edit an existing map set, click **Edit** ( ) next to the map set you want to modify. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Enter a **Mapping Set Name**.

Step 5 Enter a **Description**.

Step 6 You have two choices:

- Create — To map a third-party product, click **Add Product Map**.
- Edit — To edit an existing third-party product map, **Edit** ( ) next to the map set you want to modify. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 7 Enter the **Vendor String** used by the third-party product.

Step 8 Enter the **Product String** used by the third-party product.

Step 9 Enter the **Version String** used by the third-party product.

Step 10 In the Product Mappings section, choose the operating system, product, and versions you want to use for vulnerability mapping from the **Vendor**, **Product**, **Major Version**, **Minor Version**, **Revision Version**, **Build**, **Patch**, and **Extension** fields.

Example:

If you want a host running a product whose name consists of third-party strings to use the vulnerabilities from Red Hat Linux 9, choose **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the version.

Step 11 Click **Save**.

Related Topics

Mapping Vulnerabilities for Servers, on page 918

Mapping Third-Party Product Fixes

If you map a fix name to a particular set of fixes in the database, you can then import data from a third-party patch management application and apply the fix to a set of hosts. When the fix name is imported to a host, the system marks all vulnerabilities addressed by the fix as invalid for that host.

Procedure

Step 1 Choose **Policies** > **Application Detectors**.

Step 2 Click **User Third-Party Mappings**.

Step 3 You have two choices:

- Create — To create a new map set, click **Create Product Map Set**.
- Edit — To edit an existing map set, click **Edit** ( ) next to the map set you want to modify. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Enter a **Mapping Set Name**.

Step 5 Enter a **Description**.

Step 6 You have two choices:

- Create — To map a third-party product, click **Add Fix Map.**
- Edit — To edit an existing third-party product map, click **Edit** ( ) next to it. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 7 Enter the name of the fix you want to map in the **Third-Party Fix Name** field.

Step 8 In the **Product Mappings** section, choose the operating system, product, and versions you want to use for fix mapping from the following fields:

- **Vendor**
- **Product**
- **Major Version**
- **Minor Version**
- **Revision Version**
- **Build**
- **Patch**
- **Extension**

Example:

If you want your mapping to assign the fixes from Red Hat Linux 9 to hosts where the patch is applied, choose **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the version.

Step 9 Click **Save** to save the fix map.

Mapping Third-Party Vulnerabilities

To add vulnerability information from a third party to the VDB, you must map the third-party identification string for each imported vulnerability to any existing SVID, Bugtraq, or SID. After you create a mapping for the vulnerability, the mapping works for all vulnerabilities imported to hosts in the network map and allows impact correlation for those vulnerabilities.

You must enable impact correlation for third-party vulnerabilities to allow correlation to occur. For versionless or vendorless applications, you must also map vulnerabilities for the application types in the Firepower Management Center configuration.

Although many clients have associated vulnerabilities, and clients are used for impact assessment, you cannot use third-party client vulnerabilities for impact assessment.

Tip If you have already created a third-party mapping on another Firepower Management Center, you can export it and then import it onto this FMC. You can then edit the imported mapping to suit your needs.

Procedure

Step 1 Choose **Policies** > **Application Detectors**.

Step 2 Click **User Third-Party Mappings**.

Step 3 You have two choices:

- Create — To create a new vulnerability set, click **Create Vulnerability Map Set**.
- Edit — To edit an existing vulnerability set, click **Edit** ( ) next to the vulnerability set. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Click **Add Vulnerability Map**.

Step 5 Enter the third-party identification for the vulnerability in the **Vulnerability ID** field.

Step 6 Enter a **Vulnerability Description**.

Step 7 Optionally:

- Enter a Snort ID in the **Snort Vulnerability ID Mappings** field.
- Enter a legacy vulnerability ID in the **SVID Mappings** field.
- Enter a Bugtraq identification number in the **Bugtraq Vulnerability ID Mappings** field.

Step 8 Click **Add**.

Related Topics

Enabling Network Discovery Vulnerability Impact Assessment, on page 1932
Mapping Vulnerabilities for Servers, on page 918

Custom Product Mappings

You can use product mappings to ensure that servers input by a third party are associated with the appropriate Cisco definitions. After you define and activate the product mapping, all servers or clients on monitored hosts that have the mapped vendor strings use the custom product mappings. For this reason, you may want to map vulnerabilities for all servers in the network map with a particular vendor string instead of explicitly setting the vendor, product, and version for the server.

Creating Custom Product Mappings

If the system cannot map a server to a vendor and product in the VDB, you can manually create the mapping. When you activate a custom product mapping, the system maps vulnerabilities for the specified vendor and product to all servers in the network map where that vendor string occurs.



Note Custom product mappings apply to all occurrences of an application protocol, regardless of the source of the application data (such as Nmap, the host input feature, or the Firepower System itself). However, if third-party vulnerability mappings for data imported using the host input feature conflicts with the mappings you set through a custom product mapping, the third-party vulnerability mapping overrides the custom product mapping and uses the third-party vulnerability mapping settings when the input occurs.

You create lists of product mappings and then enable or disable use of several mappings at once by activating or deactivating each list. When you specify a vendor to map to, the system updates the list of products to include only those made by that vendor.

After you create a custom product mapping, you must activate the custom product mapping list. After you activate a list of custom product mappings, the system updates all servers with occurrences of the specified vendor strings. For data imported through the host input feature, vulnerabilities update unless you have already explicitly set the product mappings for this server.

If, for example, your company modifies the banner for your Apache Tomcat web servers to read `Internal Web Server`, you can map the vendor string `Internal Web Server` to the vendor **Apache** and the product **Tomcat**, then activate the list containing that mapping, all hosts where a server labeled `Internal Web Server` occurs have the vulnerabilities for Apache Tomcat in the database.



Tip You can use this feature to map vulnerabilities to local intrusion rules by mapping the SID for the rule to another vulnerability.

Procedure

Step 1 Choose **Policies** > **Application Detectors**.

Step 2 Click **Custom Product Mappings**

Step 3 Click **Create Custom Product Mapping List**.

Step 4 Enter a **Custom Product Mapping List Name**.

Step 5 Click **Add Vendor String**.

Step 6 In the **Vendor String** field, enter the vendor string that identifies the applications that should map to the chosen vendor and product values.

Step 7 Choose the vendor you want to map to from the **Vendor** drop-down list.

Step 8 Choose the product you want to map to from the **Product** drop-down list.

Step 9 Click **Add** to add the mapped vendor string to the list.

Step 10 Optionally, repeat steps 4 to 8 as needed to add additional vendor string mappings to the list.

Step 11 Click **Save**.

What to do next

- Activate the custom product mapping list. For more information, see Activating and Deactivating Custom Product Mappings, on page 1851.

Editing Custom Product Mapping Lists

You can modify existing custom product mapping lists by adding or removing vendor strings or changing the list name.

Procedure

Step 1 Choose **Policies** > **Application Detectors**.

Step 2 Click **Custom Product Mappings**.

Step 3 Click **Edit** () next to the product mapping list you want to edit.

If **View** () appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Make changes to the list as described in Creating Custom Product Mappings, on page 1850.

Step 5 When you finish, click **Save**.

Activating and Deactivating Custom Product Mappings

You can enable or disable use of an entire list of custom product mappings at once. After you activate a custom product mapping list, each mapping on that list applies to all applications with the specified vendor string, whether detected by managed devices or imported through the host input feature.

Procedure

Step 1 Choose **Policies** > **Application Detectors**.

Step 2 Click **Custom Product Mappings**.

Step 3 Click the slider next to the custom product mapping list to activate or deactivate it.

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

eStreamer Server Streaming

The Event Streamer (eStreamer) allows you to stream several kinds of event data from a Firepower Management Center or 7000 or 8000 Series device to a custom-developed client application. For more information, see *Firepower System Event Streamer Integration Guide*.

Before the appliance you want to use as an eStreamer server can begin streaming eStreamer events to an external client, you must configure the eStreamer server to send events to clients, provide information about the client, and generate a set of authentication credentials to use when establishing communication. You can perform all of these tasks from the appliance's user interface. Once your settings are saved, the events you selected will be forwarded to eStreamer clients when requested.

You can control which types of events the eStreamer server is able to transmit to clients that request them.

Table 218: Event Types Transmittable by the eStreamer Server

Event Type	Description	Available on FMC	Available on 7000 & 8000 Series Devices
Intrusion Events	intrusion events generated by managed devices	yes	yes
Intrusion Event Packet Data	packets associated with intrusion events	yes	yes
Intrusion Event Extra Data	additional data associated with an intrusion event such as the originating IP addresses of a client connecting to a web server through an HTTP proxy or load balancer	yes	yes
Discovery Events	Network discovery events	yes	no
Correlation and White List Events	correlation and white list events	yes	no
Impact Flag Alerts	impact alerts generated by the FMC	yes	no
User Events	user events	yes	no
Malware Events	malware events	yes	no
File Events	file events	yes	no
Connection Events	information about the session traffic between your monitored hosts and all other hosts.	yes	yes

Choosing eStreamer Event Types

The **eStreamer Event Configuration** check boxes control which events the eStreamer server can transmit. Your client must still specifically request the types of events you want it to receive in the request message it

sends to the eStreamer server. For more information, see the *Firepower System Event Streamer Integration Guide*.

In a multidomain deployment, you can configure eStreamer Event Configuration at any domain level. However, if an ancestor domain has enabled a particular event type, you cannot disable that event type in the descendant domains.

You must be an Admin user to perform this task, for FMC and 7000 & 8000 Series devices.

Procedure

Step 1 Choose **System** > **Integration**.

Step 2 Click **eStreamer**.

Step 3 Under **eStreamer Event Configuration**, check or clear the check boxes next to the types of events you want eStreamer to forward to requesting clients, described in eStreamer Server Streaming, on page 1852.

Step 4 Click **Save**.

Configuring eStreamer Client Communications

Before eStreamer can send eStreamer events to a client, you must add the client to the eStreamer server's peers database from the eStreamer page. You must also copy the authentication certificate generated by the eStreamer server to the client. After completing these steps you do not need to restart the eStreamer service to enable the client to connect to the eStreamer server.

In a multidomain deployment, you can create an eStreamer client in any domain. The authentication certificate allows the client to request events only from the client certificate's domain and any descendant domains. The eStreamer configuration page shows only clients associated with the current domain, so if you want to download or revoke a certificate, switch to the domain where the client was created.

You must be an Admin or Discovery Admin user to perform this task, for FMC and 7000 & 8000 Series devices.

Procedure

Step 1 Choose **System** > **Integration**.

Step 2 Click **eStreamer**.

Step 3 Click **Create Client**.

Step 4 In the **Hostname** field, enter the host name or IP address of the host running the eStreamer client.

Note If you have not configured DNS resolution, use an IP address.

Step 5 If you want to encrypt the certificate file, enter a password in the **Password** field.

Step 6 Click **Save**.
The eStreamer server now allows the host to access port 8302 on the eStreamer server and creates an authentication certificate to use during client-server authentication.

Step 7 Click **Download** ( ) next to the client hostname to download the certificate file.

Step 8 Save the certificate file to the appropriate directory used by your client for SSL authentication.

Step 9 To revoke access for a client, click **Delete** ( ) next to the host you want to remove.

Note that you do not need to restart the eStreamer service; access is revoked immediately.

Configuring the Host Input Client

The host input feature allows you to update the Firepower Management Center's network map from a client program running on another appliance. For example, you can add or delete hosts from the network map, or update the host OS and service information. For more information, see *Firepower System Host Input API Guide*.

Before you can run a remote client, you must add the client to the Firepower Management Center's peers database from the Host Input Client page. You must also copy the authentication certificate generated by the FMC to the client. After completing these steps the client can connect to the FMC.

In a multidomain deployment, you can create a client in any domain. The authentication certificate allows the client to submit network map updates for any leaf domains associated with the client certificate's domain. If you create a certificate for an ancestor domain (or if your certificate domain later becomes an ancestor domain after adding descendant domains), any clients using that certificate must specify a target leaf domain with every transaction, as described in the *Firepower System Host Input API Guide*.

The Host Input Client shows only clients associated the current domain, so if you want to download or revoke a certificate, switch to the domain where the client was created.

Procedure

Step 1 Choose **System** > **Integration**.

Step 2 Click **Host Input Client**.

Step 3 Click **Create Client**.

Step 4 In the **Hostname** field, enter the host name or IP address of the host running the host input client.

Note If you have not configured DNS resolution, use an IP address.

Step 5 If you want to encrypt the certificate file, enter a password in the **Password** field.

Step 6 Click **Save**.
The host input service allows the host to access port 8307 on the Firepower Management Center and creates an authentication certificate to use during client-server authentication.

Step 7 Click **Download** ( ) next to the certificate file.

Step 8 Save the certificate file to the directory used by your client for SSL/TLS authentication.

Step 9 To revoke access for a client, click **Delete** ( ) next to the host you want to remove.

/firepower/fmc/fmc_config_guide/discovery-host-profiles/t_editing_server_identities.xml

Nmap Scanning

The Firepower System builds network maps through passive analysis of traffic on your network. Information obtained through this passive analysis can occasionally be incomplete, depending on system conditions. However, you can actively scan a host to obtain complete information. For example, if a host has a server running on an open port but the server has not received or sent traffic during the time that the system has been monitoring your network, the system does not add information about that server to the network map. If you directly scan that host using an active scanner, however, you can detect the presence of the server.

The Firepower System integrates with Nmap™, an open source active scanner for network exploration and security auditing.

When you scan a host using Nmap, the system:

- Adds servers on previously undetected open ports to the Servers list in the host profile for that host. The host profile lists any servers detected on filtered or closed TCP ports or on UDP ports in the Scan Results section. By default, Nmap scans more than 1660 TCP ports.

  If the system recognizes a server identified in an Nmap scan and has a corresponding server definition, the system maps the names Nmap uses for servers to the corresponding Cisco server definitions.

- Compares the results of the scan to over 1500 known operating system fingerprints to determine the operating system and assigns scores to each. The operating system assigned to the host is the operating system fingerprint with the highest score.

  The system maps Nmap operating system names to Cisco operating system definitions.

- Assigns vulnerabilities to the host for the added servers and operating systems.

Note:

- A host must exist in the network map before Nmap can append its results to the host profile.
- If the host is deleted from the network map, any Nmap scan results for that host are discarded.

Tip Some scanning options (such as portscans) may place a significant load on networks with low bandwidths. Schedule scans like these to run during periods of low network use.

For more information on the underlying Nmap technology used to scan, refer to the Nmap documentation at http://insecure.org/.

Related Topics

Nmap Scan Automation, on page 184

Nmap Remediation Options

You define the settings for an Nmap scan by creating an Nmap remediation. An Nmap remediation can be used as a response in a correlation policy, run on demand, or scheduled to run at a specific time.

Note that Nmap-supplied server and operating system data remain static until you run another Nmap scan. If you plan to scan a host for operating system and server data using Nmap, you may want to set up regularly scheduled scans to keep any Nmap-supplied operating system and server data up-to-date.

The following table explains the options configurable in Nmap remediations on a Firepower System.

Table 219: Nmap Remediation Options

Option	Description	Corresponding Nmap Option
Scan Which Address(es) From Event?	When you use an Nmap scan as a response to a correlation rule, select one of the following options to control which address in the event is scanned, that of the source host, the destination host, or both: • **Scan Source and Destination Addresses** scans the hosts represented by the source IP address and the destination IP address in the event. • **Scan Source Address Only** scans the host represented by the event's source IP address. • **Scan Destination Address Only** scans the host represented by the event's destination IP address.	N/A
Scan Types	Select how Nmap scans ports: • The **TCP Syn** scan connects quickly to thousand of ports without using a complete TCP handshake. This options allows you to scan quickly in stealth mode on hosts where the `admin` account has raw packet access or where IPv6 is not running, by initiating TCP connections but not completing them. If a host acknowledges the Syn packet sent in a TCP Syn scan, Nmap resets the connection. • The **TCP Connect** scan uses the `connect()` system call to open connections through the operating system on the host. You can use the TCP Connect scan if the `admin` user on the Firepower Management Center or managed device does not have raw packet privileges on a host or you are scanning IPv6 networks. In other words, use this option in situations where the TCP Syn scan cannot be used. • The **TCP ACK** scan sends an ACK packet to check whether ports are filtered or unfiltered. • The **TCP Window** scan works in the same way as a TCP ACK scan but can also determine whether a port is open or closed. • The **TCP Maimon** scan identifies BSD-derived systems using a FIN/ACK probe.	**TCP Syn**: `-sS` **TCP Connect**: `-sT` **TCP ACK**: `-sA` **TCP Window**: `-sW` **TCP Maimon**: `-sM`
Scan for UDP ports	Enable to scan UDP ports in addition to TCP ports. Note that scanning UDP ports may be time-consuming, so avoid using this option if you want to scan quickly.	`-sU`

Option	Description	Corresponding Nmap Option
Use Port From Event	If you plan to use the remediation as a response in a correlation policy, enable to cause the remediation to scan only the port specified in the event that triggers the correlation response. • Select **On** to scan the port in the correlation event, rather than the ports you specify during Nmap remediation configuration. If you scan the port in the correlation event, note that the remediation scans the port on the IP addresses that you specify during Nmap remediation configuration. These ports are also added to the remediation's dynamic scan target. • Select **Off** to scan only the ports you specify Nmap remediation configuration. You can also control whether Nmap collects information about operating system and server information. Enable the **Use Port From Event** option to scan the port associated with the new server.	N/A
Scan from reporting detection engine	Enable to scan a host from the appliance where the detection engine that reported the host resides. • To scan from the appliance running the reporting detection engine, select **On**. • To scan from the appliance configured in the remediation, select **Off**.	N/A
Fast Port Scan	Enable to scan only the TCP ports listed in the `nmap-services` file located in the `/var/sf/nmap/share/nmap/nmap-services` directory on the device that does the scanning, ignoring other port settings. Note that you cannot use this option with the **Port Ranges and Scan Order** option. • To scan only the ports listed in the `nmap-services` file located in the `/var/sf/nmap/share/nmap/nmap-services` directory on the device that does the scanning, ignoring other port settings, select **On**. • To scan all TCP ports, select **Off**.	`-F`
Port Ranges and Scan Order	Set the specific ports you want to scan, using Nmap port specification syntax, and the order you want to scan them. Note that you cannot use this option with the **Fast Port Scan** option.	`-p`
Probe open ports for vendor and version information	Enable to detect server vendor and version information. If you probe open ports for server vendor and version information, Nmap obtains server data that it uses to identify servers. It then replaces the Cisco server data for that server. • Select **On** to scan open ports on the host for server information to identify server vendors and versions. • Select **Off** to continue using Cisco server information for the host.	`-sV`

Option	Description	Corresponding Nmap Option
Service Version Intensity	Select the intensity of Nmap probes for service versions. • To use more probes for higher accuracy with a longer scan, select a higher number. • To use fewer probes for less accuracy with a faster scan, select a lower number.	`--version-intensity` <intensity>
Detect Operating System	Enable to detect operating system information for the host. If you configure detection of the operating system for a host, Nmap scans the host and uses the results to create a rating for each operating system that reflects the likelihood that the operating system is running on the host. • Select **On** to scan the host for information to identify the operating system. • Select **Off** to continue using Cisco operating system information for the host.	`-O`
Treat All Hosts As Online	Enable to skip the host discovery process and run a port scan on every host in the target range. Note that when you enable this option, Nmap ignores settings for **Host Discovery Method** and **Host Discovery Port List**. • To skip the host discovery process and run a port scan on every host in the target range, select **On**. • To perform host discovery using the settings for **Host Discovery Method** and **Host Discovery Port List** and skip the port scan on any host that is not available, select **Off**.	`-PN`
Host Discovery Method	Select to perform host discovery for all hosts in the target range, over the ports listed in the **Host Discovery Port Lis**t, or if no ports are listed, over the default ports for that host discovery method. Note that if you also enabled **Treat All Hosts As Online**, however, the **Host Discovery Method** option has no effect and host discovery is not performed. Select the method to be used when Nmap tests to see if a host is present and available: • The **TCP SYN** option sends an empty TCP packet with the SYN flag set and recognizes the host as available if a response is received. TCP SYN scans port 80 by default. Note that TCP SYN scans are less likely to be blocked by a firewall with stateful firewall rules. • The **TCP ACK** option sends an empty TCP packet with the ACK flag set and recognizes the host as available if a response is received. TCP ACK also scans port 80 by default. Note that TCP ACK scans are less likely to be blocked by a firewall with stateless firewall rules. • The **UDP** option sends a UDP packet and assumes host availability if a port unreachable response comes back from a closed port. UDP scans port 40125 by default.	**TCP SYN**: `-PS` **TCP ACK**: `-PA` **UDP**: `-PU`

Option	Description	Corresponding Nmap Option
Host Discovery Port List	Specify a customized list of ports, separated by commas, that you want to scan when doing host discovery.	port list for host discovery method
Default NSE Scripts	Enable to run the default set of Nmap scripts for host discovery and server and operating system and vulnerability detection. See https://nmap.org/nsedoc/categories/default.html for the list of default scripts. • To run the default set of Nmap scripts, select **On**. • To skip the default set of Nmap scripts, select **Off**.	`-sC`
Timing Template	Select the timing of the scan process; the higher the number you select, the faster and less comprehensive the scan.	**0**: `T0` (paranoid) **1**: `T1` (sneaky) **2**: `T2` (polite) **3**: `T3` (normal) **4**: `T4` (aggressive) **5**: `T5` (insane)

Nmap Scanning Guidelines

While active scanning can obtain valuable information, overuse of a tool such as Nmap may overload your network resources or even crash important hosts. When using any active scanner, you should create a scanning strategy following these guidelines to make sure that you are scanning only the hosts and ports that you need to scan.

Selecting Appropriate Scan Targets

When you configure Nmap, you can create scan targets that identify which hosts you want to scan. A scan target includes a single IP address, a CIDR block or octet range of IP addresses, an IP address range, or a list of IP addresses or ranges to scan, as well as the ports on the host or hosts.

You can specify targets in the following ways:

- For IPv6 hosts:
  - an exact IP address (for example, `192.168.1.101`)
- For IPv4 hosts:
  - an exact IP address (for example, `192.168.1.101`) or a list of IP addresses separated by commas or spaces
  - an IP address block using CIDR notation (for example, `192.168.1.0/24` scans the 254 hosts between 192.168.1.1 and 192.168.1.254, inclusive).
  - an IP address range using octet range addressing (for example, `192.168.0-255.1-254` scans all addresses in the 192.168.x.x range, except those that end in .0 and or .255)

- an IP address range using hyphenation (for example, `192.168.1.1 - 192.168.1.5` scans the six hosts between 192.168.1.1 and 192.168.1.5, inclusive)
- a list of addresses or ranges separated by commas or spaces (for example, for example, `192.168.1.0/24, 194.168.1.0/24` scans the 254 hosts between 192.168.1.1 and 192.168.1.254, inclusive and the 254 hosts between 194.168.1.1 and 194.168.1.254, inclusive)

Ideal scan targets for Nmap scans include hosts with operating systems that the system is unable to identify, hosts with unidentified servers, or hosts recently detected on your network. Remember that Nmap results cannot be added to the network map for hosts that do not already exist in the network map.

Caution

- Nmap-supplied server and operating system data remains static until you run another Nmap scan. If you plan to scan a host using Nmap, regularly schedule scans.
- If a host is deleted from the network map, any Nmap scan results are discarded.
- Make sure you have permission to scan your targets. Using Nmap to scan hosts that do not belong to you or your company may be illegal.

Selecting Appropriate Ports to Scan

For each scan target you configure, you can select the ports you want to scan. You can designate individual port numbers, port ranges, or a series of port numbers and port ranges to identify the exact set of ports that should be scanned on each target.

By default, Nmap scans TCP ports 1 through 1024. If you plan to use the remediation as a response in a correlation policy, you can cause the remediation to scan only the port specified in the event that triggers the correlation response. If you run the remediation on demand or as a scheduled task, or if you do not use the port from the event, you can use other port options to determine which ports are scanned. You can choose to scan only the TCP ports listed in the `nmap-services` file, ignoring other port settings. You can also scan UDP ports in addition to TCP ports. Note that scanning for UDP ports may be time-consuming, so avoid using that option if you want to scan quickly. To select the specific ports or range of ports to scan, use Nmap port specification syntax to identify ports.

Setting Host Discovery Options

You can decide whether to perform host discovery before starting a port scan for a host, or you can assume that all the hosts you plan to scan are online. If you choose not to treat all hosts as online, you can choose what method of host discovery to use and, if needed, customize the list of ports scanned during host discovery. Host discovery does not probe the ports listed for operating system or server information; it uses the response over a particular port only to determine whether a host is active and available. If you perform host discovery and a host is not available, Nmap does not scan ports on that host.

Related Topics

Firepower System IP Address Conventions, on page 15
Nmap Scan Automation, on page 184

Example: Using Nmap to Resolve Unknown Operating Systems

This example walks through an Nmap configuration designed to resolve unknown operating systems. For a complete look at Nmap configuration, see Managing Nmap Scanning, on page 1863.

If the system cannot determine the operating system on a host on your network, you can use Nmap to actively scan the host. Nmap uses the information it obtains from the scan to rate the possible operating systems. It then uses the operating system that has the highest rating as the host operating system identification.

Using Nmap to challenge new hosts for operating system and server information deactivates the system's monitoring of that data for scanned hosts. If you use Nmap to discover host and server operating system for hosts the system marks as having unknown operating systems, you may be able to identify groups of hosts that are similar. You can then create a custom fingerprint based on one of them to cause the system to associate the fingerprint with the operating system you know is running on the host based on the Nmap scan. Whenever possible, create a custom fingerprint rather than inputting static data through a third-party source like Nmap because the custom fingerprint allows the system to continue to monitor the host operating system and update it as needed.

In this example, you would:

1. Configure a scan instance as described in Adding an Nmap Scan Instance, on page 1863.
2. Create an Nmap remediation using the following settings:
   - Enable **Use Port From Event** to scan the port associated with the new server.
   - Enable **Detect Operating System** to detect operating system information for the host.
   - Enable **Probe open ports for vendor and version information** to detect server vendor and version information.
   - Enable **Treat All Hosts as Online**, because you know the host exists.
3. Create a correlation rule that triggers when the system detects a host with an unknown operating system. The rule should trigger when **a discovery event occurs** and **the OS information for a host has changed** and it meets the following conditions: **OS Name is unknown**.
4. Create a correlation policy that contains the correlation rule.
5. n the correlation policy, add the Nmap remediation you created in step 2 as a response to the rule you created in step 3.
6. Activate the correlation policy.
7. Purge the hosts on the network map to force network discovery to restart and rebuild the network map.
8. After a day or two, search for events generated by the correlation policy. Analyze the Nmap results for the operating systems detected on the hosts to see if there is a particular host configuration on your network that the system does not recognize.
9. If you find hosts with unknown operating systems whose Nmap results are identical, create a custom fingerprint for one of those hosts and use it to identify similar hosts in the future.

Related Topics

Creating an Nmap Remediation, on page 1867
Configuring Correlation Rules, on page 1989
Nmap Scan Results, on page 1870
Creating a Custom Fingerprint for Clients, on page 1840
Configuring Correlation Policies, on page 1987

Example: Using Nmap to Respond to New Hosts

This example walks through an Nmap configuration designed to respond to new hosts. For a complete look at Nmap configuration, see Managing Nmap Scanning, on page 1863.

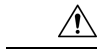
When the system detects a new host in a subnet where intrusions may be likely, you may want to scan that host to make sure you have accurate vulnerability information for it.

You can accomplish this by creating and activating a correlation policy that detects when a new host appears in this subnet, and that launches a remediation that performs an Nmap scan on the host.

To do this, you would:

1. Configure a scan instance as described in Adding an Nmap Scan Instance, on page 1863.
2. Create an Nmap remediation using the following settings:
   - Enable **Use Port From Event** to scan the port associated with the new server.
   - Enable **Detect Operating System** to detect operating system information for the host.
   - Enable **Probe open ports for vendor and version information** to detect server vendor and version information.
   - Enable **Treat All Hosts as Online**, because you know the host exists.
3. Create a correlation rule that triggers when the system detects a new host on a specific subnet. The rule should trigger when **a discovery event occurs** and **a new host is detected**.
4. Create a correlation policy that contains the correlation rule.
5. In the correlation policy, add the Nmap remediation you created in step 2 as a response to the rule you created in step 3.
6. Activate the correlation policy.
7. When you are notified of a new host, check the host profile to see the results of the Nmap scan and address any vulnerabilities that apply to the host.

After you activate the policy, you can periodically check the remediation status view (**Analysis** > **Correlation** > **Status**) to see when the remediation launched. The remediation's dynamic scan target should include the IP addresses of the hosts it scanned as a result of the server detection. Check the host profile for those hosts to see if there are vulnerabilities that need to be addressed for the host, based on the operating system and servers detected by Nmap.



Caution If you have a large or dynamic network, detection of a new host may be too frequent an occurrence to respond to using a scan. To prevent resource overload, avoid using Nmap scans as a response to events that occur frequently. In addition, note that using Nmap to challenge new hosts for operating system and server information deactivates Cisco monitoring of that data for scanned hosts.

Related Topics

Creating an Nmap Remediation, on page 1867
Configuring Correlation Rules, on page 1989
Configuring Correlation Policies, on page 1987

Managing Nmap Scanning

To use Nmap scanning, you must, at minimum, configure an Nmap scan instance and an Nmap remediation. Configuring an Nmap scan target is optional.

Procedure

Step 1 Configure the Nmap scan:

- Add an Nmap scan instance as described in Adding an Nmap Scan Instance, on page 1863.
- Create an Nmap remediation as described in Creating an Nmap Remediation, on page 1867.
- Optionally, add an Nmap scan target as described in Adding an Nmap Scan Target, on page 1865.

Step 2 Run the Nmap scan:

- Run an on-demand Nmap scan as described in Running an On-Demand Nmap Scan, on page 1869.
- Configure automatic Nmap scans as described in Nmap Scan Automation, on page 184.
- Schedule automatic Nmap scans as described in Scheduling an Nmap Scan, on page 184.

What to do next

- Monitor the Nmap scan in progress by viewing the related task; see Viewing Task Messages, on page 316.
- Optionally, refine the scan:
  - Edit an Nmap scan instance as described in Editing an Nmap Scan Instance, on page 1864.
  - Edit an Nmap scan target as described in Editing an Nmap Scan Target, on page 1866.
  - Edit an Nmap remediation as described in Editing an Nmap Remediation, on page 1869.

Adding an Nmap Scan Instance

You can set up a separate scan instance for each Nmap module that you want to use to scan your network for vulnerabilities. You can set up scan instances for the local Nmap module on the Firepower Management Center and for any devices you want to use to run scans remotely. The results of each scan are always stored on the Firepower Management Center where you configure the scan, even if you run the scan from a remote device. To prevent accidental or malicious scanning of mission-critical hosts, you can create a blacklist for the instance to indicate the hosts that should never be scanned with the instance.

You cannot add a scan instance with the same name as any existing scan instance.

In a multidomain deployment, the system displays scan instances created in the current domain, which you can edit. It also displays scan instances created in ancestor domains, which you cannot edit. To view and edit scan instances in a lower domain, switch to that domain.

Procedure

Step 1 Access the list of Nmap scan instances using either of the following methods:

- Choose **Policies** > **Actions** > **Instances**.
- Choose **Policies** > **Actions** > **Scanners**.

Step 2 Add the remediation:

- If you accessed the list via the first method above, locate the Add a New Instance section, choose the Nmap Remediation module from the drop-down list, and click **Add**.
- If you accessed the list via the second method above, click **Add Nmap Instance**.

Step 3 Enter an **Instance Name**.

Step 4 Enter a **Description**.

Step 5 Optionally, in the **Blacklisted Scan hosts** field, specify any hosts or networks that should *never* be scanned with this scan instance, using the following syntax:

- For IPv6 hosts, an exact IP address (for example, `2001:DB8::fedd:eeff`)
- For IPv4 hosts, an exact IP address (for example, `192.168.1.101`) or an IP address block using CIDR notation (for example, `192.168.1.0/24` scans the 254 hosts between `192.168.1.1` and `192.168.1.254`, inclusive)
- Note that you cannot use an exclamation mark (!) to negate an address value.

Note If you specifically target a scan to a host that is in a blacklisted network, that scan will not run.

Step 6 Optionally, to run the scan from a remote device instead of the Firepower Management Center, specify the IP address or name of the device as it appears in the Information page for the device in the FMC web interface, in the **Remote Device Name** field.

Step 7 Click **Create**.
When the system is done creating the instance, it displays it in edit mode.

Step 8 Optionally, add an Nmap remediation to the instance. To do so, locate the Configured Remediations section of the instance, click **Add**, and create a remediation as described in Creating an Nmap Remediation, on page 1867.

Step 9 Click **Cancel** to return to the list of instances.

Note If you accessed the list of Nmap scan instances via the **Scanners** option, the system does not display the instance you added unless you also added a remediation to the instance. To view any instances to which you have not yet added remedations, use the **Instances** menu option to access the list.

Editing an Nmap Scan Instance

When you edit a scan instance, you can view, add, and delete remediations associated with the instance. Delete an Nmap scan instance when you no longer want to use the Nmap module profiled in the instance. Note that when you delete the scan instance, you also delete any remediations that use that instance.

In a multidomain deployment, the system displays scan instances created in the current domain, which you can edit. It also displays scan instances created in ancestor domains, which you cannot edit. To view and edit scan instances in a lower domain, switch to that domain.

Procedure

Step 1 Access the list of Nmap scan instances using either of the following methods:

- Choose **Policies** > **Actions** > **Instances**.
- Choose **Policies** > **Actions** > **Scanners**.

Step 2 Click **View** () next to the instance you want to edit.

Step 3 Make changes to the scan instance settings as described in Adding an Nmap Scan Instance, on page 1863.

Step 4 Click **Save**.

Step 5 Click **Done**.

What to do next

- Optionally, add a new remediation to the scan instance; see Creating an Nmap Remediation, on page 1867
- Optionally, edit a remediation associated with the instance; see Editing an Nmap Remediation, on page 1869.
- Optionally, delete a remediation associated with the instance; see Running an On-Demand Nmap Scan, on page 1869.
- Optionally, delete the scan instance by clicking **Delete** () next to it.

Adding an Nmap Scan Target

When you configure an Nmap module, you can create and save scan targets that identify the hosts and ports you want to target when you perform an on-demand or a scheduled scan, so that you do not have to construct a new scan target every time. A scan target includes a single IP address or a block of IP addresses to scan, as well as the ports on the host or hosts. For Nmap targets, you can also use Nmap octet range addressing or IP address ranges. For more information on Nmap octet range addressing, refer to the Nmap documentation at http://insecure.org.

Note:

- Scans for scan targets containing a large number of hosts can take an extended period of time. As a workaround, scan fewer hosts at a time.
- Nmap-supplied server and operating system data remains static until you run another Nmap scan. If you plan to scan a host using Nmap, regularly schedule scans. If a host is deleted from the network map, any Nmap scan results are discarded.
- In a multidomain deployment, the system displays scan targets created in the current domain, which you can edit. It also displays scan targets created in ancestor domains, which you cannot edit. To view and edit scan targets in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Policies** > **Actions** > **Scanners**.

Step 2 On the toolbar, click **Targets**.

Step 3 Click **Create Scan Target**.

Step 4 In the **Name** field, enter the name you want to use for this scan target.

Step 5 In the **IP Range** text box, specify the host or hosts you want to scan using the syntax described in Nmap Scanning Guidelines, on page 1859.

Note If you use a comma in a list of IP addresses or ranges in a scan target, the comma converts to a space when you save the target.

Step 6 In the **Ports** field, specify the ports you want to scan.

You can enter any of the following, using values from 1 to 65535:

- a port number
- a list of ports separated by commas
- a range of port numbers separated by a dash
- ranges of port numbers separated by dashes, separated by commas

Step 7 Click **Save**.

Related Topics

Nmap Scan Automation, on page 184

Editing an Nmap Scan Target

Tip You might want to edit a remediation's dynamic scan target if you do not want to use the remediation to scan a specific IP address, but the IP address was added to the target because the host was involved in a correlation policy violation that launched the remediation.

Delete a scan target if you no longer want to scan the hosts listed in it.

In a multidomain deployment, the system displays scan targets created in the current domain, which you can edit. It also displays scan targets created in ancestor domains, which you cannot edit. To view and edit scan targets in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Policies** > **Actions** > **Scanners**.

Step 2 On the toolbar, click **Targets**.

Step 3 Click **Edit** ( ) next to the scan target you want to edit.

If **View** () appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Make modifications as necessary. For more information, see Adding an Nmap Scan Target, on page 1865.

Step 5 Click **Save**.

Step 6 Optionally, delete the scan target by clicking **Delete** () next to it.

Creating an Nmap Remediation

An Nmap remediation can only be created by adding it to an existing Nmap scan instance. The remediation defines the settings for the scan. It can be used as a response in a correlation policy, run on demand, or run as a scheduled task at a specific time.

Nmap-supplied server and operating system data remains static until you run another Nmap scan. If you plan to scan a host using Nmap, regularly schedule scans. If a host is deleted from the network map, any Nmap scan results are discarded.

For general information about Nmap functionality, refer to the Nmap documentation at http://insecure.org.

In a multidomain deployment, the system displays Nmap remediations created in the current domain, which you can edit. It also displays Nmap remediations created in ancestor domains, which you cannot edit. To view and edit Nmap remediations in a lower domain, switch to that domain.

Before you begin

- Add an Nmap scan instance as described in Adding an Nmap Scan Instance, on page 1863.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Click **View** () next to the instance to which you want to add the remediation.

Step 3 In the Configured Remediations section, click **Add**.

Step 4 Enter a **Remediation Name**.

Step 5 Enter a **Description**.

Step 6 If you plan to use this remediation in response to a correlation rule that triggers on an intrusion event, a connection event, or a user event, configure the **Scan Which Address(es) From Event?** option.

Tip If you plan to use this remediation in response to a correlation rule that triggers on a discovery event or a host input event, by default the remediation scans the IP address of the host involved in the event; you do not need to configure this option.

Note Do **not** assign an Nmap remediation as a response to a correlation rule that triggers on a traffic profile change.

Step 7 Configure the **Scan Type** option.

Step 8 Optionally, to scan UDP ports in addition to TCP ports, choose **On** for the **Scan for UDP ports** option.

Tip A UDP portscan takes more time than a TCP portscan. To speed up your scans, leave this option disabled.

Step 9 If you plan to use this remediation in response to correlation policy violations, configure the **Use Port From Event** option.

Step 10 If you plan to use this remediation in response to correlation policy violations and want to run the scan using the appliance running the detection engine that detected the event, configure the **Scan from reporting detection engine** option.

Step 11 Configure the **Fast Port Scan** option.

Step 12 In the **Port Ranges and Scan Order** field, enter the ports you want to scan by default, using Nmap port specification syntax, in the order you want to scan those ports.

Use the following format:

- Specify values from 1 to 65535.
- Separate ports using commas or spaces.
- Use a hyphen to indicate a port range.
- When scanning for both TCP and UDP ports, preface the list of TCP ports you want to scan with a T and the list of UDP ports with a U.

Note The **Use Port From Event** option overrides this setting when the remediation is launched in response to a correlation policy violation, as described in step 8.

Example:

To scan ports 53 and 111 for UDP traffic, then scan ports 21-25 for TCP traffic, enter `U:53,111,T:21-25`.

Step 13 To probe open ports for server vendor and version information, configure **Probe open ports for vendor and version information**.

Step 14 If you choose to probe open ports, set the number of probes used by choosing a number from the **Service Version Intensity** drop-down list.

Step 15 To scan for operating system information, configure **Detect Operating System** settings.

Step 16 To determine whether host discovery occurs and whether port scans are only run against available hosts, configure **Treat All Hosts As Online**.

Step 17 To set the method you want Nmap to use when it tests for host availability, choose a method from the **Host Discovery Method** drop-down list.

Step 18 If you want to scan a custom list of ports during host discovery, enter a list of ports appropriate for the host discovery method you chose, separated by commas, in the **Host Discovery Port List** field.

Step 19 Configure the **Default NSE Scripts** option to control whether to use the default set of Nmap scripts for host discovery and server, operating system, and vulnerability discovery.

Tip See http://nmap.org/nsedoc/categories/default.html for the list of default scripts.

Step 20 To set the timing of the scan process, choose a timing template number from the **Timing Template** drop-down list.

Choose a higher number for a faster, less comprehensive scan and a lower number for a slower, more comprehensive scan.

Step 21 Click **Create**.

When the system is done creating the remediation, it displays it in edit mode.

Step 22 Click **Done** to return to the related instance.

Step 23 Click **Cancel** to return to the instance list.

Related Topics

Nmap Scan Automation, on page 184
Nmap Remediation Options, on page 1855

Editing an Nmap Remediation

Modifications you make to Nmap remediations do not affect scans in progress. The new settings take effect when the next scan starts. Delete an Nmap remediation if you no longer need it.

In a multidomain deployment, the system displays Nmap remediations created in the current domain, which you can edit. It also displays Nmap remediations created in ancestor domains, which you cannot edit. To view and edit Nmap remediations in a lower domain, switch to that domain.

Procedure

Step 1 Access the list of Nmap scan instances using either of the following methods:

- Choose **Policies** > **Actions** > **Instances**.
- Choose **Policies** > **Actions** > **Scanners**.

Step 2 Access the remediation you want to edit:

- If you accessed the list via the first method above, click **View** ( ) next to the relevant instance, and then click it again next to the remediation you want to edit in the Configured Remediations section.
- If you accessed the list via the second method above, click **View** ( ) next to the remediation you want to edit.

Step 3 Make modifications as necessary as described in Creating an Nmap Remediation, on page 1867.

Step 4 Click **Save** if you want to save your changes, or **Done** if you want to exit without saving.

Step 5 Optionally, delete the remediation by clicking **Delete** ( ) next to it.

Running an On-Demand Nmap Scan

You can launch on-demand Nmap scans whenever needed. You can specify the target for an on-demand scan by entering the IP addresses and ports you want to scan or by choosing an existing scan target.

Nmap-supplied server and operating system data remains static until you run another Nmap scan. If you plan to scan a host using Nmap, regularly schedule scans. If a host is deleted from the network map, any Nmap scan results are discarded.

Before you begin

- Optionally, add an Nmap scan target; see Adding an Nmap Scan Target, on page 1865.

Procedure

Step 1 Choose **Policies** > **Actions** > **Scanners**.

Step 2 Next to the Nmap remediation you want to use to perform the scan, click **Scan** ( ).

Step 3 Optionally, to scan using a saved scan target, choose a target from the **Saved Targets** drop-down list, and click **Load**.

Step 4 In the **IP Range(s)** field, specify the IP address for hosts you want to scan or modify the loaded list.

Note:

- For hosts with IPv4 addresses, you can specify multiple IP addresses separated by commas or use CIDR notation. You can also negate IP addresses by preceding them with an exclamation point (!).
- For hosts with IPv6 addresses, use an exact IP address. Ranges are not supported.

Step 5 In the **Ports** field, specify the ports you want to scan or modify the loaded list.

You can enter a port number, a list of ports separated by commas, or a range of port numbers separated by a dash.

Step 6 In a multidomain deployment, use the **Domain** field to specify the leaf domain where you want to perform the scan.

Step 7 Click **Scan Now**.

What to do next

- Optionally, monitor the task status; see Viewing Task Messages, on page 316.

Related Topics

Nmap Scan Automation, on page 184
Firepower System IP Address Conventions, on page 15
Ports in Searches, on page 2181

Nmap Scan Results

You can monitor Nmap scans in progress, import results from scans previously performed through the Firepower System or results preformed outside the Firepower System, and view and analyze scan results.

You can view scan results that you create using the local Nmap module as a rendered page in a pop-up window. You can also download the Nmap results file in raw XML format.

You can also view operating system and server information detected by Nmap in host profiles and in the network map. If a scan of a host produces server information for servers on filtered or closed ports, or if a scan collects information that cannot be included in the operating system information or the servers section, the host profile includes those results in an Nmap Scan Results section.

Viewing Nmap Scan Results

When an Nmap scan is complete, you can view a table of scan results.

You can manipulate the results view depending on the information you are looking for. The page you see when you access scan results differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of scan results. You can also create a custom workflow that displays only the information that matches your specific needs.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

You can download and view the Nmap results using the Nmap Version 1.01 DTD, available at http://insecure.org.

You can also clear scan results.

Procedure

Step 1 Choose **Policies** > **Actions** > **Scanners**.

Step 2 On the toolbar, click **Scan Results**.

Step 3 You have the following choices:

- Adjust the time range as described in Event Time Constraints, on page 2165.
- To use a different workflow, including a custom workflow, click **(switch workflows)** by the workflow title.
- To view the scan results as a rendered page in a pop-up window, click **View** next to the scan job.
- To save a copy of the scan results file so that you can view the raw XML code in any text editor, click **Download** next to the scan job.
- To sort scan results, click the column title. Click the column title again to reverse the sort order.
- To constrain the columns that appear, click **Close** ( ) in the column heading that you want to hide. In the pop-up window that appears, click **Apply**.

  Tip To hide or show other columns, check or clear the appropriate check boxes before you click **Apply**. To add a disabled column back to the view, Click the expand arrow to expand the search constraints, then click the column name under **Disabled Columns**.

- To drill down to the next page in the workflow, see Using Drill-Down Pages, on page 2158.
- To configure scan instances and remediations, click **Scanners** in the toolbar and see Managing Nmap Scanning, on page 1863.
- To navigate within and between workflow pages, see Workflow Page Navigation Tools, on page 2155.
- To navigate to other event views to view associated events, choose the name of the event view you want to see from the **Jump to** drop-down list.
- To search for scan results, enter your search criteria in the appropriate fields.

Related Topics

Nmap Scan Results Fields, on page 1871

Nmap Scan Results Fields

When you run an Nmap scan, the Firepower Management Center collects the scan results in a database. The following table describes the fields in the scan results table that can be viewed and searched.

Table 220: Scan Results Fields

Field	Description
Start Time	The date and time that the scan that produced the results started.
End Time	The date and time that the scan that produced the results ended.
Target	The IP address (or host name, if DNS resolution is enabled) of the scan target for the scan that produced the results.
Scan Type	Either `Nmap` or the name of the third-party scanner to indicate the type of the scan that produced the results.
Scan Mode	The mode of the scan that produced the results: • `On Demand` — results from scans run on demand. • `Imported` — results from scans on a different system and imported onto the Firepower Management Center. • `Scheduled` — results from scans run as a scheduled task.
Results	The results of the scan.
Domain	The domain of the scan target. This field is only present in a multidomain deployment.

Related Topics

Event Searches, on page 2177

Importing Nmap Scan Results

You can import XML results files created by an Nmap scan performed outside of the Firepower System. You can also import XML results files that you previously downloaded from the Firepower System. To import Nmap scan results, the results file must be in XML format and adhere to the Nmap Version 1.01 DTD. For more information on creating Nmap results and on the Nmap DTD, refer to the Nmap documentation at http://insecure.org.

A host must already exist in the network map before Nmap can append its results to the host profile.

Procedure

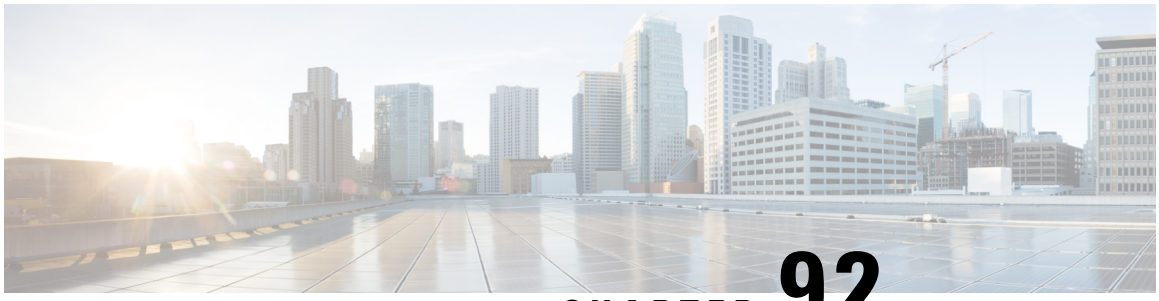
Step 1 Choose **Policies** > **Actions** > **Scanners**.

Step 2 On the toolbar, click **Import Results**.

Step 3 In a multidomain deployment, choose a leaf domain from the **Domain** drop-down list to specify where you want to store the imported results.

Step 4 Click **Browse** to navigate to the results file.

Step 5 After you return to the Import Results page, click **Import** to import the results.



CHAPTER 92

Application Detection

The following topics describe Firepower System application detection :

- Overview: Application Detection, on page 1873
- Requirements and Prerequisites for Application Detection, on page 1878
- Custom Application Detectors, on page 1878
- Viewing or Downloading Detector Details, on page 1886
- Sorting the Detector List, on page 1887
- Filtering the Detector List, on page 1887
- Navigating to Other Detector Pages, on page 1888
- Activating and Deactivating Detectors, on page 1889
- Editing Custom Application Detectors, on page 1889
- Deleting Detectors, on page 1890

Overview: Application Detection

When the Firepower System analyzes IP traffic, it attempts to identify the commonly used applications on your network. Application awareness is crucial to application control.

There are three types of applications that the system detects:

- *application protocols* such as HTTP and SSH, which represent communications between hosts
- *clients* such as web browsers and email clients, which represent software running on the host
- *web applications* such as MPEG video and Facebook, which represent the content or requested URL for HTTP traffic

The system identifies applications in your network traffic according to the characteristics specified in the detector. For example, the system can identify an application by an ASCII pattern in the packet header. In addition, Secure Socket Layers (SSL) protocol detectors use information from the secured session to identify the application from the session.

There are two sources of application detectors in the Firepower System:

- *System-provided detectors* detect web applications, clients, and application protocols.

  The availability of system-provided detectors for applications (and operating systems) depends on the version of the Firepower System and the version of the VDB you have installed. Release notes and

advisories contain information on new and updated detectors. You can also import individual detectors authored by Professional Services.

- *Custom application protocol detectors* are user-created and detect web applications, clients, and application protocols.

You can also detect application protocols through *implied application protocol detection*, which infers the existence of an application protocol based on the detection of a client.

The system identifies only those application protocols running on hosts in your monitored networks, as defined in the network discovery policy. For example, if an internal host accesses an FTP server on a remote site that you are not monitoring, the system does not identify the application protocol as FTP. On the other hand, if a remote or internal host accesses an FTP server on a host you are monitoring, the system can positively identify the application protocol.

If the system can identify the client used by a monitored host to connect to a non-monitored server, the system identifies the client's corresponding application protocol, but does not add the protocol to the network map. Note that client sessions must include a response from the server for application detection to occur.

The system characterizes each application that it detects; see Application Characteristics, on page 378. The system uses these characteristics to create groups of applications, called *application filters*. Application filters are used to perform access control and to constrain search results and data used in reports and dashboard widgets.

You can also supplement application detector data using exported NetFlow records, Nmap active scans, and the host input feature.

Related Topics

Best Practices for Configuring Application Control, on page 381
Application Detector Fundamentals, on page 1874

Application Detector Fundamentals

The Firepower System uses *application detectors* to identify the commonly used applications on your network. Use the Detectors page (**Policies** > **Application Detectors**) to view the detector list and customize detection capability.

Whether you can modify a detector or change its state (active or inactive) depends on its type. The system uses only active detectors to analyze application traffic.

Note Cisco-provided detectors may change with Firepower System and VDB updates. See the release notes and advisories for information on updated detectors.

Cisco-Provided Internal Detectors

Internal detectors are a special category of detectors for client, web application, and application protocol traffic. Internal detectors are delivered with system updates and are always on.

If an application matches against internal detectors designed to detect client-related activity and no specific client detector exists, a generic client may be reported.

Cisco-Provided Client Detectors

Client detectors detect client traffic and are delivered via VDB or system update, or are provided for import by Cisco Professional Services. You can activate and deactivate client detectors. You can export a client detector only if you import it.

Cisco-Provided Web Application Detectors

Web application detectors detect web applications in HTTP traffic payloads and are delivered via VDB or system update. Web application detectors are always on.

Cisco-Provided Application Protocol (Port) Detectors

Port-based application protocol detectors use well-known ports to identify network traffic. They are delivered via VDB or system update, or are provided for import by Cisco Professional Services. You can activate and deactivate application protocol detectors, and view a detector definition to use it as the basis for a custom detector.

Cisco-Provided Application Protocol (Firepower) Detectors

Firepower-based application protocol detectors analyze network traffic using Firepower application fingerprints and are delivered via VDB or system update. You can activate and deactivate application protocol detectors.

Custom Application Detectors

Custom application detectors are pattern-based. They detect patterns in packets from client, web application, or application protocol traffic. You have full control over imported and custom detectors.

Identification of Application Protocols in the Web Interface

The following table outlines how the Firepower System identifies detected application protocols:

Table 221: Firepower System Identification of Application Protocols

Identification	Description
application protocol name	The Firepower Management Center identifies an application protocol with its name if the application protocol was: • positively identified by the system • identified using NetFlow data and there is a port-application protocol correlation in `/etc/sf/services` • manually identified using the host input feature • identified by Nmap or another active source

Identification	Description
`pending`	The Firepower Management Center identifies an application protocol as `pending` if the system can neither positively nor negatively identify the application. Most often, the system needs to collect and analyze more connection data before it can identify a pending application. In the Application Details and Servers tables and in the host profile, the `pending` status appears only for application protocols where specific application protocol traffic was detected (rather than inferred from detected client or web application traffic).
`unknown`	The Firepower Management Center identifies an application protocol as `unknown` if: • the application does not match any of the system's detectors • the application protocol was identified using NetFlow data, but there is no port-application protocol correlation in `/etc/sf/services` • Snort has closed the session but it still persists in the device. Here, the traffic is allowed to pass through the firewall, but the application is not detected.
blank	All available detected data has been examined, but no application protocol was identified. In the Application Details and Servers tables and in the host profile, the application protocol is left blank for non-HTTP generic client traffic with no detected application protocol.

Implied Application Protocol Detection from Client Detection

If the system can identify the client used by a monitored host to access a non-monitored server, the Firepower Management Center infers that the connection is using the application protocol that corresponds with the client. (Because the system tracks applications only on monitored networks, connection logs usually do not include application protocol information for connections where a monitored host is accessing a non-monitored server.)

This process, or *implied application protocol detection*, has the following consequences:

- Because the system does not generate a New TCP Port or New UDP Port event for these servers, the server does not appear in the Servers table. In addition, you cannot trigger either discovery event alerts or correlation rules using the detection of these application protocol as a criterion.
- Because the application protocol is not associated with a host, you cannot view its details in host profiles, set its server identity, or use its information in host profile qualifications for traffic profiles or correlation rules. In addition, the system does not associate vulnerabilities with hosts based on this type of detection.

You can, however, trigger correlation events on whether the application protocol information is present in a connection. You can also use the application protocol information in connection logs to create connection trackers and traffic profiles.

Host Limits and Discovery Event Logging

When the system detects a client, server, or web application, it generates a discovery event unless the associated host has already reached its maximum number of clients, servers, or web applications.

Host profiles display up to 16 clients, 100 servers, and 100 web applications per host.

Note that actions dependent on the detection of clients, servers, or web applications are unaffected by this limit. For example, access control rules configured to trigger on a server will still log connection events.

Special Considerations for Application Detection

SFTP

In order to detect SFTP traffic, the same rule must also detect SSH.

Squid

The system positively identifies Squid server traffic when either:

- the system detects a connection from a host on your monitored network to a Squid server where proxy authentication is enabled, or
- the system detects a connection from a Squid proxy server on your monitored network to a target system (that is, the destination server where the client is requesting information or another resource).

However, the system cannot identify Squid service traffic if:

- a host on your monitored network connects to a Squid server where proxy authentication is disabled, or
- the Squid proxy server is configured to strip Via: header fields from its HTTP responses

SSL Application Detection

The system provides application detectors that can use session information from a Secure Socket Layers (SSL) session to identify the application protocol, client application, or web application in the session.

When the system detects an encrypted connection, it marks that connection as either a generic HTTPS connection or as a more specific secure protocol, such as SMTPS, when applicable. When the system detects an SSL session, it adds `SSL client` to the **Client** field in connection events for the session. If it identifies a web application for the session, the system generates discovery events for the traffic.

For SSL application traffic, managed devices can also detect the common name from the server certificate and match that against a client or web application from an SSL host pattern. When the system identifies a specific client, it replaces `SSL client` with the name of the client.

Because the SSL application traffic is encrypted, the system can use only information in the certificate for identification, not application data within the encrypted stream. For this reason, SSL host patterns can sometimes only identify the company that authored the application, so SSL applications produced by the same company may have the same identification.

In some instances, such as when an HTTPS session is launched from within an HTTP session, managed devices detect the server name from the client certificate in a client-side packet.

To enable SSL application identification, you must create access control rules that monitor responder traffic. Those rules must have either an application condition for the SSL application or URL conditions using the URL from the SSL certificate. For network discovery, the responder IP address does not have to be in the networks to monitor in the network discovery policy; the access control policy configuration determines whether the traffic is identified. To identify detections for SSL applications, you can filter by the `SSL protocol` tag, in the application detectors list or when adding application conditions in access control rules.

Referred Web Applications

Web servers sometimes refer traffic to other websites, which are often advertisement servers. To help you better understand the context for referred traffic occurring on your network, the system lists the web application that referred the traffic in the **Web Application** field in events for the referred session. The VDB contains a list of known referred sites. When the system detects traffic from one of those sites, the referring site is stored with the event for that traffic. For example, if an advertisement accessed via Facebook is actually hosted on Advertising.com, the detected Advertising.com traffic is associated with the Facebook web application. The system can also detect referring URLs in HTTP traffic, such as when a website provides a simple link to another site; in this case, the referring URL appears in the HTTP Referrer event field.

In events, if a referring application exists, it is listed as the web application for the traffic, while the URL is that for the referred site. In the example above, the web application for the connection event for that traffic would be Facebook, but the URL would be Advertising.com. A referred application may appear as the web application if no referring web application is detected, if the host refers to itself, or if there is a chain of referrals. In the dashboard, connection and byte counts for web applications include sessions where the web application is associated with traffic referred by that application.

Note that if you create a rule to act specifically on referred traffic, you should add a condition for the referred application, rather than the referring application. To block Advertising.com traffic referred from Facebook, for example, add an application condition to your access control rule for the Advertising.com application.

Requirements and Prerequisites for Application Detection

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Discovery Admin

Custom Application Detectors

If you use a custom application on your network, you can create a custom web application, client, or application protocol detector that provides the system with the information it needs to identify the application. The type of application detector is determined by your selections in the **Protocol**, **Type**, and **Direction** fields.

Client sessions must include a responder packet from the server for the system to begin detecting and identifying application protocols in server traffic. Note that, for UDP traffic, the system designates the source of the responder packet as the server.

If you have already created a detector on another Firepower Management Center, you can export it and then import it onto this Firepower Management Center. You can then edit the imported detector to suit your needs.

You can export and import custom detectors as well as detectors provided by Cisco Professional Services. However, you **cannot** export or import any other type of Cisco-provided detectors.

Custom Application Detector and User-Defined Application Fields

You can use the following fields to configure custom application detectors and user-defined applications.

Custom Application Detector Fields: General

Use the following fields to configure basic and advanced custom application detectors.

Application Protocol

The application protocol you want to detect. This can be a system-provided application or a user-defined application.

If you want the application to be available for exemption from active authentication (configured in your identity rules), you must select or create an application protocol with the **`User-Agent Exclusion`** tag.

Description

A description for the application detector.

Name

A name for the application detector.

Detector Type

The type of detector, **`Basic`** or **`Advanced`**. Basic application detectors are created in the web interface as a series of fields. Advanced application detectors are created externally and uploaded as custom .lua files.

Custom Application Detector Fields: Detection Patterns

Use the following fields to configure the detection patterns for basic custom application detectors.

Direction

The source of the traffic the detector should inspect, **`Client`** or **`Server`**.

Offset

The location in a packet, in bytes from the beginning of the packet payload, where the system should begin searching for the pattern.

Because packet payloads start at byte 0, calculate the offset by subtracting 1 from the number of bytes you want to move forward from the beginning of the packet payload. For example, to look for the pattern in the fifth bit of the packet, type `4` in the **Offset** field.

Pattern

The pattern string associated with the **Type** you selected.

Ports

The port of the traffic the detector should inspect.

Protocol

The protocol you want to detect. Your protocol selection determines whether the **Type** or the **URL** field displays.

The protocol (and, in some cases, your subsequent selections in the **Type** and **Direction** fields) determine the type of application detector you create: web application, client, or application protocol.

Detector Type	Protocol	Type or Direction
Web Application	HTTP	**Type** is **`Content Type`** or **`URL`**
	RTMP	Any
	SSL	Any
Client	HTTP	**Type** is **`User Agent`**
	SIP	Any
	TCP or UDP	**Direction** is **`Client`**
Application Protocol	TCP or UDP	**Direction** is **`Server`**

Type

The type of pattern string you entered. The options you see are determined by the **Protocol** you selected. If you selected **`RTMP`** as the protocol, the **URL** field displays instead of the **Type** field.



Note If you select **`User Agent`** as the **Type**, the system automatically sets the **Tag** for the application to **`User-Agent Exclusion`**.

Type Selection	String Characteristics
`Ascii`	The string is ASCII encoded.
`Common Name`	The string is the value in the commonName field within the server response message.
`Content Type`	The string is the value in the content-type field within the server response header.
`Hex`	The string is in hexadecimal notation.
`Organizational Unit`	The string is the value in the organizationName field within the server response message.
`SIP Server`	The string is the value in the From field within the message header.
`SSL Host`	The string is the value in the server_name field within the ClientHello message.

Type Selection	String Characteristics
URL	The string is a URL. **Note** The detector assumes that the string you enter is a complete section of the URL. For example, entering **cisco.com** would match **www.cisco.com/support** and **www.cisco.com**, but not **www.wearecisco.com**.
User Agent	The string is the value in the user-agent field within the GET request header. It is also available for the SIP protocol and indicates that the string is the value in the User-Agent field within the SIP message header.

URL

Either a full URL or a section of a URL from the swfURL field within the C2 message of a RTMP packet. This field displays instead of the **Type** field when you select **RTMP** as the **Protocol**.



Note The detector assumes that the string you enter is a complete section of the URL. For example, entering **cisco.com** would match **www.cisco.com/support** and **www.cisco.com**, but not **www.wearecisco.com**.

User-Defined Application Fields

Use the following fields to configure user-defined applications within basic and advanced custom application detectors.

Business Relevance

The likelihood that the application is used within the context of your organization's business operations, as opposed to recreationally: **Very High**, **High**, **Medium**, **Low**, or **Very Low**. Select the option that best describes the application.

Categories

A general classification for the application that describes its most essential function.

Description

A description for the application.

Name

A name for the application.

Risk

The likelihood that the application is used for purposes that might be against your organization's security policy: **Very High**, **High**, **Medium**, **Low**, or **Very Low**. Select the option that best describes the application.

Tags

One or more predefined tags that provide additional information about the application. If you want an application to be available for exemption from active authentication (configured in your identity rules), you must add the **`User-Agent Exclusion`** tag to your application.

Configuring Custom Application Detectors

You can configure basic or advanced custom application detectors.

Procedure

Step 1 Select **Policies** > **Application Detectors**.

Step 2 Click **Create Custom Detector**.

Step 3 Enter a **Name** and a **Description**.

Step 4 Select an **Application Protocol**. You have the following options:

- If you are creating a detector for an existing application protocol (for example, if you want to detect a particular application protocol on a non-standard port), select the application protocol from the drop-down list.
- If you are creating a detector for a user-defined application, follow the procedure outlined in Creating a User-Defined Application, on page 1883.

Step 5 Select a **Detector Type**.

Step 6 Click **OK**.

Step 7 Configure **Detection Patterns** or **Detection Criteria**:

- If you are configuring a basic detector, specify preset **Detection Patterns** as described in Specifying Detection Patterns in Basic Detectors, on page 1883.
- If you are configuring an advanced detector, specify custom **Detection Criteria** as described in Specifying Detection Criteria in Advanced Detectors, on page 1884.

Caution Advanced custom detectors are complex and require outside knowledge to construct valid .lua files. Incorrectly configured detectors could have a negative impact on performance or detection capability.

Step 8 Optionally, use **Packet Captures** to test the new detector as described in Testing a Custom Application Protocol Detector, on page 1885.

Step 9 Click **Save**.

Note If you include the application in an access control rule, the detector is automatically activated and cannot be deactivated while in use.

What to do next

- Activate the detector as described in Activating and Deactivating Detectors, on page 1889.

Related Topics

Custom Application Detector and User-Defined Application Fields, on page 1879

Creating a User-Defined Application

Applications, categories, and tags created here are available in access control rules and in the application filter object manager as well.

Caution Creating a user-defined application immediately restarts the Snort process without going through the deploy process. Restarting the Snort process temporarily interrupts traffic inspection on all managed devices. Whether traffic drops during this interruption or passes without further inspection depends on the model of the managed device and how it handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Before you begin

- Begin configuring your custom application protocol detector as described in Configuring Custom Application Detectors, on page 1882.

Procedure

Step 1 On the Create Detector page, click **Add**.

Step 2 Type a **Name**.

Step 3 Type a **Description**.

Step 4 Select a **Business Relevance**.

Step 5 Select a **Risk**.

Step 6 Click **Add** next to Categories to add a category and type a new category name, or select an existing category from the **Categories** drop-down list.

Step 7 Optionally, click **Add** next to Tags to add a tag and type a new tag name, or select an existing tag from the **Tags** drop-down list.

Step 8 Click **OK**.

What to do next

- Continue configuring your custom application protocol detector as described in Configuring Custom Application Detectors, on page 1882. You must save and activate the detector before the system can use it to analyze traffic.

Related Topics

Custom Application Detector and User-Defined Application Fields, on page 1879

Specifying Detection Patterns in Basic Detectors

You can configure a custom application protocol detector to search application protocol packet headers for a particular pattern string. You can also configure detectors to search for multiple patterns; in that case the

application protocol traffic must match all of the patterns for the detector to positively identify the application protocol.

Application protocol detectors can search for ASCII or hexadecimal patterns, using any offset.

Before you begin

- Begin configuring your custom application protocol detector as described in Configuring Custom Application Detectors, on page 1882.

Procedure

Step 1 On the Create Detector page, in the Detection Patterns section, click **Add**.

Step 2 Select a **Protocol** for traffic the detector should inspect.

Step 3 Specify the pattern **Type** you want to detect.

Step 4 Type a **Pattern String** that matches the **Type** you specified.

Step 5 Optionally, type the **Offset** (in bytes).

Step 6 Optionally, to identify application protocol traffic based on the port it uses, type a port from 1 to 65535 in the **Port(s)** field. To use multiple ports, separate them by commas.

Step 7 Optionally, select a **Direction**: `Client` or `Server`.

Step 8 Click **OK**.

Tip If you want to delete a pattern, click **Delete** ( ) next to the pattern you want to delete.

What to do next

- Continue configuring your custom application protocol detector as described in Configuring Custom Application Detectors, on page 1882. You must save and activate the detector before the system can use it to analyze traffic.

Related Topics

Specifying Detection Criteria in Advanced Detectors, on page 1884

Specifying Detection Criteria in Advanced Detectors

Caution Advanced custom detectors are complex and require outside knowledge to construct valid .lua files. Incorrectly configured detectors could have a negative impact on performance or detection capability.

Caution Do not upload .lua files from untrusted sources.

Custom .lua files contain your custom application detector settings. Creating custom .lua files requires advanced knowledge of the lua programming language and experience with Cisco's C-lua API. Cisco strongly recommends you use the following to prepare .lua files:

- third-party instruction and reference material for the lua programming language
- The Open Source Detectors Developers Guide: https://www.snort.org/downloads
- OpenAppID Snort community resources: http://blog.snort.org/search/label/openappid



Note The system does not support .lua files that reference system calls or file I/O.

Before you begin

- Begin configuring your custom application protocol detector as described in Configuring Custom Application Detectors, on page 1882.
- Prepare to create a valid .lua file by downloading and studying the .lua files for comparable detectors. For more information about downloading detector files, see Viewing or Downloading Detector Details, on page 1886.
- Create a valid .lua file that contains your custom application detector settings.

Procedure

Step 1 On the Create Detector page for an advanced custom application detector, in the Detection Criteria section, click **Add**.

Step 2 Click **Browse...** to navigate to the `.lua` file and upload it.

Step 3 Click **OK**.

What to do next

- Continue configuring your custom application protocol detector as described in Configuring Custom Application Detectors, on page 1882. You must save and activate the detector before the system can use it to analyze traffic.

Related Topics

Specifying Detection Patterns in Basic Detectors, on page 1883

Testing a Custom Application Protocol Detector

If you have a packet capture (pcap) file that contains packets with traffic from the application protocol you want to detect, you can test a custom application protocol detector against that pcap file. Cisco recommends using a simple, clean pcap file without unnecessary traffic.

Pcap files must be 256 KB or smaller; if you try to test your detector against a larger pcap file, the Firepower Management Center automatically truncates it and tests the incomplete file. You must fix the unresolved checksums in a pcap before using the file to test a detector.

Before you begin

- Configure your custom application protocol detector as described in Configuring Custom Application Detectors, on page 1882.

Procedure

Step 1 On the Create Detector page, in the Packet Captures section, click **Add**.

Step 2 Browse to the pcap file in the pop-up window and click **OK**.

Step 3 To test your detector against the contents of the pcap file, click evaluate next to the pcap file. A message indicates whether the test succeeded.

Step 4 Optionally, repeat steps 1 to 3 to test the detector against additional pcap files.

Tip To delete a pcap file, click **Delete** ( ) next to the file you want to delete.

What to do next

- Continue configuring your custom application protocol detector as described in Configuring Custom Application Detectors, on page 1882. You must save and activate the detector before the system can use it to analyze traffic.

Viewing or Downloading Detector Details

You can use the detectors list to view application detector details (all detectors) and download detector details (custom application detectors only).

Procedure

Step 1 To view application detector details, do one of the following:

- See the *Cisco Firepower Application Detector Reference* for the relevant VDB version at https://www.cisco.com/c/en/us/support/security/defense-center/products-technical-reference-list.html
- a. Select **Policies** > **Application Detectors**.

  b. Filter the list to find a particular detector.

  c. Click **Information** ( )

Step 2 To download detector details for a custom application detector, click **Download** ( ).

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have the necessary permissions.

Sorting the Detector List

By default, the Detectors page lists detectors alphabetically by name. An up or down arrow next to a column heading indicates that the page is sorted by that column in that direction.

Procedure

Step 1 Select **Policies** > **Application Detectors**.

Step 2 Click the appropriate column heading.

Filtering the Detector List

Procedure

Step 1 Select **Policies** > **Application Detectors**.

Step 2 Expand one of the filter groups described in Filter Groups for the Detector List, on page 1887 and select the check box next to a filter. To select all filters in a group, right-click the group name and select **Check All**.

Step 3 If you want to remove a filter, click **Remove** ( ) in the name of the filter in the **Filters** field or disable the filter in the filter list. To remove all filters in a group, right-click the group name and select **Uncheck All**.

Step 4 If you want to remove all filters, click **Clear all** next to the list of filters applied to the detectors.

Filter Groups for the Detector List

You can use several filter groups, separately or in combination, to filter the list of detectors.

Name

Finds detectors with names or descriptions containing the string you type. Strings can contain any alphanumeric or special character.

Custom Filter

Finds detectors matching a custom application filter created on the object management page.

Author

Finds detectors according to who created the detector. You can filter detectors by:

- any individual user who has created or imported a custom detector
- Cisco, which represents all Cisco-provided detectors *except* individually imported add-on detectors (you are the author for any detector that you import)

- **Any User**, which represents all detectors not provided by Cisco

State

Finds detectors according to their state, that is, **Active** or **Inactive**.

Type

Finds detectors according to the detector type, as described in Application Detector Fundamentals, on page 1874.

Protocol

Finds detectors according to which traffic protocol the detector inspects.

Category

Finds detectors according to the categories assigned to the application they detect.

Tag

Finds detectors according to the tags assigned to the application they detect.

Risk

Finds detectors according to the risks assigned to the application they detect: **Very High**, **High**, **Medium**, **Low**, and **Very Low**.

Business Relevance

Finds detectors according to the business relevance assigned to the application they detect: **Very High**, **High**, **Medium**, **Low**, and **Very Low**.

Navigating to Other Detector Pages

Procedure

Step 1 Select **Policies** > **Application Detectors**.

Step 2 If you want to view the next page, click **Right Arrow** ().

Step 3 If you want to view the previous page, click **Left Arrow** ().

Step 4 If you want to view a different page, type the page number and press Enter.

Step 5 If you want to jump to the last page, click **Right End Arrow** ().

Step 6 If you want to jump to the first page, click **Left End Arrow** ().

Activating and Deactivating Detectors

You must activate a detector before you can use it to analyze network traffic. By default, all Cisco-provided detectors are activated.

You can activate multiple application detectors for each port to supplement the system's detection capability.

When you include an application in an access control rule in a policy and that policy is deployed, if there is no active detector for that application, one or more detectors automatically activate. Similarly, while an application is in use in a deployed policy, you cannot deactivate a detector if deactivating leaves no active detectors for that application.

Tip For improved performance, deactivate any application protocol, client, or web application detectors you do not intend to use.

Caution Activating or deactivating a system or custom application detector immediately restarts the Snort process without going through the deploy process. Restarting the Snort process temporarily interrupts traffic inspection on all managed devices. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Procedure

Step 1 Select **Policies** > **Application Detectors**.

Step 2 Click the slider next to the detector you want to activate or deactivate. If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Note Some application detectors are required by other detectors. If you deactivate one of these detectors, a warning appears to indicate that the detectors that depend on it are also disabled.

Editing Custom Application Detectors

Use the following procedure to modify custom application detectors.

Procedure

Step 1 Select **Policies** > **Application Detectors**.

Step 2 Click **Edit** ( ) next to the detector you want to modify. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Make changes to the detector as described in Configuring Custom Application Detectors, on page 1882.

Step 4 You have the following saving options, depending on the state of the detector:

- To save an inactive detector, click **Save**.
- To save an inactive detector as a new, inactive detector, click **Save as New**.
- To save an active detector and immediately start using it, click **Save and Reactivate**.

  Caution Saving and reactivating a custom application detector immediately restarts the Snort process without going through the deploy process. Restarting the Snort process temporarily interrupts traffic inspection on all managed devices. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

- To save an active detector as a new, inactive detector, click **Save as New**.

Deleting Detectors

You can delete custom detectors as well as individually imported add-on detectors provided by Cisco Professional Services. You cannot delete any of the other Cisco-provided detectors, though you can deactivate many of them.

Note While a detector is in use in a deployed policy, you cannot delete the detector.

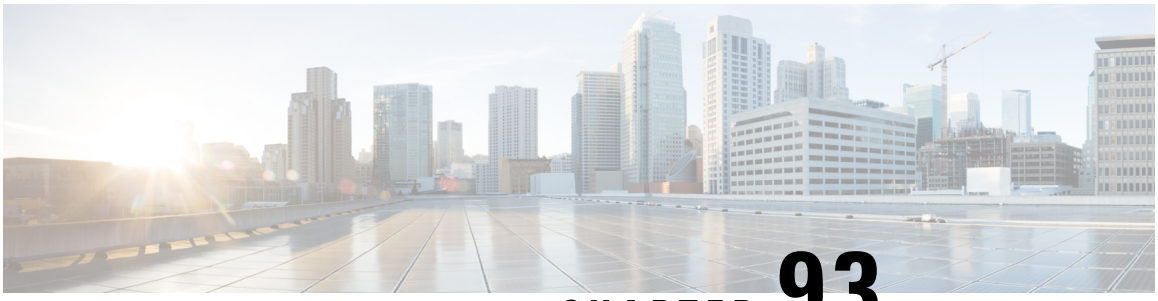
Caution Deleting an activated custom application detector immediately restarts the Snort process without going through the deploy process. Restarting the Snort process temporarily interrupts traffic inspection on all managed devices. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Procedure

Step 1 Select **Policies** > **Application Detectors**.

Step 2 Click **Delete** ( ) next to the detector you want to delete. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **OK**.



CHAPTER 93

User Identity Sources

The following topics describe Firepower System user *identity sources*, which are sources for *user awareness*. These users can be controlled with identity and access control policies:

- About User Identity Sources, on page 1891
- The User Agent Identity Source, on page 1893
- The ISE Identity Source, on page 1895
- The Terminal Services (TS) Agent Identity Source, on page 1901
- The Captive Portal Identity Source, on page 1903
- The Traffic-Based Detection Identity Source, on page 1915

About User Identity Sources

The following table provides a brief overview of the user identity sources supported by the Firepower System. Each identity source provides a store of users for user awareness. These users can then be controlled with identity and access control policies.

User Identity Source	Policy	Server Requirements	Type	Authentication Type	User Awareness?	User Control?	For more information, see...
User Agent	Identity	Microsoft Active Directory	Authoritative logins	Passive	Yes	Yes	The User Agent Identity Source, on page 1893
ISE	Identity	Microsoft Active Directory	Authoritative logins	Passive	Yes	Yes	The ISE Identity Source, on page 1895

User Identity Source	Policy	Server Requirements	Type	Authentication Type	User Awareness?	User Control?	For more information, see...
TS Agent	Identity	Microsoft Windows Terminal Server	Authoritative logins	Passive	Yes	Yes	The Terminal Services (TS) Agent Identity Source, on page 1901
Captive portal	Identity	Microsoft Active Directory	Authoritative logins	Active	Yes	Yes	The Captive Portal Identity Source, on page 1903
Identity	RADIUS	Authoritative logins	Active	Yes	No	No	The Traffic-Based Detection Identity Source, on page 1915
Traffic-based detection	Network discovery	n/a	Non-authoritative logins	n/a	Yes		

Consider the following when selecting identity sources to deploy:

- You must use traffic-based detection for non-LDAP user logins. For example, if you are using only user agents to detect user activity, restricting non-LDAP logins has no effect.
- You must use traffic-based detection or captive portal to record failed login or authentication activity. A failed login or authentication attempt does not add a new user to the list of users in the database.
- The captive portal identity source requires a managed device with a routed interface. You *cannot* use an inline (also referred to as tap mode) interface with captive portal.

Data from those identity sources is stored in the Firepower Management Center's users database and the user activity database. You can configure Firepower Management Center-server user downloads to automatically and regularly download new user data to your databases.

After you configure identity rules using the desired identity source, you must associate each rule with an access control policy and deploy the policy to managed devices for the policy to have any effect. For more information about access control policies and deployment, see User, Realm, and ISE Attribute Conditions (User Control), on page 384.

For general information about user identity in the Firepower System, see About User Identity, on page 1827.

Video icon YouTube videos for configuring identity sources.

The User Agent Identity Source

The Cisco Firepower User Agent is a passive authentication method; it is an authoritative identity source, meaning user information is supplied by a trusted Active Directory server. When integrated with the Firepower System, the user agent monitors users when they log in and out of hosts with Active Directory credentials. The data gained from the User Agent can be used for user awareness and user control.

The user agent associates each user with an IP address, which allows access control rules with user conditions to trigger. You can use one user agent to monitor user activity on up to five Active Directory servers and send encrypted data to up to five Firepower Management Centers.

The User Agent does not report failed login attempts.

Video User agent setup video on YouTube.

User Agent Guidelines

The User Agent requires a multi-step configuration that includes the following:

- At least one computer with the user agent installed.
- Connections between a Firepower Management Center and the computers or Active Directory servers with the user agent installed.
- An identity realm configured in each Firepower Management Center that receives user data from a user agent.

For detailed information about the multi-step User Agent configuration and a complete discussion of the server requirements, see the *Cisco Firepower User Agent Configuration Guide*.

Note Make sure the time on your computer or Active Directory server is synchronized with the time on the Firepower Management Center. If the appliances are not synchronized, the system might perform user timeouts at unexpected intervals.

The Firepower Management Center connection not only allows you to retrieve metadata for the users whose logins and logoffs were detected by User Agents, but also is used to specify the users and groups you want to use in access control rules. If the user agent is configured to exclude specific user names, login data for those user names are not reported to the Firepower Management Center. User agent data is stored in the user database and user activity database on the Firepower Management Center.

Note User Agents cannot transmit Active Directory user names ending with the `$` character to the Firepower Management Center. You must remove the final `$` character if you want to monitor these users.

If multiple users are logged into a host using remote sessions, the agent might not detect logins from that host properly. For information about how to prevent this, see the *Cisco Firepower User Agent Configuration Guide*.

Configure the User Agent for User Control

For more information about the User Agent, see The User Agent Identity Source, on page 1893.

Before you begin

- Configure and enable an Active Directory realm for your User Agent connection as described in Create a Realm, on page 1949.

Procedure

Step 1 Log in to the Firepower Management Center.

Step 2 Click **System** > **Integration**.

Step 3 Click **Identity Sources**.

Step 4 Click **User Agent** for the **Service Type** to enable the User Agent connection.

Note To disable the connection, click **None**.

Step 5 Click **New Agent** to add a new agent.

Step 6 Enter the **Hostname** or **Address** of the computer where you plan to install the agent. You must use an IPv4 address; you cannot configure the Firepower Management Center to connect to a User Agent using an IPv6 address.

Step 7 Click **Add**.

Step 8 To delete a connection, click **Delete** () and confirm that you want to delete it.

What to do next

- Continue User Agent setup as described in the *Cisco Firepower User Agent Configuration Guide*.
- Configure an identity rule as described in Create an Identity Rule, on page 1962.
- Associate the identity policy with an access control policy as discussed in Associating Other Policies with Access Control, on page 1213.
- Deploy your identity and access control policies to managed devices as discussed in Deploy Configuration Changes, on page 346.

Related Topics

Troubleshoot the User Agent Identity Source, on page 1894
Access Control Policies, on page 1201

Troubleshoot the User Agent Identity Source

If you experience issues with the User Agent connection, see the *Cisco Firepower User Agent Configuration Guide*.

For related troubleshooting information in this guide, see Troubleshoot Realms and User Downloads, on page 1945 and Troubleshoot User Control, on page 387.

If you experience issues with user data reported by the User Agent, note:

- After the system detects activity from a User Agent user whose data is not yet in the database, the system retrieves information about them from the server. That user's activity is not handled by rules, and is not displayed in the web interface until the system successfully retrieves information about them in a user download.
- If you have Firepower Management Center high availability configured and the primary fails, all logins reported by a User Agent cannot be identified during failover downtime, even if the users were previously seen and downloaded to the Firepower Management Center. The unidentified users are logged as Unknown users on the Firepower Management Center. After the downtime, the Unknown users are re identified and processed according to the rules in your identity policy.
- If the User Agent monitors the same users as the TS Agent, the system prioritizes the TS Agent data. If the TS Agent and the User Agent report identical activity from the same IP address, only the TS Agent data is logged.
- Active FTP sessions are displayed as the **Unknown** user in events. This is normal because, in active FTP, the server (not the client) initiates the connection and the FTP server should not have an associated user name. For more information about active FTP, see RFC 959.

The ISE Identity Source

You can integrate your Cisco Identity Services Engine (ISE) deployment with the Firepower System to use ISE for passive authentication.

ISE is an authoritative identity source, and provides user awareness data for users who authenticate using Active Directory (AD), LDAP, RADIUS, or RSA. Additionally, you can perform user control on Active Directory users. ISE does not report failed login attempts or the activity of ISE Guest Services users.

Note The Firepower System does not parse IEEE 802.1x machine authentication but it *does* parse 802.1x user authentication. If you are using 802.1x with ISE, you must include user authentication. 802.1x machine authentication will not provide a user identity to the FMC that can be used in policy.

For more information on Cisco ISE, see the *Cisco Identity Services Engine Administrator Guide*.

Note We strongly recommend you use the latest version of ISE to get the latest feature set and the most number of issue fixes.

How to Configure ISE for User Control

You can use ISE in any of the following configurations:

- With a realm, identity policy, and associated access control policy.

Use a realm to control *user* access to network resources in policy. You can still use ISE Security Group Tags (SGT) metadata in your policies.

- With an access control policy only. No realm or identity policy are necessary.

  Use this method to control network access using SGT metadata alone.

ISE Guidelines and Limitations

Use the guidelines discussed in this section when configuring ISE with the Firepower System.

ISE Version and Configuration Compatibility

Your ISE version and configuration affects its integration and interaction with Firepower, as follows:

- We strongly recommend you use the latest version of ISE to get the latest feature set.
- Synchronize the time on the ISE server and the Firepower Management Center. Otherwise, the system might perform user timeouts at unexpected intervals.
- To implement user control using ISE data, configure and enable a realm for the ISE server assuming the pxGrid persona as described in Create a Realm, on page 1949.
- Each Firepower Management Center host name that connects to an ISE server must be unique; otherwise, the connection to one of the Firepower Management Centers will be dropped.
- Version 1.3 of ISE does not include support for IPv6-enabled endpoints. With this version of ISE, you cannot gather user identity data or perform remediations on IPv6-enabled endpoints.
- If ISE Endpoint Protection Service (EPS) is enabled and configured in your ISE deployment, you can use your ISE connection to run ISE EPS remediations on the source or destination host involved in a correlation policy violation.
- If you configured your ISE deployment to update a user's SGT after the user's EPSStatus changes, your ISE EPS remediations also update the SGT on the Firepower Management Center.

For the specific versions of ISE that are compatible with this version of the system, see the *Cisco Firepower Compatibility Guide*.

IPv6 support

Version 2.0 (patch 4) and later of ISE includes support for IPv6-enabled endpoints.

Approve clients in ISE

Before a connection between the ISE server and the Firepower Management Center succeeds, you must manually approve the clients in ISE. (Typically, there are two clients: one for the connection test and another for ISE agent.)

You can also enable **Automatically approve new accounts** in ISE as discussed in the chapter on Managing users and external identity sources in the *Cisco Identity Services Engine Administrator Guide*.

Security Group Tags (SGT)

A Security Group Tag (SGT) specifies the privileges of a traffic source within a trusted network. Cisco ISE and Cisco TrustSec use a feature called Security Group Access (SGA) to apply SGT attributes to packets as they enter the network. These SGTs correspond to a user's assigned security group within ISE

or TrustSec. If you configure ISE as an identity source, the Firepower System can use these SGTs to filter traffic.

Note To implement user control using only the ISE SGT attribute tag, you do not need to configure a realm for the ISE server. ISE SGT attribute conditions can be configured in policies with or without an associated identity policy. For more information, see Configuring ISE Attribute Conditions, on page 387.

Note In some rules, custom SGT conditions can match traffic tagged with SGT attributes that were *not* assigned by ISE. This is not considered user control, and works only if you are not using ISE as an identity source; see Custom SGT Conditions, on page 389.

ISE and High Availability

When the primary Firepower Management Center fails, the following occur:

- Until the standby is promoted to primary, the user database on the secondary Firepower Management Center is read-only.

  Users added to the repository (for example, Active Directory) are not downloaded to the Firepower Management Center and those users are identified as Unknown.

  New SGTs are not used.

- After the standby is promoted to primary, all operations return to normal; that is, users are downloaded, new SGTs are used, and users are identified if possible.

When the ISE primary server fails, you must manually promote the secondary to primary; there is no automatic failover.

Endpoint Location (or Location IP)

An Endpoint Location attribute is the IP address of the network device that used ISE to authenticate the user, as identified by ISE.

You must configure and deploy an identity policy to control traffic based on **Endpoint Location (Location IP)**.

ISE Attributes

Configuring an ISE connection populates the Firepower Management Center database with ISE attribute data. You can use the following ISE attributes for user awareness and user control.

Endpoint Profile (or Device Type)

An Endpoint Profile attribute is the user's endpoint device type, as identified by ISE.

You must configure and deploy an identity policy to control traffic based on **Endpoint Profile (Device Type)**.

Configure ISE for User Control

The following procedure discusses how to configure the ISE identity source. You must be in the global domain to perform this task.

Before you begin

- To get user sessions from a Microsoft Active Directory Server or supported LDAP server, configure and enable a realm for the ISE server, assuming the pxGrid persona, as discussed in Create a Realm, on page 1949.

Procedure

Step 1 Log in to the Firepower Management Center.

Step 2 Click **System** > **Integration**.

Step 3 Click **Identity Sources**.

Step 4 Click **Identity Services Engine** for the **Service Type** to enable the ISE connection.

Note To disable the connection, click **None**.

Step 5 Enter a **Primary Host Name/IP Address** and, optionally, a **Secondary Host Name/IP Address**.

Step 6 Click the appropriate certificate authorities from the **pxGrid Server CA** and **MNT Server CA** lists, and the appropriate certificate from the **FMC Server Certificate** list. You can also click **Add** ( ) to add a certificate.

Note The **FMC Server Certificate** must include the **clientAuth** extended key usage value, or it must not include any extended key usage values.

Step 7 (Optional.) Enter an **ISE Network Filter** using CIDR block notation.

Step 8 To test the connection, click **Test**.

If the test fails, click **Additional Logs** for more information about the connection failure.

Note When you run two ISE pxGrid 1.0 nodes, it is normal for one host to show Success and one to show Failure. Because pxGrid 1.0 only runs actively on one ISE node at a time, the likelihood of success depends on which node in ISE is the active pxGrid node.

What to do next

- Specify users to control and other options using an identity policy as described in Create an Identity Policy, on page 1961.
- Associate the identity rule with an access control policy, which filters and optionally inspects traffic, as discussed in Associating Other Policies with Access Control, on page 1213.
- Deploy your identity and access control policies to managed devices as discussed in Deploy Configuration Changes, on page 346.
- Monitor user activity as discussed in Using Workflows, on page 2150.

Related Topics

Troubleshoot the Captive Portal Identity Source, on page 1914
Trusted Certificate Authority Objects, on page 452
Internal Certificate Objects, on page 455

ISE Configuration Fields

The following fields are used to configure a connection to ISE.

Primary and Secondary Host Name/IP Address

The hostname or IP address for the primary and, optionally, the secondary pxGrid ISE servers.

The ports used by the host names you specify must be reachable by both ISE and the Firepower Management Center.

pxGrid Server CA

The certificate authority for the pxGrid framework. If your deployment includes a primary and a secondary pxGrid node, the certificates for both nodes must be signed by the same certificate authority.

MNT Server CA

The certificate authority for the ISE certificate when performing bulk downloads. If your deployment includes a primary and a secondary MNT node, the certificates for both nodes must be signed by the same certificate authority.

FMC Server Certificate

The certificate and key that the Firepower Management Center must provide to ISE to connect to ISE or to perform bulk downloads.

Note The **FMC Server Certificate** must include the clientAuth extended key usage value, or it must not include any extended key usage values.

ISE Network Filter

An optional filter you can set to restrict the data that ISE reports to the Firepower Management Center. If you provide a network filter, ISE reports data from the networks within that filter. You can specify a filter in the following ways:

- Leave the field blank to specify **`any`**.
- Enter a single IPv4 address block using CIDR notation.
- Enter a list of IPv4 address blocks using CIDR notation, separated by commas.

Note This version of the Firepower System does not support filtering using IPv6 addresses, regardless of your ISE version.

Related Topics

Trusted Certificate Authority Objects, on page 452
Internal Certificate Objects, on page 455

Troubleshoot ISE or Cisco TrustSec Issues

Troubleshoot Cisco TrustSec issues

A device interface can be configured to propagate Security Group Tags (SGTs) either from ISE or from a Cisco device on the network (referred to as Cisco TrustSec.) On the device management page (**Devices** > **Device Management**), the **Propagate Security Group Tag** check box for an interface is checked after a device reboot. If you do not want the interface to propagate TrustSec data, uncheck the box.

FMC health monitor issue

The ISE Connection Status Monitor (health monitor) displays `check connectivity error` if ISE uses pxgrid v1 even though there is nothing wrong with the connection.

Troubleshoot ISE issues

For other related troubleshooting information, see Troubleshoot Realms and User Downloads, on page 1945 and Troubleshoot User Control, on page 387.

If you experience issues with the ISE connection, check the following:

- The pxGrid Identity Mapping feature in ISE must be enabled before you can successfully integrate ISE with the Firepower System.
- When the primary server fails, you must manually promote the secondary to primary; there is no automatic failover.
- Before a connection between the ISE server and the Firepower Management Center succeeds, you must manually approve the clients in ISE. (Typically, there are two clients: one for the connection test and another for ISE agent.)

  You can also enable **Automatically approve new accounts** in ISE as discussed in the chapter on Managing users and external identity sources in the *Cisco Identity Services Engine Administrator Guide*.
- The **FMC Server Certificate** must include the **clientAuth** extended key usage value, or it must not include any extended key usage values.
- The time on your ISE server must be synchronized with the time on the Firepower Management Center. If the appliances are not synchronized, the system may perform user timeouts at unexpected intervals.
- If your deployment includes a primary and a secondary pxGrid node,
  - The certificates for both nodes must be signed by the same certificate authority.
  - The ports used by the host name must be reachable by both the ISE server and by the Firepower Management Center.
- If your deployment includes a primary and a secondary MNT node, the certificates for both nodes must be signed by the same certificate authority.

To exclude subnets from receiving user-to-IP and Security Group Tag (SGT)-to-IP mappings from ISE, use the **configure identity-subnet-filter** {**add** | **remove**} command. You should typically do this for lower-memory managed devices to prevent Snort identity health monitor memory errors.

If you experience issues with user data reported by ISE, note the following:

- After the system detects activity from an ISE user whose data is not yet in the database, the system retrieves information about them from the server. Activity seen by the ISE user is *not* handled by access control rules, and is *not* displayed in the web interface until the system successfully retrieves information about them in a user download.
- You cannot perform user control on ISE users who were authenticated by an LDAP, RADIUS, or RSA domain controller.
- The Firepower Management Center does not receive user data for ISE Guest Services users.
- If ISE monitors the same users as TS Agent, the Firepower Management Center prioritizes the TS Agent data. If the TS Agent and ISE report identical activity from the same IP address, only the TS Agent data is logged to the Firepower Management Center.
- Your ISE version and configuration impact how you can use ISE in the Firepower System. For more information, see The ISE Identity Source, on page 1895.
- If you have Firepower Management Center high availability configured and the primary fails, see the section on ISE and High Availability in ISE Guidelines and Limitations, on page 1896.
- Active FTP sessions are displayed as the **Unknown** user in events. This is normal because, in active FTP, the server (not the client) initiates the connection and the FTP server should not have an associated user name. For more information about active FTP, see RFC 959.

If you experience issues with supported functionality, see The ISE Identity Source, on page 1895 for more information about version compatibility.

Troubleshoot ISE user timeout

If you're setting up ISE/ISE-PIC without a realm, be aware there is a user session timeout that affects how users are seen by the Firepower Management Center. For more information, see Realm Fields, on page 1950.

The Terminal Services (TS) Agent Identity Source

The TS Agent is a passive authentication method and one of the authoritative identity sources supported by the Firepower System. A Windows Terminal Server performs the authentication, and the TS Agent reports it to a standalone or high availability Firepower Management Center.

When installed on Windows Terminal Servers, the TS Agent assigns a unique port range to individual users as they log in or log out of a monitored network. The Firepower Management Center uses the unique port to identify individual users in the Firepower System. You can use one TS Agent to monitor user activity on one Windows Terminal Server and send encrypted data to a Firepower Management Center.

The TS Agent does not report failed login attempts. The data gained from the TS Agent can be used for user awareness and user control.

video TS Agent setup video on YouTube.

TS Agent Guidelines

The TS Agent requires a multi-step configuration, and includes the following:

1. A Windows Terminal Server with the TS Agent installed and configured.

2. One or more identity realms targeting the users your server is monitoring.

You install the TS Agent on a Microsoft Windows Terminal Server. For detailed information about the multi-step TS Agent installation and configuration and a complete discussion of the server and Firepower System requirements, see the *Cisco Terminal Services (TS) Agent Guide*.

TS Agent data is visible in the Users, User Activity, and Connection Event tables and can be used for user awareness and user control.



Note If the TS Agent monitors the same users as another passive authentication identity source (the user agent or ISE), the Firepower Management Center prioritizes the TS Agent data. If the TS Agent and another passive identity source report activity by the same IP address, only the TS Agent data is logged to the Firepower Management Center.

Configure the TS Agent for User Control

To use the TS Agent as an identity source for user awareness and user control, install and configure the TS Agent software as discussed in the *Cisco Terminal Services (TS) Agent Guide*.

What to do next:

- Specify users to control and other options using an identity policy as described in Create an Identity Policy, on page 1961.
- Associate the identity rule with an access control policy, which filters and optionally inspects traffic, as discussed in Associating Other Policies with Access Control, on page 1213.
- Deploy your identity and access control policies to managed devices as discussed in Deploy Configuration Changes, on page 346.
- Monitor user activity as discussed in Using Workflows, on page 2150.

Troubleshoot the TS Agent Identity Source

For other related troubleshooting information, see Troubleshoot Realms and User Downloads, on page 1945 and Troubleshoot User Control, on page 387.

If you experience issues with the TS Agent-Firepower System integration, check the following:

- You must synchronize the time on your TS Agent server with the time on the Firepower Management Center.
- If the TS Agent monitors the same users as another passive authentication identity source (the User Agent or ISE), the Firepower Management Center prioritizes the TS Agent data. If the TS Agent and a passive identity source report activity by the same IP address, only the TS Agent data is logged to the Firepower Management Center.
- Active FTP sessions are displayed as the **Unknown** user in events. This is normal because, in active FTP, the server (not the client) initiates the connection and the FTP server should not have an associated user name. For more information about active FTP, see RFC 959.

For complete troubleshooting information, see the *Cisco Terminal Services (TS) Agent Configuration Guide*.

The Captive Portal Identity Source

Captive portal is one of the authoritative identity sources supported by the Firepower System. It is the only active authentication method supported by the Firepower System, where users can authenticate onto the network using a managed device.

You typically use captive portal to require authentication to access the internet or to access restricted internal resources; you can optionally configure guest access to resources. After the system authenticates captive portal users, it handles their user traffic according to access control rules. Captive portal performs authentication on HTTP and HTTPS traffic only.

Note HTTPS traffic must be decrypted before captive portal can perform authentication.

Captive portal also records failed authentication attempts. A failed attempt does not add a new user to the list of users in the database. The user activity type for failed authentication activity reported by captive portal is **Failed Auth User**.

The authentication data gained from captive portal can be used for user awareness and user control.

Related Topics

How to Configure the Captive Portal for User Control, on page 1905

Captive Portal Guidelines and Limitations

When you configure and deploy captive portal in an identity policy, users from specified realms authenticate through the following device to access your network:

- Virtual routers on 7000 and 8000 Series devices
- ASA FirePOWER devices in routed mode running Version 9.5(2) or later
- Firepower Threat Defense devices in routed mode

Routed Interface Required

Captive portal active authentication can be performed only by a device with a routed interface configured. If you are configuring the rule for captive portal and your captive portal device contains inline and routed interfaces, you must configure an Interface Conditions to target only the routed interfaces on the device.

If the identity policy referenced by your access control policy contains one or more captive portal identity rules and you deploy the policy on a Firepower Management Center that manages:

- One or more devices with routed interfaces configured, the policy deployment succeeds and the routed interfaces perform active authentication.

  The system does not validate the type of interface in ASA with FirePOWER devices. If you apply a captive portal policy to an inline (tap mode) interface on an ASA with FirePOWER device, the policy deployment succeeds but users in traffic matching those rules are identified as Unknown.

- One or more NGIPSv devices, the policy deployment fails.

Captive Portal and Policies

You configure captive portal in your identity policy and invoke active authentication in your identity rules. Identity policies are associated with access control policies.

You configure some captive portal identity policy settings on the access control policy's **Active Authentication** tab page and configure the rest in an identity rule associated with the access control policy.

An active authentication rule has either an **Active Authentication** rule action, or a **Passive Authentication** rule action with **Use active authentication if passive authentication cannot identify user** selected. In each case the system transparently enables or disables SSL decryption, which restarts the Snort process.

Captive portal authenticates any user in the associated realm, even if that user does not belong to a downloaded group. The system identifies users in non-downloaded groups as Unknown; Unknown users match no identity rules. To avoid that, configure the realm to download users in all groups you expect to authenticate with captive portal.

To make sure the system downloads all users in a realm, make sure the groups are in the Available Groups list in the realm's configuration.

For more information about users and groups, see Download Users and Groups, on page 1960.



Caution Adding the first or removing the last active authentication rule when SSL decryption is disabled (that is, when the access control policy does not include an SSL policy) restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Captive Portal Requirements and Limitations

Note the following requirements and limitations:

- The system supports up to 20 captive portal logins per second.
- There is a maximum five minute limit between failed login attempts for a failed login attempt to be counted toward the count of maximum login attempts. The five minute limit is not configurable.

  (Maximum login attempts are displayed in connection events: **Analysis** > **Connections** > **Events**.)

  If more than five minutes elapse between failed logins, the user will continue to be redirected to captive portal for authentication, will not be designated a failed login user or a guest user, and will not be reported to the Firepower Management Center.
- The only way to be sure a user logs out is to close and reopen the browser. Unless that happens, in some cases, the user can log out of captive portal and be able to access the network without authenticating again using the same browser.
- If a realm is created for a parent domain and the managed device detects a login to a child of that parent domain, the user's subsequent logout is not detected by the managed device.
- To use an ASA FirePOWER device (in routed mode and running ASA version 9.5(2) or later) for captive portal, use the **captive-portal** ASA CLI command to enable captive portal for active authentication and define the port as described in the *ASA Firewall Configuration Guide* (Version 9.5(2) or later): https://www.cisco.com/c/en/us/support/security/adaptive-security-appliance-asa-software/products-installation-and-configuration-guides-list.html.

- You must allow traffic destined for the IP address and port of the device you plan to use for captive portal.
- To perform captive portal active authentication on HTTPS traffic, you must use an SSL policy to decrypt the traffic from the users you want to authenticate. You cannot decrypt the traffic in the connection between a captive portal user's web browser and the captive portal daemon on the managed device; this connection is used to authenticate the captive portal user.
- To limit the amount of non-HTTP or HTTPS traffic that is allowed through the managed device, you should enter typical HTTP and HTTPS ports in the identity policy's **Ports** tab page.

  The managed device changes a previously unseen user from **`Pending`** to **`Unknown`** when it determines that the incoming request does not use the HTTP or HTTPS protocol. As soon as the managed device changes a user from **`Pending`** to another state, access control, Quality of Service, and SSL policies can be applied to that traffic. If your other policies don't permit non-HTTP or HTTPS traffic, configuring ports on the captive portal identity policy can prevent undesired traffic from being allowed through the managed device.

How to Configure the Captive Portal for User Control

High-level overview of how to control user activity with captive portal:

Before you begin

To use the captive portal for active authentication, you must set up an AD or LDAP realm, access control policy, an identity policy, an SSL policy, and associate the identity and SSL policies with the access control policy. Finally, you must deploy the policies to managed devices. This topic provides a high-level summary of those tasks.

An example of the entire procedure begins in Configure the Captive Portal Part 1: Create an Identity Policy, on page 1907.

Perform the following tasks first:

- Confirm that your Firepower Management Center manages one or more devices with a routed interface configured.

  In particular, if your Firepower Management Center manages ASA with FirePOWER devices, see Captive Portal Guidelines and Limitations, on page 1903.
- To use encrypted authentication with the captive portal, either create a PKI object or have your certificate data and key available on the machine from which you're accessing the Firepower Management Center. To create a PKI object, see PKI Objects, on page 447.

Procedure

Step 1 Create and enable a realm as discussed in the following topics:

- Configure a Realm Directory, on page 1959
- Download Users and Groups, on page 1960

Captive portal authenticates any user in the associated realm, even if that user does not belong to a downloaded group. The system identifies users in non-downloaded groups as Unknown; Unknown users match no identity

rules. To avoid that, configure the realm to download users in all groups you expect to authenticate with captive portal.

To make sure the system downloads all users in a realm, make sure the groups are in the Available Groups list in the realm's configuration.

For more information about users and groups, see Download Users and Groups, on page 1960.

Step 2 Create an active authentication identity policy for captive portal.

The identity policy enables selected users in your realm access resources after authenticating with the captive portal.

For more information, see Configure the Captive Portal Part 1: Create an Identity Policy, on page 1907.

Step 3 Configure an access control rule for the captive portal that allows traffic on the captive portal port (by default, TCP 885).

You can choose any available TCP port for the captive portal to use. Whatever your choice, you must create a rule that allows traffic on that port.

For more information, see Configure the Captive Portal Part 2: Create a TCP Port Access Control Rule, on page 1908.

Step 4 Add another access control rule to allow users in the selected realms to access resources using the captive portal.

This enables users to authenticate with captive portal.

For more information, see Configure the Captive Portal Part 3: Create a User Access Control Rule, on page 1909.

Step 5 Configure an SSL decrypt - resign policy for the **Unknown** user so captive portal users can access web pages using the HTTPS protocol.

The captive portal can authenticate users only if the HTTPS traffic is decrypted before the traffic is sent to the captive portal. Captive portal is seen by the system as the **Unknown** user.

For more information, see Configure Captive Portal Part 4: Create an SSL Decrypt-Resign Policy, on page 1910.

Step 6 Asscociate the identity and SSL policies with the access control policy from step 2.

This final step enables the system to authenticate users with the captive portal.

For more information, see Configure Captive Portal Part 5: Associate Identity and SSL Policies with the Access Control Policy, on page 1911.

What to do next

See Configure the Captive Portal Part 1: Create an Identity Policy, on page 1907.

Related Topics

Exclude Applications from Captive Portal, on page 1913
Internal Certificate Objects, on page 455
Troubleshoot the Captive Portal Identity Source, on page 1914
Snort® Restart Scenarios, on page 348

Configure the Captive Portal Part 1: Create an Identity Policy

Before you begin

This five-part procedure shows how to set up the captive portal using the default TCP port 885 and using a Firepower Management Center server certificate for both the captive portal and for SSL decryption. Each part of this example explains one task required to enable the captive portal to perform active authentication.

If you follow all the steps in this procedure, you can configure captive portal to work for users in your domains. You can optionally perform additional tasks, which are discussed in each part of the procedure.

For an overview of the entire procedure, see How to Configure the Captive Portal for User Control, on page 1905.

Procedure

Step 1 Log in to the Firepower Management Center if you have not already done so.

Step 2 Click **Policies** > **Access Control** > **Identity** and create or edit an identity policy.

Step 3 (Optional.) Click **Add Category** to add a category for the captive portal identity rules and enter a **Name** for the category.

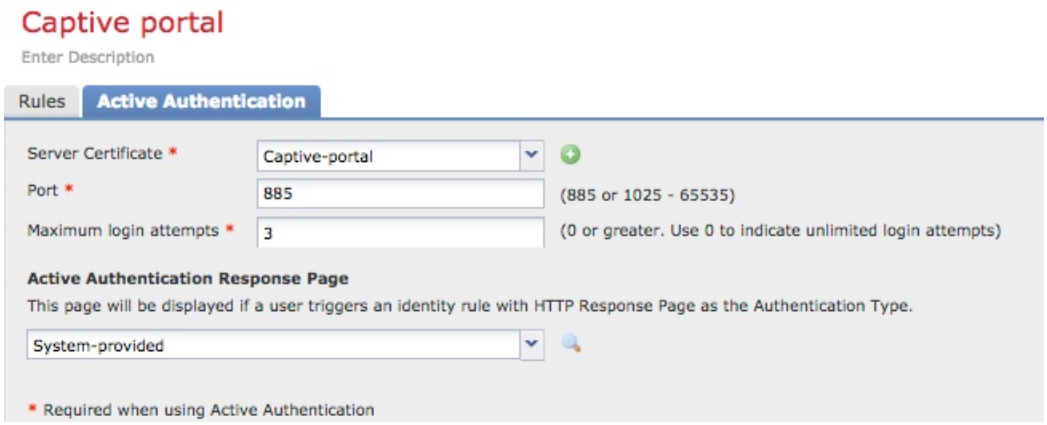
Step 4 Click **Active Authentication**.

Step 5 Choose the appropriate **Server Certificate** from the list or click **Add** ( ) to add a certificate.

Note Captive portal does *not* support the use of Digital Signature Algorithm (DSA) or Elliptic Curve Digital Signature Algorithm (ECDSA) certificates.

Step 6 Enter **885** in the **Port** field and specify the **Maximum login attempts**.

Step 7 (Optional.) Choose an **Active Authentication Response Page** as described in Captive Portal Fields, on page 1912.
The following figure shows an example.



Step 8 Click **Save**.

Step 9 Click **Rules**.

Step 10 Click **Add Rule** to add a new captive portal identity policy rule, or click **Edit** ( ) to edit an existing rule.

Step 11 Enter a **Name** for the rule.

Step 12 From the **Action** list, choose **Active Authentication**.

The system can enforce captive portal active authentication on HTTP and HTTPS traffic only. If an identity rule **Action** is **Active Authentication** (you are using captive portal) or if you are using passive authentication and you check the option on **Realms & Settings** page to **Use active authentication if passive authentication cannot identify user**, use TCP ports constraints only.

Step 13 Click **Realm & Settings**.

Step 14 From the **Realms** list, choose a realm to use for user authentication.

Step 15 (Optional.) Check **Identify as Guest if authentication cannot identify user**. For more information, see Captive Portal Fields, on page 1912.

Step 16 Choose an **Authentication Type** from the list.

Step 17 (Optional.) To exempt specific application traffic from captive portal, see Exclude Applications from Captive Portal, on page 1913.

Step 18 Add conditions to the rule (port, network, and so on) as discussed in Rule Condition Types, on page 363.

Step 19 Click **Add**.

Step 20 At the top of the page, click **Save**.

What to do next

Continue with Configure the Captive Portal Part 2: Create a TCP Port Access Control Rule, on page 1908.

Configure the Captive Portal Part 2: Create a TCP Port Access Control Rule

This part of the procedure shows how to create an access control rule that allows the captive portal to communicate with clients using TCP port 885, which is the captive portal's default port. You can choose another port if you wish, but the port must match the one you chose in Configure the Captive Portal Part 1: Create an Identity Policy, on page 1907.

Before you begin

For an overview of the entire captive portal configuration, see How to Configure the Captive Portal for User Control, on page 1905.

Procedure

Step 1 Log in to the Firepower Management Center if you have not already done so.

Step 2 If you haven't done so already, create a certificate for the captive portal as discussed in PKI Objects, on page 447.

Step 3 Click **Policies** > **Access Control** > **Access Control** and create or edit an access control policy.

Step 4 Click **Add Rule**.

Step 5 Enter a **Name** for the rule.

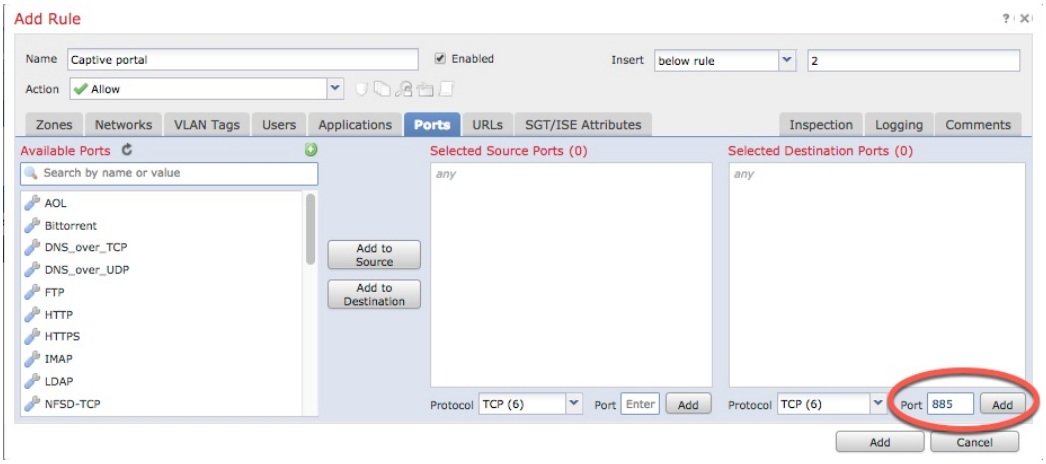
Step 6 Choose **Allow** from the **Action** list.

Step 7 Click **Ports**.

Step 8 From the **Protocol** list under the **Selected Destination Ports** field, choose **TCP**.

Step 9 In the **Port** field, enter **885**.

Step 10 Click **Add** next to the **Port** field.
The following figure shows an example.



Step 11 Click **Add** at the bottom of the page.

What to do next

Continue with Configure the Captive Portal Part 3: Create a User Access Control Rule, on page 1909.

Configure the Captive Portal Part 3: Create a User Access Control Rule

This part of the procedure discusses how to add an access control rule that enables users in a realm to authenticate using captive portal.

Before you begin

For an overview of the entire captive portal configuration, see How to Configure the Captive Portal for User Control, on page 1905.

Procedure

Step 1 In the rule editor, click **Add Rule**.

Step 2 Enter a **Name** for the rule.

Step 3 Choose **Allow** from the **Action** list.

Step 4 Click **Users**.

Step 5 In the **Available Realms** list, click the realms to allow.

Step 6 If no realms display, click **Refresh** ( ).

Step 7 In the **Available Users** list, choose the users to add to the rule and click **Add to Rule**.

Step 8 (Optional.) Add conditions to the access control policy as discussed in Rule Condition Types, on page 363.

Step 9 Click **Add**.

Step 10 On the access control rule page, click **Save**.

Step 11 In the policy editor, set the rule position. Click and drag or use the right-click menu to cut and paste. Rules are numbered starting at 1. The system matches traffic to rules in top-down order by ascending rule number.

The first rule that traffic matches is the rule that handles that traffic. Proper rule order reduces the resources required to process network traffic and prevents rule preemption.

What to do next

Continue with Configure Captive Portal Part 4: Create an SSL Decrypt-Resign Policy, on page 1910.

Configure Captive Portal Part 4: Create an SSL Decrypt-Resign Policy

This part of the procedure discusses how to create an SSL access policy to decrypt and resign traffic before the traffic reaches the captive portal. The captive portal can authenticate traffic only after it has been decrypted.

Before you begin

For an overview of the entire captive portal configuration, see How to Configure the Captive Portal for User Control, on page 1905.

Procedure

Step 1 If you haven't done so already, create a certificate object to decrypt SSL traffic as discussed in PKI Objects, on page 447.

Step 2 Click **Policies** > **Access Control** > **SSL**.

Step 3 Click **New Policy**.

Step 4 Enter a **Name** and choose a **Default Action** for the policy. Default actions are discussed in SSL Policy Default Actions, on page 1338.

Step 5 Click **Save**.

Step 6 Click **Add Rule**.

Step 7 Enter a **Name** for the rule.

Step 8 From the **Action** list, choose **Decrypt - Resign**.

Step 9 From the **with** list, choose your PKI object.

Step 10 Click **Users**.

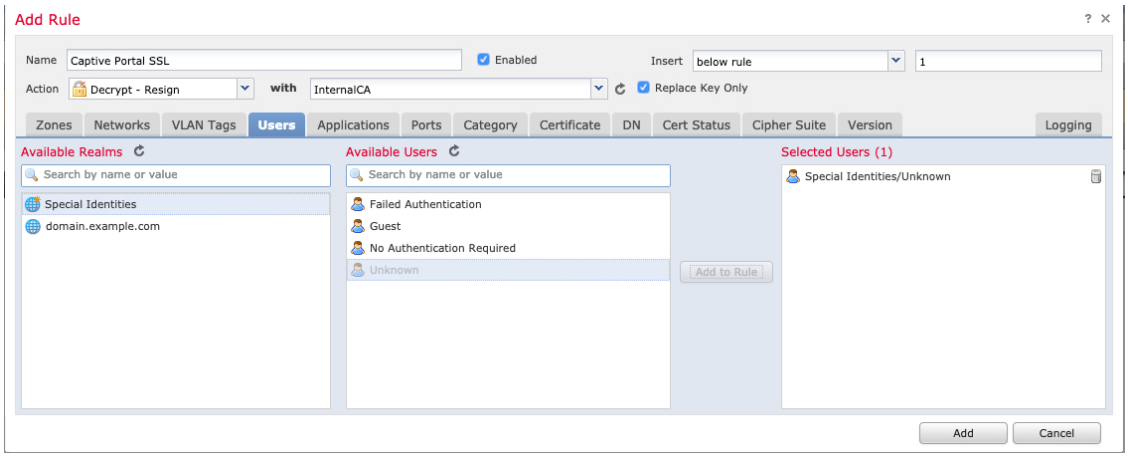
Step 11 Above the **Available Realms** list, click **Refresh** ( ).

Step 12 In the **Available Realms** list, click **Special Identities**.

Step 13 In the **Available Users** list, click **Unknown**.

Step 14 Click **Add to Rule**.

The following figure shows an example.



Step 15 (Optional.) Set other options as discussed in TLS/SSL Rule Conditions, on page 1357.

Step 16 Click **Add**.

Step 17 At the top of the page, click **Save**.

What to do next

Continue with Configure Captive Portal Part 5: Associate Identity and SSL Policies with the Access Control Policy, on page 1911.

Configure Captive Portal Part 5: Associate Identity and SSL Policies with the Access Control Policy

This part of the procedure discusses how to associate the identity policy and SSL **Decrypt - Resign** rule with the access control policy you created earlier. After this, users can authenticate using the captive portal.

Before you begin

For an overview of the entire captive portal configuration, see How to Configure the Captive Portal for User Control, on page 1905.

Procedure

Step 1 Click **Policies** > **Access Control** > **Access Control** and edit the access control policy you created as discussed in Configure the Captive Portal Part 2: Create a TCP Port Access Control Rule, on page 1908. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 2 Either create a new access control policy or edit an existing policy.

Step 3 At the top of the page, click the link next to **Identity Policy**.

Step 4 From the list, choose the name of your identity policy and, at the top of the page, click **Save**.

Step 5 Repeat the preceding steps to associate your captive portal SSL policy with the access control policy.

Step 6 If you haven't done so already, target the policy at managed devices as discussed in Setting Target Devices for an Access Control Policy, on page 1210.

What to do next

- Deploy your identity and access control policies to managed devices as discussed in Deploy Configuration Changes, on page 346.
- Monitor user activity as discussed in Using Workflows, on page 2150.

Captive Portal Fields

Use the following fields to configure captive portal on the **Active Authentication** tab page of your identity policy. See also Identity Rule Fields, on page 1963 and Exclude Applications from Captive Portal, on page 1913.

Server Certificate

The server certificate presented by the captive portal daemon.

Note Captive portal does *not* support the use of Digital Signature Algorithm (DSA) or Elliptic Curve Digital Signature Algorithm (ECDSA) certificates.

Port

The port number to use for the captive portal connection. If you plan to use an ASA FirePOWER device for captive portal, the port number in this field must match the port number you configured on the ASA FirePOWER device using the **captive-portal** CLI command.

Maximum login attempts

The maximum allowed number of failed login attempts before the system denies a user's login request.

Active Authentication Response Page

The system-provided or custom HTTP response page you want to display to captive portal users. After you select an **Active Authentication Response Page** in your identity policy active authentication settings, you also must configure one or more identity rules with **`HTTP Response Page`** as the .

The system-provided HTTP response page includes **Username** and **Password** fields, as well as a **Login as guest** button to allow users to access the network as guests. To display a single login method, configure a custom HTTP response page.

Choose the following options:

- To use a generic response, click **System-provided**. You can click **View** ( ) to view the HTML code for this page.
- To create a custom response, click **Custom**. A window with system-provided code is displayed that you can replace or modify. When you are done, save your changes. You can edit a custom page by clicking **Edit** ( ).

Related Topics

Internal Certificate Objects, on page 455

Exclude Applications from Captive Portal

You can select applications (identified by their HTTP `User-Agent` strings) and exempt them from captive portal active authentication. This allows traffic from the selected applications to pass through the identity policy without authenticating.

Note Only applications with the **`User-Agent Exclusion` Tag** are displayed in this list.

Procedure

Step 1 If you haven't done so already, log in to the Firepower Management Center.

Step 2 Click **Policies** > **Access Control** > **Identity** .

Step 3 Edit the identity policy that contains the captive portal rule.

Step 4 On **Realm & Settings** tab page, use the filters in the **Application Filters** list to narrow the applications you want to add to the filter.

- Click the arrow next to each filter type to expand and collapse the list.
- Right-click a filter type and click **Check All** or **Uncheck All**. Note that the list indicates how many filters you have selected of each type.
- To narrow the filters that are displayed, type a search string in the **Search by name** field; this is especially useful for categories and tags. To clear the search, click **Clear** ( ).
- To refresh the filters list and clear any selected filters, click **Reload** ( ).
- To clear all filters and search fields, click **Clear All Filters**.

Note The list displays 100 applications at a time.

Step 5 Choose the applications that you want to add to the filter from the **Available Applications** list:

- To narrow the individual applications that appear, enter a search string in the **Search by name** field. To clear the search, click **Clear** ( ).
- Use paging at the bottom of the list to browse the list of individual available applications.
- To refresh the applications list and clear any selected applications, click **Reload** ( ).

Step 6 Add the selected applications to exclude from external authentication. You can click and drag, or you can click **Add to Rule**. The result is the combination of the application filters you selected.

What to do next

- Continue configuring the identity rule as described in Create an Identity Rule, on page 1962.

Troubleshoot the Captive Portal Identity Source

For other related troubleshooting information, see Troubleshoot Realms and User Downloads, on page 1945 and Troubleshoot User Control, on page 387.

If you experience issues with captive portal, check the following:

- The time on your captive portal server must be synchronized with the time on the Firepower Management Center.
- If you have DNS resolution configured and you create an identity rule to perform **`Kerberos`** (or **`HTTP Negotiate`**, if you want Kerberos as an option) captive portal, you must configure your DNS server to resolve the fully qualified domain name (FQDN) of the captive portal device. The FQDN must match the hostname you provided when configuring DNS.

  For ASA with FirePOWER Services and Firepower Threat Defense devices, the FQDN must resolve to the IP address of the routed interface used for captive portal.
- If you're using Kerberos authentication, the managed device's host name must be less than 15 characters (it's a NetBIOS limitation set by Windows); otherwise, captive portal authentication fails. You set the managed device host name when you set up the device. For more information, see an article like this one on the Microsoft documentation site: Naming conventions in Active Directory for computers, domains, sites, and OUs.
- DNS must return a response of 512 bytes or less to the hostname; otherwise, testing the connection the AD connection fails. This limit applies in both directions and is discussed in RFC 6891 section-6.2.5.
- If you select **`Kerberos`** (or **`HTTP Negotiate`**, if you want Kerberos as an option) as the **Authentication Type** in an identity rule, the **Realm** you select must be configured with an **AD Join Username** and **AD Join Password** to perform Kerberos captive portal active authentication.
- If you select **`HTTP Basic`** as the **Authentication Type** in an identity rule, users on your network might not notice their sessions time out. Most web browsers cache the credentials from **`HTTP Basic`** logins and use the credentials to seamlessly begin a new session after an old session times out.
- If the connection between your Firepower Management Center and a managed device fails, no captive portal logins reported by the device can be identified during the downtime, unless the users were previously seen and downloaded to the Firepower Management Center. The unidentified users are logged as Unknown users on the Firepower Management Center. After the downtime, the Unknown users are reidentified and processed according to the rules in your identity policy.
- If the device you want to use for captive portal contains both inline and routed interfaces, you must configure a zone condition in your captive portal identity rules to target only the routed interfaces on the captive portal device.
- The system does not validate the type of interface in ASA with FirePOWER devices. If you apply a captive portal policy to an inline (tap mode) interface on an ASA with FirePOWER device, the policy deployment succeeds but users in traffic matching those rules are identified as Unknown.
- The host name of the managed device must be less than 15 characters for Kerberos authentication to succeed.
- The only way to be sure a user logs out is to close and reopen the browser. Unless that happens, in some cases, the user can log out of captive portal and be able to access the network without authenticating again using the same browser.

- Active FTP sessions are displayed as the **Unknown** user in events. This is normal because, in active FTP, the server (not the client) initiates the connection and the FTP server should not have an associated user name. For more information about active FTP, see RFC 959.
- Captive portal authenticates any user in the associated realm, even if that user does not belong to a downloaded group. The system identifies users in non-downloaded groups as Unknown; Unknown users match no identity rules. To avoid that, configure the realm to download users in all groups you expect to authenticate with captive portal.

  To make sure the system downloads all users in a realm, make sure the groups are in the Available Groups list in the realm's configuration.

  For more information about users and groups, see Download Users and Groups, on page 1960.

The Traffic-Based Detection Identity Source

Traffic-based detection is the only non-authoritative identity source supported by the Firepower System. When configured, managed devices detect LDAP, AIM, POP3, IMAP, Oracle, SIP (VoIP), FTP, HTTP, MDNS, and SMTP logins on the networks you specify. The data gained from traffic-based detection can be used only for user awareness. Unlike authoritative identity sources, you configure traffic-based detection in your network discovery policy as described in Configuring Traffic-Based User Detection, on page 1928.

Note the following limitations:

- Traffic-based detection interprets only Kerberos logins for LDAP connections as LDAP authentications. Managed devices cannot detect encrypted LDAP authentications using protocols such as SSL or TLS.
- Traffic-based detection detects AIM logins using the OSCAR protocol only. They cannot detect AIM logins using TOC2.
- Traffic-based detection cannot restrict SMTP logging. This is because users are not added to the database based on SMTP logins; although the system detects SMTP logins, the logins are not recorded unless there is already a user with a matching email address in the database.

Traffic-based detection also records failed login attempts. A failed login attempt does not add a new user to the list of users in the database. The user activity type for detected failed login activity detected by traffic-based detection is **Failed User Login**.

Note The system cannot distinguish between failed and successful HTTP logins. To see HTTP user information, you must enable **Capture Failed Login Attempts** in the traffic-based detection configuration.

Caution Enabling or disabling non-authoritative, traffic-based user detection over the HTTP, FTP, or MDNS protocols, using the network discovery policy restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Traffic-Based Detection Data

When a device detects a login using traffic-based detection, it sends the following information to the Firepower Management Center to be logged as user activity:

- the user name identified in the login
- the time of the login
- the IP address involved in the login, which can be the IP address of the user's host (for LDAP, POP3, IMAP, and AIM logins), the server (for HTTP, MDNS, FTP, SMTP and Oracle logins), or the session originator (for SIP logins)
- the user's email address (for POP3, IMAP, and SMTP logins)
- the name of the device that detected the login

If the user was previously detected, the Firepower Management Center updates that user's login history. Note that the Firepower Management Center can use the email addresses in POP3 and IMAP logins to correlate with LDAP users. This means that, for example, if the Firepower Management Center detects a new IMAP login, and the email address in the IMAP login matches that for an existing LDAP user, the IMAP login does not create a new user; rather, it updates the LDAP user's history.

If the user was previously undetected, the Firepower Management Center adds the user to the users database. Unique AIM, SIP, and Oracle logins always create new user records, because there is no data in those login events that the Firepower Management Center can correlate with other login types.

The Firepower Management Center does **not** log user activity or user identities in the following cases:

- if you configured the network discovery policy to ignore that login type
- if a managed device detects an SMTP login, but the users database does not contain a previously detected LDAP, POP3, or IMAP user with a matching email address

The user data is added to the users table.

Traffic-Based Detection Strategies

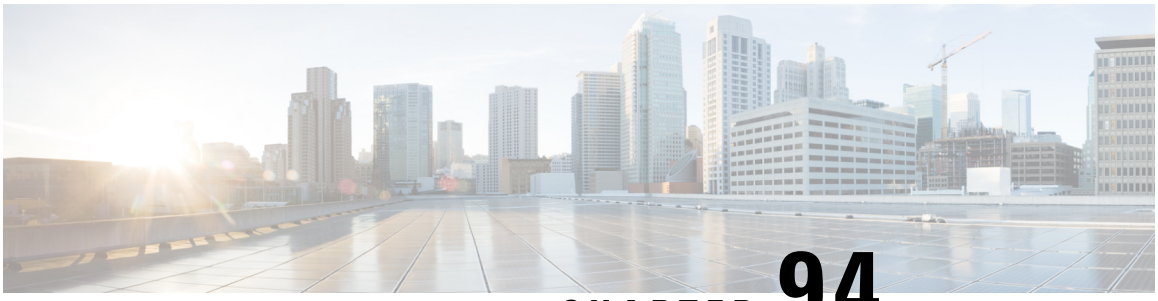
You can restrict the protocols where user activity is discovered to reduce the total number of detected users so you can focus on users likely to provide the most complete user information. Limiting protocol detection helps minimize user name clutter and preserve storage space on your Firepower Management Center.

Consider the following when selecting traffic-based detection protocols:

- Obtaining user names through protocols such as AIM, POP3, and IMAP may introduce user names not relevant to your organization due to network access from contractors, visitors, and other guests.
- AIM, Oracle, and SIP logins may create extraneous user records. This occurs because these login types are not associated with any of the user metadata that the system obtains from an LDAP server, nor are they associated with any of the information contained in the other types of login that your managed devices detect. Therefore, the Firepower Management Center cannot correlate these users with other types of users.

Related Topics

Configuring Traffic-Based User Detection, on page 1928



CHAPTER 94

Network Discovery Policies

The following topics describe how to create, configure, and manage network discovery policies:

- Overview: Network Discovery Policies, on page 1917
- Requirements and Prerequisites for Network Discovery Policies, on page 1918
- Network Discovery Customization, on page 1918
- Network Discovery Rules, on page 1919
- Configuring Advanced Network Discovery Options, on page 1928
- Troubleshooting Your Network Discovery Strategy, on page 1938

Overview: Network Discovery Policies

The network discovery policy on the Firepower Management Center controls how the system collects data on your organization's network assets and which network segments and ports are monitored.

In a multidomain deployment, each leaf domain has an independent network discovery policy. Network discovery policy rules and other settings cannot be shared, inherited, or copied between domains. Whenever you create a new domain, the system creates a network discovery policy for the new domain, using default settings. You must explicitly apply any desired customizations to the new policy.

Discovery rules within the policy specify which networks and ports the Firepower System monitors to generate discovery data based on network data in traffic, and the zones to which the policy is deployed. Within a rule, you can configure whether hosts, applications, and non-authoritative users are discovered. You can create rules to exclude networks and zones from discovery. You can configure discovery of data from NetFlow exporters and restrict the protocols for traffic where user data is discovered on your network.

The network discovery policy has a single default rule in place, configured to discover applications from all observed traffic. The rule does not exclude any networks, zones, or ports, host and user discovery is not configured, and the rule is not configured to monitor a NetFlow exporter. This policy is deployed by default to any managed devices when they are registered to the Firepower Management Center. To begin collecting host or user data, you must add or modify discovery rules and re-deploy the policy to a device.

If you want to adjust the scope of network discovery, you can create additional discovery rules and modify or remove the default rule.

Remember that the access control policy for each managed device defines the traffic that you permit for that device and, therefore, the traffic you can monitor with network discovery. If you block certain traffic using access control, the system cannot examine that traffic for host, user, or application activity. For example, if an access control policy blocks access to social networking applications, the system cannot provide any discovery data on those applications.

If you enable traffic-based user detection in your discovery rules, you can detect non-authoritative users through user login activity in traffic over a set of application protocols. You can disable discovery in particular protocols across all rules if needed. Disabling some protocols can help avoid reaching the user limit associated with your Firepower Management Center model, reserving available user count for users from the other protocols.

Advanced network discovery settings allow you to manage what data is logged, how discovery data is stored, what indications of compromise (IOC) rules are active, what vulnerability mappings are used for impact assessment, and what happens when sources offer conflicting discovery data. You can also add sources for host input and NetFlow exporters to monitor.

Requirements and Prerequisites for Network Discovery Policies

Model Support

Any.

Supported Domains

Leaf

User Roles

- Admin
- Discovery Admin

Network Discovery Customization

The information about your network traffic collected by the Firepower System is most valuable to you when the system can correlate this information to identify the hosts on your network that are most vulnerable and most important.

As an example, if you have several devices on your network running a customized version of SuSE Linux, the system cannot identify that operating system and so cannot map vulnerabilities to the hosts. However, knowing that the system has a list of vulnerabilities for SuSE Linux, you may want to create a custom fingerprint for one of the hosts that can then be used to identify the other hosts running the same operating system. You can include a mapping of the vulnerability list for SuSE Linux in the fingerprint to associate that list with each host that matches the fingerprint.

The system also allows you to input host data from third-party systems directly into the network map, using the host input feature. However, third-party operating system or application data does not automatically map to vulnerability information. If you want to see vulnerabilities and perform impact correlation for hosts using third-party operating system, server, and application protocol data, you must map the vendor and version information from the third-party system to the vendor and version listed in the vulnerability database (VDB). You also may want to maintain the host input data on an ongoing basis. Note that even if you map application data to Firepower System vendor and version definitions, imported third-party vulnerabilities are not used for impact assessment for clients or web applications.

If the system cannot identify application protocols running on hosts on your network, you can create user-defined application protocol detectors that allow the system to identify the applications based on a port

or a pattern. You can also import, activate, and deactivate certain application detectors to further customize the application detection capability of the Firepower System.

You can also replace detection of operating system and application data using scan results from the Nmap active scanner or augment the vulnerability lists with third-party vulnerabilities. The system may reconcile data from multiple sources to determine the identity for an application.

Configuring the Network Discovery Policy

In a multidomain deployment, each domain has a separate network discovery policy. If your user account can manage multiple domains, switch to the leaf domain where you want to configure the policy.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Configure the following components of your policy:

- Discovery rules — See Configuring Network Discovery Rules, on page 1920.
- Traffic-based detection for users — See Configuring Traffic-Based User Detection, on page 1928.
- Advanced network discovery options — See Configuring Advanced Network Discovery Options, on page 1928.
- Custom operating system definitions (fingerprints) — See Creating a Custom Fingerprint for Clients, on page 1840 and Creating a Custom Fingerprint for Servers, on page 1842.

Network Discovery Rules

Network discovery rules allow you to tailor the information discovered for your network map to include only the specific data you want. Rules in your network discovery policy are evaluated sequentially. You can create rules with overlapping monitoring criteria, but doing so may affect your system performance.

When you exclude a host or a network from monitoring, the host or network does not appear in the network map and no events are reported for it. However, when the host discovery rules for the local IP are disabled, the detection engine instances are impacted by a higher processing load, as it builds data from each flow afresh rather than using the existing host data.

We recommend that you exclude load balancers (or specific ports on load balancers) and NAT devices from monitoring. These devices may create excessive and misleading events, filling the database and overloading the Firepower Management Center. For example, a monitored NAT device might exhibit multiple updates of its operating system in a short period of time. If you know the IP addresses of your load balancers and NAT devices, you can exclude them from monitoring.

Tip The system can identify many load balancers and NAT devices by examining your network traffic.

In addition, if you need to create a custom server fingerprint, you should temporarily exclude from monitoring the IP address that you are using to communicate with the host you are fingerprinting. Otherwise, the network map and discovery event views will be cluttered with inaccurate information about the host represented by that IP address. After you create the fingerprint, you can configure your policy to monitor that IP address again.

Cisco also recommends that you **not** monitor the same network segment with NetFlow exporters and Firepower System managed devices. Although ideally you should configure your network discovery policy with non-overlapping rules, the system does drop duplicate connection logs generated by managed devices. However, you **cannot** drop duplicate connection logs for connections detected by both a managed device and a NetFlow exporter.

Configuring Network Discovery Rules

You can configure discovery rules to tailor the discovery of host and application data to your needs.

Before you begin

- Make sure you are logging connections for the traffic where you want to discover network data; see Best Practices for Connection Logging, on page 2216.
- If you want to collect exported NetFlow records, add a NetFlow Exporter as described in Adding NetFlow Exporters to a Network Discovery Policy, on page 1934.
- If you will want to view discovery performance graphs, you must enable hosts, users, and applications in your discovery rule. Note that this may impact system performance.



Tip In most cases, Cisco suggests restricting discovery to the addresses in RFC 1918.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Add Rule**.

Step 3 Set the **Action** for the rule as described in Actions and Discovered Assets, on page 1921.

Step 4 Set optional discovery parameters:

- Restrict the rule action to specific networks; see Restricting the Monitored Network, on page 1922.
- Restrict the rule action to traffic in specific zones; see Configuring Zones in Network Discovery Rules, on page 1926.
- Exclude ports from monitoring; see Excluding Ports in Network Discovery Rules, on page 1924.
- Configure the rule for NetFlow data discovery; see Configuring Rules for NetFlow Data Discovery, on page 1922.

Step 5 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Actions and Discovered Assets

When you configure a discovery rule, you must select an action for the rule. The effect of that action depends on whether you are using the rule to discover data from a managed device or from a NetFlow exporter.

The following table describes what assets are discovered by rules with the specified action settings in those two scenarios.

Table 222: Discovery Rule Actions

Action	Option	Managed Device	NetFlow Exporter
Exclude	--	Excludes the specified network from monitoring. If the source or destination host for a connection is excluded from discovery, the connection is recorded but discovery events are not created for excluded hosts.	Excludes the specified network from monitoring. If the source or destination host for a connection is excluded from discovery, the connection is recorded but discovery events are not created for excluded hosts.
Discover	Hosts	Adds hosts to the network map based on discovery events. (Optional, unless user discovery is enabled, then required.)	Adds hosts to the network map and logs connections based on NetFlow records. (Required)
Discover	Applications	Adds applications to the network map based on application detectors. Note that you cannot discover hosts or users in a rule without also discovering applications. (Required)	Adds application protocols to the network map based on NetFlow records and the port-application protocol correlation in `/etc/sf/services`. (Optional)
Discover	Users	Adds users to the users table and logs user activity based on traffic-based detection on the user protocols configured in the network discovery policy. (Optional)	n/a
Log NetFlow Connections	--	n/a	Logs NetFlow connections only. Does not discover hosts or applications.

If you want the rule to monitor managed device traffic, application logging is required. If you want the rule to monitor users, host logging is required. If you want the rule to monitor exported NetFlow records, you cannot configure it to log users, and logging applications is optional.

Note The system detects connections in exported NetFlow records based on the **Action** settings in the network discovery policy. The system detects connections in managed device traffic based on access control policy settings.

Monitored Networks

A discovery rule causes discovery of monitored assets only in traffic to and from hosts in the specified networks. For a discovery rule, discovery occurs for connections that have at least one IP address within the networks

specified, with events generated only for IP addresses within the networks to monitor. The default discovery rule discovers applications from all observed traffic (`0.0.0.0/0` for all IPv4 traffic, and `::/0` for all IPv6 traffic).

If you configure a rule to handle NetFlow discovery and log only connections data, the system also logs connections to and from IP addresses in the specified networks. Note that network discovery rules provide the only way to log NetFlow network connections.

You can also use network object or object groups to specify the networks to monitor.

Restricting the Monitored Network

Every discovery rule must include at least one network.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Add Rule**.

Step 3 Click **Networks**, if it is not already open.

Step 4 Optionally, add network objects to the Available Networks list as described in Creating Network Objects During Discovery Rule Configuration, on page 1923.

Note If you modify a network object used in the network discovery policy, the changes do not take effect for discovery until you deploy the configuration changes.

Step 5 Specify a network:

- Choose a network from the **Available Networks** list.

  Tip If the network does not immediately appear on the list, click **Reload** ( ).

- Enter the IP address into the text box below the Available Networks label.

Step 6 Click **Add**.

Step 7 Optionally, repeat the previous two steps to add additional networks.

Step 8 Click **Save** to save the changes you made.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configuring Rules for NetFlow Data Discovery

The Firepower System can use data from NetFlow exporters to generate connection and discovery events, and to add host and application data to the network map.

If you choose a NetFlow exporter in a discovery rule, the rule is limited to discovery of NetFlow data for the specified networks. Choose the NetFlow device to monitor before you configure other aspects of rule behavior,

as the available rule actions change when you choose a NetFlow device. You cannot configure port exclusions for monitoring NetFlow exporters.

Before you begin

- Add NetFlow-enabled devices to the network discovery policy; see Adding NetFlow Exporters to a Network Discovery Policy, on page 1934.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Add Rule**.

Step 3 Choose **NetFlow Device**.

Step 4 From the **Netflow Device** drop-down list, choose the IP address of the NetFlow exporter to be monitored.

Step 5 Specify the type of NetFlow data you want the Firepower System managed device to collect:

- Connection only — Choose `Log NetFlow Connections` from the **Action** drop-down list.
- Host, Application, and Connection — Choose `Discover` from the **Action** drop-down list. The system automatically checks the **Hosts** check box and enables collection of connection data. Optionally, you can check the **Application** check box to collect application data.

Step 6 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Creating Network Objects During Discovery Rule Configuration

You can add new network objects to the list of available networks that appears in a discovery rule by adding them to the list of reusable network objects and groups.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 In **Networks**, click **Add Rule**.

Step 3 Click **Add** () next to **Available Networks**.

Step 4 Create a network object as described in Creating Network Objects, on page 403.

Step 5 Finish adding the network discovery rule as described in Configuring Network Discovery Rules, on page 1920.

Port Exclusions

Just as you can exclude hosts from monitoring, you can exclude specific ports from monitoring. For example:

- Load balancers can report multiple applications on the same port in a short period of time. You can configure your network discovery rules so that they exclude that port from monitoring, such as excluding port 80 on a load balancer that handles a web farm.
- Your organization may use a custom client that uses a specific range of ports. If the traffic from this client generates excessive and misleading events, you can exclude those ports from monitoring. Similarly, you may decide that you do not want to monitor DNS traffic. In that case, you could configure your rules so that your discovery policy does not monitor port 53.

When adding ports to exclude, you can decide whether to use a reusable port object from the Available Ports list, add ports directly to the source or destination exclusion lists, or create a new reusable port and then move it into the exclusion lists.

Note You cannot exclude ports in rules handling NetFlow data discovery.

Excluding Ports in Network Discovery Rules

You cannot exclude ports in rules handling NetFlow data discovery.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Add Rule**.

Step 3 Click **Port Exclusions**.

Step 4 Optionally, add port objects to the Available Ports list as described in Creating Port Objects During Discovery Rule Configuration, on page 1925.

Step 5 Exclude specific source ports from monitoring, using either of the following methods:

- Choose a port or ports from the **Available Ports** list and click **Add to Source**.
- To exclude traffic from a specific source port without adding a port object, under the **Selected Source Ports** list, choose a **Protocol**, enter a **Port** number (a value from 1 to 65535), and click **Add**.

Step 6 Exclude specific destination ports from monitoring, using either of the following methods:

- Choose a port or ports from the **Available Ports** list and click **Add to Destination**.
- To exclude traffic from a specific destination port without adding a port object, under the **Selected Destination Ports** list, choose a **Protocol**, enter a **Port** number, and click **Add**.

Step 7 Click **Save** to save the changes you made.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Creating Port Objects During Discovery Rule Configuration

You can add new port objects to the list of available ports that appears in a discovery rule by adding them to the list of reusable port objects and groups that can be used anywhere in the Firepower System.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 In Networks, click **Add Rule**.

Step 3 Click **Port Exclusions**.

Step 4 To add a port to the Available Ports list, click **Add** ( ).

Step 5 Supply a **Name**.

Step 6 In the **Protocol** field, specify the protocol of the traffic you want to exclude.

Step 7 In the **Port** field, enter the ports you want to exclude from monitoring.

You can specify a single port, a range of ports using the dash (-), or a comma-separated list of ports and port ranges. Allowed port values are from 1 to 65535.

Step 8 Click **Save**.

Step 9 If the port does not immediately appear on the list, click **Refresh**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Zones in Network Discovery Rules

To improve performance, discovery rules can be configured so that the zones in the rule include the sensing interfaces on your managed devices that are physically connected to the networks-to-monitor in the rule.

Unfortunately, you may not always be kept informed of network configuration changes. A network administrator may modify a network configuration through routing or host changes without informing you, which may make it challenging to stay on top of proper network discovery policy configurations. If you do not know how the sensing interfaces on your managed devices are physically connected to your network, leave the zone configuration as the default. This default causes the system to deploy the discovery rule to all zones in your deployment. (If no zones are excluded, the system deploy the discovery policy to all zones.)

Configuring Zones in Network Discovery Rules

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Add Rule**.

Step 3 Click **Zones**.

Step 4 Choose a zone or zones from the **Available Zones** list.

Step 5 Click **Save** to save the changes you made.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

The Traffic-Based Detection Identity Source

Traffic-based detection is the only non-authoritative identity source supported by the Firepower System. When configured, managed devices detect LDAP, AIM, POP3, IMAP, Oracle, SIP (VoIP), FTP, HTTP, MDNS, and SMTP logins on the networks you specify. The data gained from traffic-based detection can be used only for user awareness. Unlike authoritative identity sources, you configure traffic-based detection in your network discovery policy as described in Configuring Traffic-Based User Detection, on page 1928.

Note the following limitations:

- Traffic-based detection interprets only Kerberos logins for LDAP connections as LDAP authentications. Managed devices cannot detect encrypted LDAP authentications using protocols such as SSL or TLS.
- Traffic-based detection detects AIM logins using the OSCAR protocol only. They cannot detect AIM logins using TOC2.
- Traffic-based detection cannot restrict SMTP logging. This is because users are not added to the database based on SMTP logins; although the system detects SMTP logins, the logins are not recorded unless there is already a user with a matching email address in the database.

Traffic-based detection also records failed login attempts. A failed login attempt does not add a new user to the list of users in the database. The user activity type for detected failed login activity detected by traffic-based detection is **Failed User Login**.



Note The system cannot distinguish between failed and successful HTTP logins. To see HTTP user information, you must enable **Capture Failed Login Attempts** in the traffic-based detection configuration.

Caution Enabling or disabling non-authoritative, traffic-based user detection over the HTTP, FTP, or MDNS protocols, using the network discovery policy restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Traffic-Based Detection Data

When a device detects a login using traffic-based detection, it sends the following information to the Firepower Management Center to be logged as user activity:

- the user name identified in the login
- the time of the login
- the IP address involved in the login, which can be the IP address of the user's host (for LDAP, POP3, IMAP, and AIM logins), the server (for HTTP, MDNS, FTP, SMTP and Oracle logins), or the session originator (for SIP logins)
- the user's email address (for POP3, IMAP, and SMTP logins)
- the name of the device that detected the login

If the user was previously detected, the Firepower Management Center updates that user's login history. Note that the Firepower Management Center can use the email addresses in POP3 and IMAP logins to correlate with LDAP users. This means that, for example, if the Firepower Management Center detects a new IMAP login, and the email address in the IMAP login matches that for an existing LDAP user, the IMAP login does not create a new user; rather, it updates the LDAP user's history.

If the user was previously undetected, the Firepower Management Center adds the user to the users database. Unique AIM, SIP, and Oracle logins always create new user records, because there is no data in those login events that the Firepower Management Center can correlate with other login types.

The Firepower Management Center does **not** log user activity or user identities in the following cases:

- if you configured the network discovery policy to ignore that login type
- if a managed device detects an SMTP login, but the users database does not contain a previously detected LDAP, POP3, or IMAP user with a matching email address

The user data is added to the users table.

Traffic-Based Detection Strategies

You can restrict the protocols where user activity is discovered to reduce the total number of detected users so you can focus on users likely to provide the most complete user information. Limiting protocol detection helps minimize user name clutter and preserve storage space on your Firepower Management Center.

Consider the following when selecting traffic-based detection protocols:

- Obtaining user names through protocols such as AIM, POP3, and IMAP may introduce user names not relevant to your organization due to network access from contractors, visitors, and other guests.
- AIM, Oracle, and SIP logins may create extraneous user records. This occurs because these login types are not associated with any of the user metadata that the system obtains from an LDAP server, nor are

they associated with any of the information contained in the other types of login that your managed devices detect. Therefore, the Firepower Management Center cannot correlate these users with other types of users.

Related Topics

Configuring Traffic-Based User Detection, on page 1928

Configuring Traffic-Based User Detection

When you enable traffic-based user detection in a network discovery rule, host discovery is automatically enabled. For more information about traffic-based detection, see The Traffic-Based Detection Identity Source, on page 1915.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Users**.

Step 3 Click **Edit** ( ).

Step 4 Check the check boxes for protocols where you want to detect logins or clear check boxes for protocols where you do not want to detect logins.

Step 5 Optionally, to record failed login attempts detected in LDAP, POP3, FTP, or IMAP traffic, or to capture user information for HTTP logins, enable **Capture Failed Login Attempts**.

Step 6 Click **Save**.

What to do next

Caution Enabling or disabling non-authoritative, traffic-based user detection over the HTTP, FTP, or MDNS protocols, using the network discovery policy restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

- Configure network discovery rules to discover users as described in Configuring Network Discovery Rules, on page 1920.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configuring Advanced Network Discovery Options

The Advanced of the network discovery policy allows you to configure policy-wide settings for what events are detected, how long discovery data is retained and how often it is updated, what vulnerability mappings

are used for impact correlation, and how operating system and server identity conflicts are resolved. In addition, you can add host input sources and NetFlow exporters to allow import of data from other sources.

Note Database event limits for discovery and user activity events are set in system configuration.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Advanced**.

Step 3 Click **Edit** ( ) or **Add** ( ) next to the setting you want to modify:

- Data Storage Settings — Update the settings as described in Configuring Network Discovery Data Storage, on page 1936.
- Event Logging Settings — Update the settings as described in Configuring Network Discovery Event Logging, on page 1936.
- General Settings — Update the settings as described in Configuring Network Discovery General Settings, on page 1930.
- Identity Conflict Settings — Update the settings as described in Configuring Network Discovery Identity Conflict Resolution, on page 1931.
- Indications of Compromise Settings — Update the settings as described in Enabling Indications of Compromise Rules, on page 1933.
- NetFlow Exporters — Update the settings as described in Adding NetFlow Exporters to a Network Discovery Policy, on page 1934.
- OS and Server Identity Sources — Update the settings as described in Adding Network Discovery OS and Server Identity Sources, on page 1937.
- Vulnerabilities to use for Impact Assessment — Update the settings as described in Enabling Network Discovery Vulnerability Impact Assessment, on page 1932.

Step 4 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Database Event Limits, on page 879

Network Discovery General Settings

The general settings control how often the system updates network maps and whether server banners are captured during discovery.

Capture Banners

Select this check box if you want the system to store header information from network traffic that advertises server vendors and versions (“banners”). This information can provide additional context to the information gathered. You can access server banners collected for hosts by accessing server details.

Update Interval

The interval at which the system updates information (such as when any of a host’s IP addresses was last seen, when an application was used, or the number of hits for an application). The default setting is 3600 seconds (1 hour).

Note that setting a lower interval for update timeouts provides more accurate information in the host display, but generates more network events.

Configuring Network Discovery General Settings

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Advanced**.

Step 3 Click **Edit** () next to **General Settings**.

Step 4 Update the settings as described in Network Discovery General Settings, on page 1929.

Step 5 Click **Save** to save the general settings.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Network Discovery Identity Conflict Settings

The system determines which operating system and applications are running on a host by matching fingerprints for operating systems and servers against patterns in traffic. To provide the most reliable operating system and server identity information, the system collates fingerprint information from several sources.

The system uses all passive data to derive operating system identities and assign a confidence value.

By default, unless there is an identity conflict, identity data added by a scanner or third-party application overrides identity data detected by the Firepower System. You can use the Identity Sources settings to rank scanner and third-party application fingerprint sources by priority. The system retains one identity for each source, but only data from the highest priority third-party application or scanner source is used as the current identity. Note, however, that user input data overrides scanner and third-party application data regardless of priority.

An identity conflict occurs when the system detects an identity that conflicts with an existing identity that came from either the active scanner or third-party application sources listed in the Identity Sources settings

or from a Firepower System user. By default, identity conflicts are not automatically resolved and you must resolve them through the host profile or by rescanning the host or re-adding new identity data to override the passive identity. However, you can set your system to automatically resolve the conflict by keeping either the passive identity or the active identity.

Generate Identity Conflict Event

Specifies whether the system generates an event when an identity conflict occurs.

Automatically Resolve Conflicts

From the **Automatically Resolve Conflicts** drop-down list, choose one of the following:

- **Disabled** if you want to force manual conflict resolution of identity conflicts
- **Identity** if you want to the system to use the passive fingerprint when an identity conflict occurs
- **Keep Active** if you want the system to use the current identity from the highest priority active source when an identity conflict occurs

Configuring Network Discovery Identity Conflict Resolution

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Advanced**.

Step 3 Click **Edit** ( ) next to **Identity Conflict Settings**.

Step 4 Update the settings in the Edit Identity Conflict Settings pop-up window as described in Network Discovery Identity Conflict Settings, on page 1930.

Step 5 Click **Save** to save the identity conflict settings.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Network Discovery Vulnerability Impact Assessment Options

You can configure how the Firepower System performs impact correlation with intrusion events. Your choices are as follows:

- Check the **Use Network Discovery Vulnerability Mappings** check box if you want to use system-based vulnerability information to perform impact correlation.
- Check the **Use Third-Party Vulnerability Mappings** check box if you want to use third-party vulnerability references to perform impact correlation. For more information, see the *Firepower System Host Input API Guide*.

You can check either or both of the check boxes. If the system generates an intrusion event and the host involved in the event has servers or an operating system with vulnerabilities in the selected vulnerability mapping sets, the intrusion event is marked with the Vulnerable (level 1: red) impact icon. For any servers which do not have vendor or version information, note that you need to enable vulnerability mapping in the Firepower Management Center configuration.

If you clear both check boxes, intrusion events will **never** be marked with the Vulnerable (level 1: red) impact icon.

Related Topics

Mapping Third-Party Vulnerabilities, on page 1849
Mapping Vulnerabilities for Servers, on page 918

Enabling Network Discovery Vulnerability Impact Assessment

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Advanced**.

Step 3 Click **Edit** ( ) next to **Vulnerabilities to use for Impact Assessment**.

Step 4 Update the settings in the Edit Vulnerability Settings pop-up window as described in Network Discovery Vulnerability Impact Assessment Options, on page 1931.

Step 5 Click **Save** to save the vulnerability settings.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Indications of Compromise

The Firepower System uses IOC rules in the network discovery policy to identify a host as likely to be compromised by malicious means. When a host meets the conditions specified in these system-provided rules, the system tags it with an *indication of compromise* (IOC). The related rules are known as *IOC rules*. Each IOC rule corresponds to one type of IOC tag. The *IOC tags* specify the nature of the likely compromise.

The Firepower Management Center can tag the host and user involved when one of the following things occurs:

- The system correlates data gathered about your monitored network and its traffic, using intrusion, connection, Security Intelligence, and file or malware events, and determines that a potential IOC has occurred.
- The Firepower Management Center can import IOC data from your AMP for Endpoints deployments via the AMP cloud. Because this data examines activity on a host itself—such as actions taken by or on individual programs—it can provide insights into possible threats that network-only data cannot. For

your convenience, the Firepower Management Center automatically obtains any new IOC tags that Cisco develops from the AMP cloud.

To configure this feature, see Enabling Indications of Compromise Rules, on page 1933.

You can also write correlation rules against host IOC data and compliance white lists that account for IOC-tagged hosts.

To investigate and work with tagged IOCs, see Indications of Compromise Data, on page 2373 and its subtopics.

Enabling Indications of Compromise Rules

For your system to detect and tag indications of compromise (IOC), you must first activate at least one IOC rule in your network discovery policy. Each IOC rule corresponds to one type of IOC tag, and all IOC rules are predefined by Cisco; you cannot create original rules. You can enable any or all rules, depending on the needs of your network and organization. For example, if hosts using software such as Microsoft Excel never appear on your monitored network, you may decide not to enable the IOC tags that pertain to Excel-based threats.



Tip To disable IOC rules for individual hosts or their associated users, see Editing Indication of Compromise Rule States for a Single Host or User, on page 2376.

Before you begin

Because IOC rules trigger based on data provided by other components of the Firepower System and by AMP for Endpoints, those components must be correctly licensed and configured for IOC rules to set IOC tags. Enable the Firepower System features associated with the IOC rules you will enable, such as intrusion detection and prevention (IPS) and Advanced Malware Protection (AMP). If an IOC rule's associated feature is not enabled, no relevant data is collected and the rule cannot trigger.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Advanced**.

Step 3 Click **Edit** () next to **Indications of Compromise Settings**.

Step 4 To toggle the entire IOC feature off or on, click the slider next to **Enable IOC**.

Step 5 To globally enable or disable individual IOC rules, click the slider in the rule's **Enabled** column.

Step 6 Click **Save** to save your IOC rule settings.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Adding NetFlow Exporters to a Network Discovery Policy

Before you begin

- Configure the NetFlow exporters you plan to use as described in Netflow Data in the Firepower System, on page 1823.
- Review the other NetFlow prerequisites described in Requirements for Using NetFlow Data, on page 1824.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Advanced**.

Step 3 Click **Add** ( ) next to **NetFlow Devices**.

Step 4 In the **IP Address** field, enter the IP address of the network device from which you want the managed device to collect NetFlow data.

Step 5 Optionally:

- Repeat the previous two steps to add additional NetFlow exporters.
- Remove a NetFlow exporter by clicking **Delete** ( ). Keep in mind that if you use a NetFlow exporter in a discovery rule, you must delete the rule before you can delete the device from the Advanced page.

Step 6 Click **Save**.

What to do next

- Configure a network discovery rule to monitor NetFlow traffic as described in Configuring Network Discovery Rules, on page 1920.
- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Network Discovery Data Storage Settings

Discovery data storage settings include the host limit and timeout settings.

When Host Limit Reached

The number of hosts a Firepower Management Center can monitor, and therefore store in the network map, depends on its model. The **When Host Limit Reached** option controls what happens when you detect a new host after you reach the host limit. You can:

Drop hosts

The system drops the host that has remained inactive for the longest time, then adds the new host. This is the default setting.

Don't insert new hosts

The system does not track any newly discovered hosts. The system only tracks new hosts after the host count drops below the limit, such as after an administrator increases the domain's host limit or manually deletes hosts from the network map, or if the system identifies hosts as timed-out due to inactivity.

In a multidomain deployment, leaf domains share the available pool of monitored hosts. To ensure that each leaf domain can populate its network map, you can set host limits at any subdomain level in the domain's properties. Because each leaf domain has its own network discovery policy, each leaf domain governs its own behavior when the system discovers a new host, as described in the following table.

Table 223: Reaching the Host Limit with Multitenancy

Setting	Domain Host Limit Set?	Domain Host Limit Reached	Ancestor Domain Host Limit Reached
Drop hosts	yes	Drops oldest host in the constrained domain.	Drops the oldest host among all descendant leaf domains configured to drop hosts. If no host can be dropped, does not add the host.
	no	n/a	Drops the oldest host among all descendant leaf domains configured to drop hosts and that share the general pool.
Don't insert new hosts	yes or no	Does not add the host.	Does not add the host.

Host Timeout

The amount of time that passes, in minutes, before the system drops a host from the network map due to inactivity. The default setting is 10080 minutes (one week). Individual host IP and MAC addresses can time out individually, but a host does not disappear from the network map unless all its associated addresses time out.

To avoid premature timeout of hosts, make sure that the host timeout value is longer than the update interval in the network discovery policy general settings.

Server Timeout

The amount of time that passes, in minutes, before the system drops a server from the network map due to inactivity. The default setting is 10080 minutes (one week).

To avoid premature timeout of servers, make sure that the service timeout value is longer than the update interval in the network discovery policy general settings.

Client Application Timeout

The amount of time that passes, in minutes, before the system drops a client from the network map due to inactivity. The default setting is 10080 minutes (one week).

Make sure that the client timeout value is longer than the update interval in the network discovery policy general settings.

Related Topics

Firepower System Host Limit, on page 1832
Domain Properties, on page 337

Configuring Network Discovery Data Storage

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Advanced**.

Step 3 Click **Edit** ( ) next to **Data Storage Settings**.

Step 4 Update the settings in the Data Storage Settings dialog as described in Network Discovery Data Storage Settings, on page 1934.

Step 5 Click **Save** to save the data storage settings.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Configuring Network Discovery Event Logging

The Event Logging Settings control whether discovery and host input events are logged. If you do not log an event, you cannot retrieve it in event views or use it to trigger correlation rules.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Advanced**.

Step 3 Click **Edit** ( ) next to **Event Logging Settings**.

Step 4 Check or clear the check boxes next to the discovery and host input event types you want to log in the database, described in Discovery Event Types, on page 2358 and Host Input Event Types, on page 2362.

Step 5 Click **Save** to save the event logging settings.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Adding Network Discovery OS and Server Identity Sources

In Advanced of the network discovery policy, you can add new active sources or change the priority or timeout settings for existing sources.

Adding a scanner to this page does not add the full integration capabilities that exist for the Nmap scanners, but does allow integration of imported third-party application or scan results.

If you import data from a third-party application or scanner, make sure that you map vulnerabilities from the source to the vulnerabilities detected in your network.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Advanced**.

Step 3 Click **Edit** ( ) next to **OS and Server Identity Sources**.

Step 4 To add a new source, click **Add Source**.

Step 5 Enter a **Name**.

Step 6 Choose the input source **Type** from the drop-down list:

- Choose **Scanner** if you plan to import scan results using the AddScanResult function.
- Choose **Application** if you do not plan to import scan results.

Step 7 To indicate the duration of time that should elapse between the addition of an identity to the network map by this source and the deletion of that identity, choose **Hours**, **Days**, or **Weeks** from the **Timeout** drop-down list and enter the appropriate duration.

Step 8 Optionally:

- To promote a source and cause the operating system and application identities to be used in favor of sources below it in the list, choose the source and click the up arrow.
- To demote a source and cause the operating system and application identities to be used only if there are no identities provided by sources above it in the list, choose the source and click the down arrow.
- To delete a source, click **Delete** ( ) next to the source.

Step 9 Click **Save** to save the identity source settings.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Related Topics

Mapping Third-Party Vulnerabilities, on page 1849

Troubleshooting Your Network Discovery Strategy

Before you make any changes to the system's default detection capabilities, you should analyze what hosts are not being identified correctly and why, so you can decide what solution to implement.

Are Your Managed Devices Correctly Placed?

If network devices such as load balancers, proxy servers, or NAT devices reside between the managed device and the unidentified or misidentified host, place a managed device closer to the misidentified host rather than using custom fingerprinting. Cisco does not recommend using custom fingerprinting in this scenario.

Do Unidentified Operating Systems Have a Unique TCP Stack?

If the system misidentifies a host, you should investigate why the host is misidentified to help you decide between creating and activating a custom fingerprint or substituting Nmap or host input data for discovery data.



Caution If you encounter misidentified hosts, contact your support representative before creating custom fingerprints.

If a host is running an operating system that is not detected by the system by default and does not share identifying TCP stack characteristics with existing detected operating systems, you should create a custom fingerprint.

For example, if you have a customized version of Linux with a unique TCP stack that the system cannot identify, you would benefit from creating a custom fingerprint, which allows the system to identify the host and continuing monitoring it, rather than using scan results or third-party data, which require you to actively update the data yourself on an ongoing basis.

Note that many open source Linux distributions use the same kernel, and as such, the system identifies them using the Linux kernel name. If you create a custom fingerprint for a Red Hat Linux system, you may see other operating systems (such as Debian Linux, Mandrake Linux, Knoppix, and so on) identified as Red Hat Linux, because the same fingerprint matches multiple Linux distributions.

You should not use a fingerprint in every situation. For example, a modification may have been made to a host's TCP stack so that it resembles or is identical to another operating system. For example, an Apple Mac OS X host is altered, making its fingerprint identical to a Linux 2.4 host, causing the system to identify it as Linux 2.4 instead of Mac OS X. If you create a custom fingerprint for the Mac OS X host, it may cause all legitimate Linux 2.4 hosts to be erroneously identified as Mac OS X hosts. In this case, if Nmap correctly identifies the host, you could schedule regular Nmap scans for that host.

If you import data from a third-party system using host input, you must map the vendor, product, and version strings that the third party uses to describe servers and application protocols to the Cisco definitions for those products. Note that even if you map application data to Firepower System vendor and version definitions, imported third-party vulnerabilities are not used for impact assessment for clients or web applications.

The system may reconcile data from multiple sources to determine the current identity for an operating system or application.

For Nmap data, you can schedule regular Nmap scans. For host input data, you can regularly run the Perl script for the import or the command line utility. However, note that active scan data and host input data may not be updated with the frequency of discovery data.

Can the Firepower System Identify All Applications?

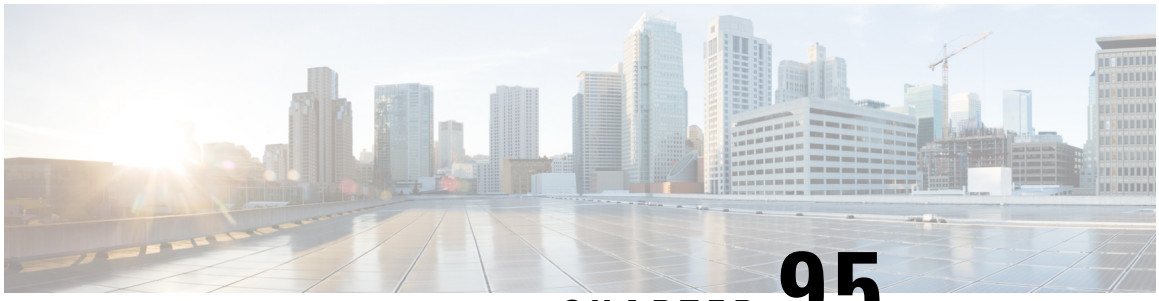
If a host is correctly identified by the system but has unidentified applications, you can create a user-defined detector to provide the system with port and pattern matching information to help identify the application.

Have You Applied Patches that Fix Vulnerabilities?

If the system correctly identifies a host but does not reflect applied fixes, you can use the host input feature to import patch information. When you import patch information, you must map the fix name to a fix in the database.

Do You Want to Track Third-Party Vulnerabilities?

If you have vulnerability information from a third-party system that you want to use for impact correlation, you can map the third-party vulnerability identifiers for servers and application protocols to vulnerability identifiers in the Cisco database and then import the vulnerabilities using the host input feature. For more information on using the host input feature, see the *Firepower System Host Input API Guide*. Note that even if you map application data to Firepower System vendor and version definitions, imported third-party vulnerabilities are not used for impact assessment for clients or web applications.



CHAPTER 95

Realms and Identity Policies

The following topics describe realms and identity policies:

- About Realms and Identity Policies, on page 1941
- License Requirements for Realms, on page 1948
- Requirements and Prerequisites for Realms, on page 1948
- Create a Realm, on page 1949
- Create an Identity Policy, on page 1961
- Create an Identity Rule, on page 1962
- Manage a Realm, on page 1965
- Manage an Identity Policy, on page 1966
- Manage an Identity Rule, on page 1967
- History for Realms, on page 1967

About Realms and Identity Policies

A *realm* consists of one or more LDAP or Microsoft Active Directory servers that share the same directory credentials. You must configure a realm to perform user and user group queries, user control, or to configure an authoritative identity source. After configuring one or more realms, you can configure an identity policy.

An *identity policy* associates traffic on your network with an authoritative identity source and a realm. After configuring one or more identity policies, you can associate one with an access control policy and deploy the access control policy to a managed device.

About Realms

Realms are connections between the Firepower Management Center and the user accounts on the servers you monitor. They specify the connection settings and authentication filter settings for the server. Realms can:

- Specify the users and user groups whose activity you want to monitor.
- Query the user repository for user metadata on authoritative users, as well as some non-authoritative users: POP3 and IMAP users detected by traffic-based detection and users detected by traffic-based detection, a user agent, a TS Agent, or ISE.

You can add multiple domain controllers as directories in a realm, but they must share the same basic realm information. The directories in a realm must be exclusively LDAP or exclusively Active Directory (AD)

servers. After you enable a realm, your saved changes take effect next time the Firepower Management Center queries the server.

To perform user awareness, you must configure a realm for any of the Supported Servers for Realms. The system uses these connections to query the servers for data associated with POP3 and IMAP users, and to collect data about LDAP users discovered through traffic-based detection.

The system uses the email addresses in POP3 and IMAP logins to correlate with LDAP users on an Active Directory, or OpenLDAP. For example, if a managed device detects a POP3 login for a user with the same email address as an LDAP user, the system associates the LDAP user's metadata with that user.

To perform user control, you can configure any of the following:



Note Configuring a realm is optional if you plan to configure SGT ISE attribute conditions but not user, group, realm, Endpoint Location, or Endpoint Profile conditions.

- A realm for an AD server for the TS Agent
- For captive portal, an LDAP realm.

  A realm sequence is not supported for LDAP.

About User Download

You can configure a realm to establish a connection between the Firepower Management Center and an LDAP or AD server to retrieve user and user group metadata for certain detected users:

- LDAP and AD users authenticated by captive portal or reported by a user agent or ISE. This metadata can be used for user awareness and user control.
- POP3 and IMAP user logins detected by traffic-based detection, if those users have the same email address as an LDAP or AD user. This metadata can be used for user awareness.

You configure LDAP server or Active Directory domain controller connections as a directory in a realm. You must check **Download users and user groups for access control** to download a realm's user and user group data for user awareness and user control.

The Firepower Management Center obtains the following information and metadata about each user:

- LDAP user name
- First and last names
- Email address
- Department
- Telephone number

About User Activity Data

User activity data is stored in the user activity database and user identity data is stored in the users database. The maximum number of users you can store and use in access control depends on your Firepower Management Center model. When choosing which users and groups to include, make sure the total number of users is less

than your model limit. If your access control parameters are too broad, the Firepower Management Center obtains information on as many users as it can and reports the number of users it failed to retrieve in the Tasks tab page of the Message Center.

Note If you remove a user that has been detected by the system from your user repository, the Firepower Management Center does *not* remove that user from its users database; you must manually delete it. However, your LDAP changes *are* reflected in access control rules when the Firepower Management Center next updates its list of authoritative users.

Video YouTube video on creating a realm.

Realms and Trusted Domains

When you configure a *realm* in the Firepower Management Center, it is associated with an Active Directory or LDAP *domain*.

A grouping of Microsoft Active Directory (AD) domains that trust each other is commonly referred to as a *forest*. This trust relationship can enable domains to access each other's resources in different ways. For example, a user account defined in domain A can be marked as a member of a group defined in domain B.

The Firepower System and trusted domains

The Firepower System does not support trusted AD domains. This means that the Firepower System does not track which configured domains trust each other, and does not know which domains are parent or child domains of each other. The Firepower System also has not been tested to assure support for environments that use cross-domain trust, even when the trust relationship is exercised outside of the Firepower System.

Supported Servers for Realms

You can configure realms to connect to the following types of servers, providing they have TCP/IP access from the Firepower Management Center:

Server Type	Supported for User Agent data retrieval?	Supported for ISE data retrieval?	Supported for TS Agent data retrieval?	Supported for captive portal data retrieval?
Microsoft Active Directory on Windows Server 2008 and Windows Server 2012	No User agent supported on Windows Server 2008 and 2012 only	Yes	Yes	Yes
OpenLDAP on Linux	No	No	No	Yes

Note the following about your server group configurations:

- To perform user control on user groups or on users in groups, you must configure user groups on the LDAP or Active Directory server.

- Group names cannot start with `s-` because it is used internally by LDAP.

  Neither group names or nor organizational unit names can contain special characters like asterisk (*), equals (=), or backslash (\); otherwise, users in those groups or organizational units are not downloaded and are not available for identity policies.

- To configure an Active Directory realm that includes or excludes users who are members of a sub-group on your server, note that Microsoft recommends that Active Directory has no more than 5000 users per group in Windows Server 2008 or 2012. For more information, see Active Directory Maximum Limits—Scalability on MSDN.

  If necessary, you can modify your Active Directory server configuration to increase this default limit and accommodate more users.

- To uniquely identify the users reported by a server in your Remote Desktop Services environment, you must configure the Cisco Terminal Services (TS) Agent. When installed and configured, the TS Agent assigns unique ports to individual users so the Firepower System can uniquely identify those users. (Microsoft changed the name *Terminal Services* to *Remote Desktop Services*.)

  For more information about the TS Agent, see the *Cisco Terminal Services (TS) Agent Guide*.

Supported Server Object Class and Attribute Names

The servers in your realms *must* use the attribute names listed in the following table for the Firepower Management Center to retrieve user metadata from the servers. If the attribute names are incorrect on your server, the Firepower Management Center cannot populate its database with the information in that attribute.

Table 224: Map of attribute names to Firepower Management Center fields

Metadata	FMC Attribute	LDAP ObjectClass	Active Directory Attribute	OpenLDAP Attribute
LDAP user name	Username	• user • inetOrgPerson	samaccountname	cn uid
first name	First Name		givenname	givenname
last name	Last Name		sn	sn
email address	Email		mail userprincipalname (if mail has no value)	mail
department	Department		department distinguishedname (if department has no value)	ou
telephone number	Phone		telephonenumber	telephonenumber

Note The LDAP ObjectClass for groups is `group`, `groupOfNames`, (`group-of-names` for Active Directory) or `groupOfUniqueNames`.

For more information about ObjectClasses and attributes, see the following references:

- Microsoft Active Directory:
  - ObjectClasses: All Classes on MSDN
  - Attributes: All Attributes on MSDN
- OpenLDAP: RFC 4512

Troubleshoot Realms and User Downloads

If you notice unexpected server connection behavior, consider tuning your realm configuration, device settings, or server settings. For other related troubleshooting information, see:

- Troubleshoot the User Agent Identity Source, on page 1894
- Troubleshoot ISE or Cisco TrustSec Issues, on page 1900
- Troubleshoot the TS Agent Identity Source, on page 1902
- Troubleshoot the Captive Portal Identity Source, on page 1914
- Troubleshoot User Control, on page 387

Symptom: Access control policy doesn't match group membership

This solution applies to an AD domain that is in a trust relationship with other AD domains. In the following discussion, *external domain* means a domain other than the one to which the user logs in.

If a user belongs to a group defined in a trusted external domain, Firepower doesn't track membership in the external domain. For example, consider the following scenario:

- Domain controllers 1 and 2 trust each other
- Group A is defined on domain controller 2
- User `mparvinder` in controller 1 is a member of Group A

Even though user `mparvinder` is in Group A, the Firepower access control policy rules specifying membership Group A don't match.

Solution: Create a similar group in domain controller 1 that contains has all domain 1 accounts that belong to group A. Change the access control policy rule to match any member of Group A or Group B.

Symptom: Access control policy doesn't match child domain membership

If a user belongs to a domain that is child of parent domain, Firepower doesn't track the parent/child relationships between domains. For example, consider the following scenario:

- Domain `child.parent.com` is child of domain `parent.com`
- User `mparvinder` is defined in `child.parent.com`

Even though user `mparvinder` is in a child domain, the Firepower access control policy matching the `parent.com` don't match `mparvinder` in the `child.parent.com` domain.

Solution: Change the access control policy rule to match membership in either `parent.com` or `child.parent.com`.

Symptom: Realm or realm directory test fails

The **Test** button on the directory page sends an LDAP query to the hostname or IP address you entered. If it fails, check the following:

- The **Hostname** you entered resolves to the IP address of an LDAP server or Active Directory domain controller.
- The **IP Address** you entered is valid.

The **Test AD Join** button on the realm configuration page verifies the following:

- DNS resolves the **AD Primary Domain** to an LDAP server or Active Directory domain controller's IP address.
- The **AD Join Username** and **AD Join Password** are correct.

  AD Join Username must be fully qualified (for example, **`administrator@mydomain.com`**, *not* **`administrator`**).
- The user has sufficient privileges to create a computer in the domain and join the Firepower Management Center to the domain as a Domain Computer.

Symptom: User timeouts are occurring at unexpected times

If you notice the system performing user timeouts at unexpected intervals, confirm that the time on your user agent or ISE server is synchronized with the time on the Firepower Management Center. If the appliances are not synchronized, the system may perform user timeouts at unexpected intervals.

If you notice the system performing user timeouts at unexpected intervals, confirm that the time on your ISE, user agent, or TS Agent server is synchronized with the time on the Firepower Management Center. If the appliances are not synchronized, the system may perform user timeouts at unexpected intervals.

Symptom: Users are not included or excluded as specified in your realm configuration

If you configure an Active Directory realm that includes or excludes users who are members of a sub-group on your server, note that Microsoft Windows servers limit the number of users they report:

- 5000 users per group on Microsoft Windows Server 2008 or 2012

If necessary, you can modify your server configuration to increase this default limit and accommodate more users.

Symptom: Users are not downloaded

Possible causes follow:

- If you have the realm **Type** configured incorrectly, users and groups cannot be downloaded because of a mismatch between the attribute the Firepower system expects and what the repository provides. For example, if you configure **Type** as **LDAP** for a Microsoft Active Directory realm, the Firepower system expects the `uid` attribute, which is set to `none` on Active Directory. (Active Directory repositories use `sAMAccountName` for the user ID.)

Solution: Set the realm **Type** field appropriately: **AD** for Microsoft Active Directory or **LDAP** for another supported LDAP repository.

- Users in Active Directory groups that have special characters in the group or organizational unit name might not be available for identity policy rules. For example, if a group or organizational unit name contains the characters asterisk (*), equals (=), or backslash (\), users in those groups are not downloaded and can't be used for identity policies.

  Solution: Remove special characters from the group or organizational unit name.

Symptom: User data for previously-unseen ISE and User Agent users is not displaying in the web interface

After the system detects activity from an ISE, user agent, or TS Agent user whose data is not yet in the database, the system retrieves information about them from the server. In some cases, the system requires additional time to successfully retrieve this information from Microsoft Windows servers. Until the data retrieval succeeds, activity seen by the ISE, User Agent, or TS Agent user is *not* displayed in the web interface.

Note that this may also prevent the system from handling the user's traffic using access control rules.

Symptom: User data in events is unexpected

If you notice user or user activity events contain unexpected IP addresses, check your realms. The system does not support configuring multiple realms with the same **AD Primary Domain** value.

Symptom: Users originating from terminal server logins are not uniquely identified by the system

If your deployment includes a terminal server and you have a realm configured for one or more servers connected to the terminal server, you must deploy the Cisco Terminal Services (TS) Agent to accurately report user logins in terminal server environments. When installed and configured, the TS Agent assigns unique ports to individual users so the Firepower System can uniquely identify those users in the web interface.

For more information about the TS Agent, see the *Cisco Terminal Services (TS) Agent Guide*.

About Identity Policies

Identity policies contain identity rules. Identity rules associate sets of traffic with a realm and an authentication method: passive authentication, active authentication, or no authentication.

With the exception noted in the following paragraphs, you must configure realms and authentication methods you plan to use before you can invoke them in your identity rules:

- You configure realms outside of your identity policy, at **System** > **Integration** > **Realms**. For more information, see Create a Realm, on page 1949.
- You configure the user agent and ISE, passive authentication identity sources, at **System** > **Integration** > **Identity Sources**. For more information, see Configure the User Agent for User Control, on page 1894 and Configure ISE for User Control, on page 1897.
- You configure the TS Agent, a passive authentication identity source, outside the Firepower System. For more information, see the *Cisco Terminal Services (TS) Agent Guide*.
- You configure captive portal, the active authentication identity source, in the identity policy. For more information, see How to Configure the Captive Portal for User Control, on page 1905.

After you add multiple identity rules to a single identity policy, order the rules. The system matches traffic to rules in top-down order by ascending rule number. The first rule that traffic matches is the rule that handles the traffic.

After you configure one or more identity policies, you must associate one identity policy with your access control policy. When traffic on your network matches the conditions in your identity rule, the system associates the traffic with the specified realm and authenticates the users in the traffic using the specified identity source.

If you do not configure an identity policy, the system does not perform user authentication.

Exception to creating an identity policy

An identity policy is not required if all of the following are true:

- You use the ISE/ISE-PIC identity source.
- You do not use users or groups in access control policies.
- You use Security Group Tags (SGT) in access control policies. For more information, see ISE SGT vs Custom SGT Rule Conditions, on page 389.

Video YouTube video on creating an identity policy and rule.

Related Topics

User Identity Sources, on page 1891

License Requirements for Realms

FTD License

Any

Classic License

Control

Requirements and Prerequisites for Realms

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Access Admin

- Network Admin

Create a Realm

For more information about realm and directory configuration fields, see Realm Fields, on page 1950 and Realm Directory and Download fields, on page 1952.

Note You must specify a unique **AD Primary Domain** for every Microsoft Active Directory (AD) realm. Although the system allows you to specify the same **AD Primary Domain** for different AD realms, the system won't function properly. This happens because system assigns a unique ID to every user and group in each *realm*; therefore, the system cannot definitively identify any particular user or group. The system prevents you from specifying more than one realm with the same **AD Primary Domain** because users and groups won't be identified properly. This happens because system assigns a unique ID to every user and group in each *realm*; therefore, the system cannot definitively identify any particular user or group.

If you're setting up ISE/ISE-PIC without a realm, be aware there is a user session timeout that affects how users are seen by the Firepower Management Center. For more information, see Realm Fields, on page 1950.

Procedure

Step 1 Log in to the Firepower Management Center.

Step 2 Click **System** > **Integration**.

Step 3 Click **Realms**.

Step 4 To create a new realm, click **Add Realm**.

Step 5 To perform other tasks (such as enable, disable, or delete a realm), see Manage a Realm, on page 1965.

Step 6 Enter realm information as discussed in Realm Fields, on page 1950.

Step 7 Click **OK**.

Step 8 Configure at least one directory as discussed in Configure a Realm Directory, on page 1959.

Step 9 Configure user and user group download (required for access control) as discussed in Download Users and Groups, on page 1960.

Step 10 Click **Realm Configuration**.

Step 11 Enter user session timeout values, in minutes, for **Authenticated Users**, **Failed Authentication Users**, and **Guest Users**.

Step 12 When you're finished configuring the realm, click **Save**.

What to do next

- Configure a Realm Directory, on page 1959
- Edit, delete, enable, or disable a realm; see Manage a Realm, on page 1965.
- Compare Realms, on page 1966.

- Optionally, monitor the task status; see Viewing Task Messages, on page 316.

Realm Fields

The following fields are used to configure a realm.

Realm Configuration Fields

These settings apply to all Active Directory servers or domain controllers (also referred to as *directories*) in a realm.

Name

A unique name for the realm. The system supports alphanumeric and special characters.

Description

(Optional.) Enter a description of the realm.

Type

The type of realm, **AD** for Microsoft Active Directory or **LDAP** for other supported LDAP repositories. For a list of supported LDAP repositories, see Supported Servers for Realms, on page 1943. You can authenticate captive portal users with an LDAP repository; all others require Active Directory.

Note Only captive portal supports an LDAP realm.

AD Primary Domain

For Microsoft Active Directory realms only. Domain for the Active Directory server where users should be authenticated.

Note You must specify a unique **AD Primary Domain** for every Microsoft Active Directory (AD) realm. Although the system allows you to specify the same **AD Primary Domain** for different AD realms, the system won't function properly. This happens because system assigns a unique ID to every user and group in each *realm*; therefore, the system cannot definitively identify any particular user or group. The system prevents you from specifying more than one realm with the same **AD Primary Domain** because users and groups won't be identified properly. This happens because system assigns a unique ID to every user and group in each *realm*; therefore, the system cannot definitively identify any particular user or group.

AD Join Username and AD Join Password

For Microsoft Active Directory realms intended for Kerberos captive portal active authentication, the distinguished username and password of any Active Directory user with appropriate rights to create a Domain Computer account in the Active Directory domain.

Keep the following in mind:

- DNS must be able to resolve the domain name to an Active Directory domain controller's IP address.
- The user you specify must be able to join computers to the Active Directory domain.

- The user name must be fully qualified (for example, **`administrator@mydomain.com`**, *not* **`administrator`**).

If you choose **`Kerberos`** (or **`HTTP Negotiate`**, if you want Kerberos as an option) as the **Authentication Type** in an identity rule, the **Realm** you select must be configured with an **AD Join Username** and **AD Join Password** to perform Kerberos captive portal active authentication.



Note The SHA-1 hash algorithm is not secure for storing passwords on your Active Directory server and should not be used. For more information, consult a reference such as Migrating your Certification Authority Hashing Algorithm from SHA1 to SHA2 on Microsoft TechNet or Password Storage Cheat Sheet on the Open Web Application Security Project website.

Directory Username and Directory Password

The distinguished username and password for a user with appropriate access to the user information you want to retrieve.

Note the following:

- For Microsoft Active Directory, the user does not need elevated privileges. You can specify any user in the domain.
- For OpenLDAP, the user's access privileges are determined by the `<level>` parameter discussed in section 8 of the OpenLDAP specification. The user's `<level>` should be `auth` or better.
- The user name must be fully qualified (for example, **`administrator@mydomain.com`**, *not* **`administrator`**).



Note The SHA-1 hash algorithm is not secure for storing passwords on your Active Directory server and should not be used. For more information, consult a reference such as Migrating your Certification Authority Hashing Algorithm from SHA1 to SHA2 on Microsoft TechNet or Password Storage Cheat Sheet on the Open Web Application Security Project website.

Base DN

The directory tree on the server where the Firepower Management Center should begin searching for user data.

Typically, the base distinguished name (DN) has a basic structure indicating the company domain name and operational unit. For example, the Security organization of the Example company might have a base DN of **`ou=security,dc=example,dc=com`**.

Group DN

The directory tree on the server where the Firepower Management Center should search for users with the group attribute. A list of supported group attributes is shown in Supported Server Object Class and Attribute Names, on page 1944.

Note Following is the list of characters the Firepower System *supports* in users, groups, DNs in your directory server. Using any characters other than the following could result in the Firepower System failing to download users and groups.

Entity	Supported characters
User name	**a-z A-Z 0-9 ! # $ % ^ & ( ) _ - { } ' . ~ `**
Group name	**a-z A-Z 0-9 ! # $ % ^ & ( ) _ - { } ' . ~ `**
Base DN and Group DN	**a-z A-Z 0-9 ! @ $ % ^ & * ( ) _ - . ~ ` [ ]**

Group Attribute

(Optional.) The group attribute for the server, **`Member`** or **`Unique Member`**.

The following fields are available when you edit an existing realm.

User Session Timeout

Enter the number of minutes before user sessions time out. The default is 1440 (24 hours) after the user's login event. After the timeout is exceeded, the user's session ends; if the user continues to access the network without logging in again, the user is seen by the Firepower Management Center as Unknown.

Note The user session timeout values apply to both active authentication (captive portal) and passive authentication (TS Agent, user agent, ISE). Setting a large value might prevent user sessions from ending, resulting in those sessions being claimed by other users.

Realm Directory and Download fields

Realm Directory Fields

These settings apply to individual servers (such as Active Directory domain controllers) in a realm.

Hostname / IP Address

Fully qualified host name of the Active Directory domain controller machine. To find the fully qualified name, see Find the Active Directory Server's Name, on page 1954.

Port

The port to use for the Firepower Management Center-controller connection.

Encryption

(Strongly recommended.) The encryption method to use for the Firepower Management Center-server connection:

- **`STARTTLS`**—encrypted LDAP connection
- **`LDAPS`**—encrypted LDAP connection
- **`None`**—unencrypted LDAP connection (unsecured traffic)

To communicate securely with an Active Directory server, see Connect Securely to Active Directory, on page 1954.

SSL Certificate

The SSL certificate to use for authentication to the server. You must configure **STARTTLS** or **LDAPS** as the **Encryption** type in order to use an SSL certificate.

If you are using a certificate to authenticate, the name of the server in the certificate must match the server **Hostname / IP Address**. For example, if you use 10.10.10.250 as the IP address but **computer1.example.com** in the certificate, the connection fails.

User Download Fields

AD Primary Domain

For Microsoft Active Directory realms only. Domain for the Active Directory server where users should be authenticated.



Note You must specify a unique **AD Primary Domain** for every Microsoft Active Directory (AD) realm. Although the system allows you to specify the same **AD Primary Domain** for different AD realms, the system won't function properly. This happens because system assigns a unique ID to every user and group in each *realm*; therefore, the system cannot definitively identify any particular user or group. The system prevents you from specifying more than one realm with the same **AD Primary Domain** because users and groups won't be identified properly. This happens because system assigns a unique ID to every user and group in each *realm*; therefore, the system cannot definitively identify any particular user or group.

Download users and groups (required for user access control)

Enables you to download users and groups for user awareness and user control.

Begin automatic download at, Repeat every

Specifies the frequency of the automatic downloads.

Download Now

Click to synchronize groups and users with AD.

Available Groups, Add to Include, Add to Exclude

Limits the groups that can be used in policy.

- Groups that are displayed in the **Available Groups** field are available for policy unless you move groups to the **Add to Include** or **Add to Exclude** field.
- If you move groups to the **Add to Include** field, only those groups are downloaded and user data is available for user awareness and user control.
- If you move groups to the **Add to Exclude** field, all groups *except* these are downloaded and available for user awareness and user control.
- To include users from groups that are not included, enter the user name in the field below **Groups to Include** and click **Add**.
- To exclude users from groups that are not excluded, enter the user name in the field below **Groups to Exclude** and click **Add**.



Note The users that are downloaded to the Firepower Management Center is calculated using the formula **R = I - (E+e) + i**, where

- R is list of downloaded users
- I is included groups
- E is excluded groups
- e is excluded users
- i is included users

Begin automatic download at
Enter the time and time interval at which to download users and groups from AD.

Realms and Identity Policies

Connect Securely to Active Directory

To create a secure connection between an Active Directory server and the FMC (which we strongly recommend), you must perform all of the following tasks:

- Export the Active Directory server's root certificate.
- Import the root certificate into the FMC as a trusted CA certificate.
- Find the Active Directory server's fully qualified name.
- Create the realm directory.

See one of the following tasks for more information.

Related Topics

Export the Active Directory Server's Root Certificate, on page 1955
Find the Active Directory Server's Name, on page 1954
Configure a Realm Directory, on page 1959

Find the Active Directory Server's Name

To configure a realm directory in the FMC, you must know the fully qualified server name, which you can find as discussed in the procedure that follows.

Before you begin

You must log in to the Active Directory server as a user with sufficient privileges to view the computer's name.

Procedure

Step 1 Log in to the Active Directory server.

Step 2 Click **Start**.

Step 3 Right-click **This PC**.

Step 4 Click **Properties**.

Step 5 Click **Advanced System Settings**.

Step 6 Click the **Computer Name** tab.

Step 7 Note the value of **Full computer name**.
You must enter this exact name when you configure the realm directory in the FMC.

What to do next

Create a realm directory.

Related Topics

Export the Active Directory Server's Root Certificate, on page 1955

Export the Active Directory Server's Root Certificate

The task that follows discusses how to export the Active Directory server's root certificate, which is required to connect securely to the FMC to obtain user identity information.

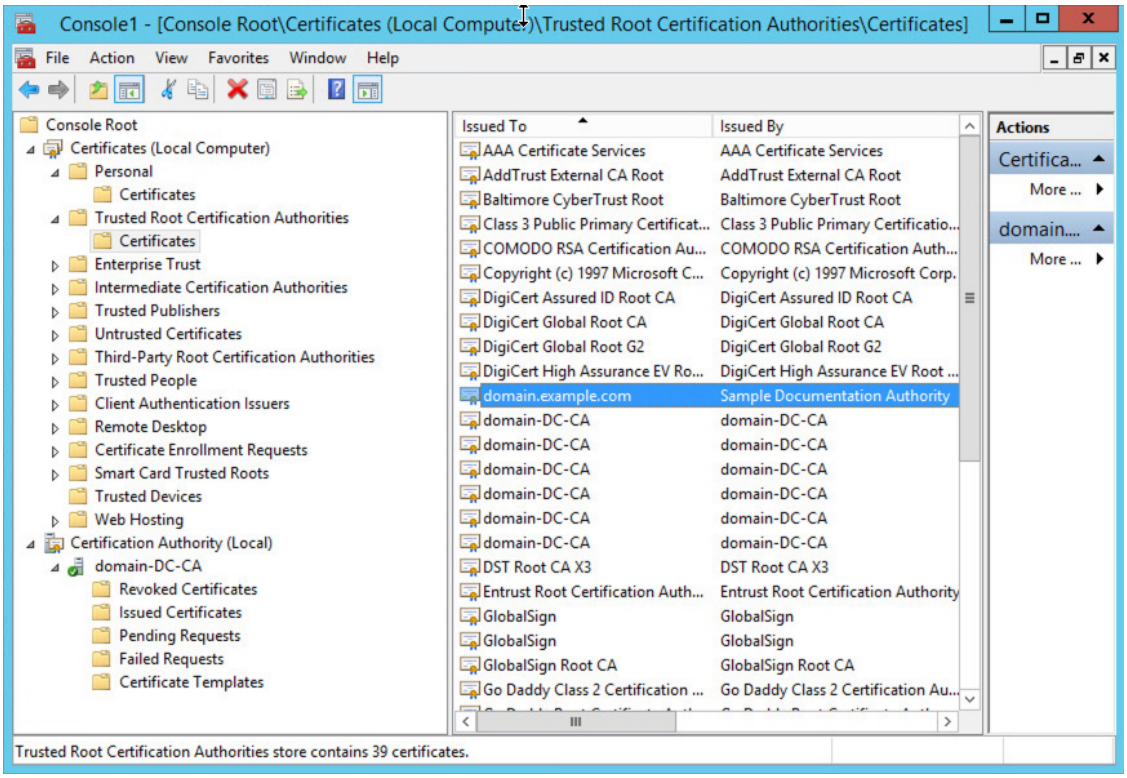
Before you begin

You must know the name of your Active Directory server's root certificate. The root certificate might have the same name as the domain or the certificate might have a different name. The procedure that follows shows one way you can find the name; there could be other ways, however.

Procedure

Step 1 Following is one way to find the name of the Active Directory Server's root certificate; consult Microsoft documentation for more information:

a) Log in to the Active Directory server as a user with privileges to run the Microsoft Management Console.
b) Click **Start** and enter **mmc**.
c) Click **File** > **Add/Remove Snap-in**
d) From the Available Snap-ins list in the left pane, click **Certificates (local)**.
e) Click **Add**.
f) At the Certificates snap-in dialog box, click **Computer Account** and click **Next**.
g) At the Select Computer dialog box, click **Local Computer** and click **Finish**.
h) *Windows Server 2012 only*. Repeat the preceding steps to add the Certification Authority snap-in.
i) Click **Console Root** > **Trusted Certification Authorities** > **Certificates**.
The server's trusted certificates are displayed in the right pane. The following figure is only an example for Windows Server 2012; yours will probably look different.



Step 2 Export the certificate using the **certutil** command.

This is only one way to export the certificate. It's a convenient way to export the certificate, especially if you can run a web browser and connect to the FMC from the Active Directory server.

a) Click **Start** and enter **cmd**.
b) Enter the command **certutil -ca.cert** *certificate-name*.
The server's certificate is displayed on the screen.
c) Copy the entire certificate to the clipboard, starting with **-----BEGIN CERTIFICATE-----** and ending with **-----END CERTIFICATE-----** (including those strings).

What to do next

Import the Active Directory server's certificate into the FMC as a Trusted CA Certificate as discussed in Adding a Trusted CA Object, on page 453.

Related Topics

Find the Active Directory Server's Name, on page 1954

Export the Active Directory Server's Root Certificate

The task that follows discusses how to export the Active Directory server's root certificate, which is required to connect securely to the FMC to obtain user identity information.

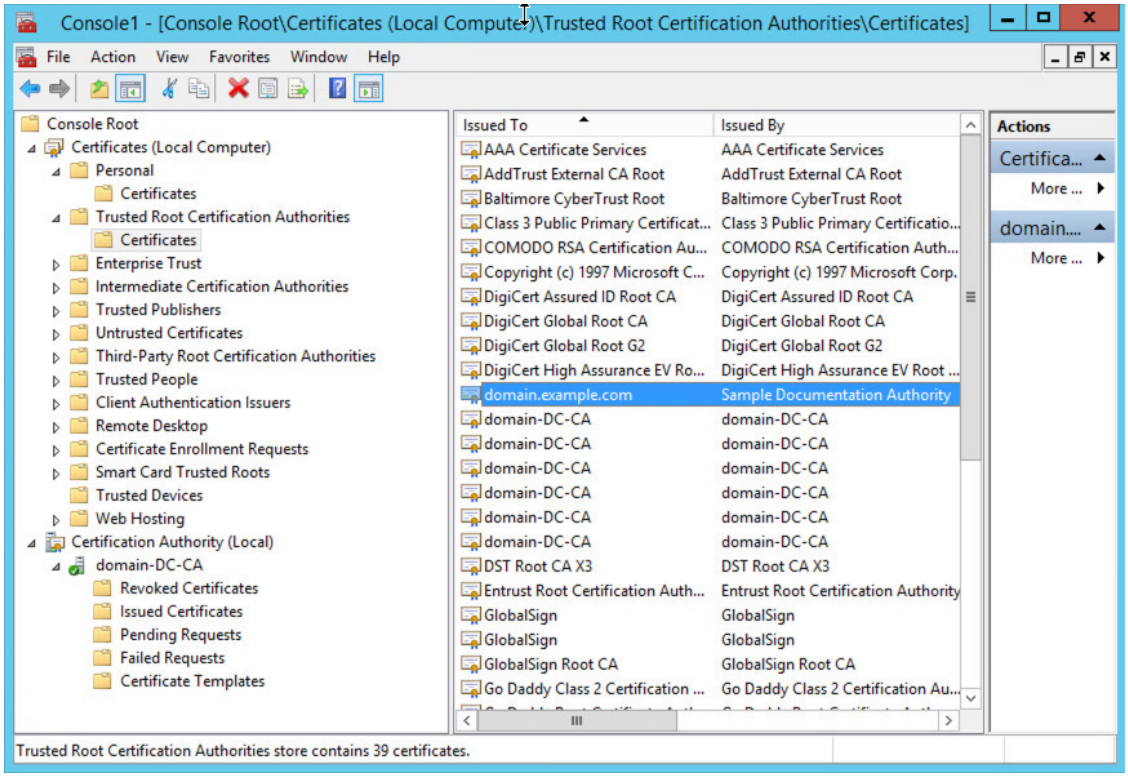
Before you begin

You must know the name of your Active Directory server's root certificate. The root certificate might have the same name as the domain or the certificate might have a different name. The procedure that follows shows one way you can find the name; there could be other ways, however.

Procedure

Step 1 Following is one way to find the name of the Active Directory Server's root certificate; consult Microsoft documentation for more information:

a) Log in to the Active Directory server as a user with privileges to run the Microsoft Management Console.
b) Click **Start** and enter **mmc**.
c) Click **File** > **Add/Remove Snap-in**
d) From the Available Snap-ins list in the left pane, click **Certificates (local)**.
e) Click **Add**.
f) At the Certificates snap-in dialog box, click **Computer Account** and click **Next**.
g) At the Select Computer dialog box, click **Local Computer** and click **Finish**.
h) *Windows Server 2012 only*. Repeat the preceding steps to add the Certification Authority snap-in.
i) Click **Console Root** > **Trusted Certification Authorities** > **Certificates**.
The server's trusted certificates are displayed in the right pane. The following figure is only an example for Windows Server 2012; yours will probably look different.



Step 2 Export the certificate using the **certutil** command.

This is only one way to export the certificate. It's a convenient way to export the certificate, especially if you can run a web browser and connect to the FMC from the Active Directory server.

a) Click **Start** and enter **cmd**.
b) Enter the command **certutil -ca.cert** *certificate-name*.
The server's certificate is displayed on the screen.
c) Copy the entire certificate to the clipboard, starting with **-----BEGIN CERTIFICATE-----** and ending with **-----END CERTIFICATE-----** (including those strings).

What to do next

Import the Active Directory server's certificate into the FMC as a Trusted CA Certificate as discussed in Adding a Trusted CA Object, on page 453.

Related Topics

Find the Active Directory Server's Name, on page 1954

Find the Active Directory Server's Name

To configure a realm directory in the FMC, you must know the fully qualified server name, which you can find as discussed in the procedure that follows.

Before you begin

You must log in to the Active Directory server as a user with sufficient privileges to view the computer's name.

Procedure

Step 1 Log in to the Active Directory server.
Step 2 Click **Start**.
Step 3 Right-click **This PC**.
Step 4 Click **Properties**.
Step 5 Click **Advanced System Settings**.
Step 6 Click the **Computer Name** tab.
Step 7 Note the value of **Full computer name**.
You must enter this exact name when you configure the realm directory in the FMC.

What to do next

Create a realm directory.

Related Topics

Export the Active Directory Server's Root Certificate, on page 1955

Configure a Realm Directory

This procedure enables you to create a realm directory, which corresponds to an LDAP server or a Microsoft Active Directory domain controller. An Active Directory server can have multiple domain controllers, each of which is capable of authenticating different users and groups.

Microsoft has announced that Active Directory servers will start enforcing LDAP binding and LDAP signing in 2020. Microsoft is making these a requirement because when using default settings, an elevation of privilege vulnerability exists in Microsoft Windows that could allow a man-in-the-middle attacker to successfully forward an authentication request to a Windows LDAP server. For more information, see 2020 LDAP channel binding and LDAP signing requirement for Windows on the Microsoft support site.

If you have not done so already, we recommend you start using TLS/SSL encryption to authenticate with an Active Directory server.

An Active Directory Global Catalog server is *not supported* as a realm directory. For more information about the Global Catalog Server, see Global Catalog on learn.microsoft.com.

For more information about realm directory configuration fields, see Realm Fields, on page 1950.

Before you begin

(Recommended.) To connect securely from the FMC to your Active Directory server, first perform the following tasks:

- Export the Active Directory Server's Root Certificate, on page 1955
- Find the Active Directory Server's Name, on page 1954

Procedure

Step 1 If you haven't done so already, log in to the Firepower Management Center and click **System** > **Integration** > **Realms**.

Step 2 On Realms page, click the name of the realm for which to configure a directory.

Step 3 On Directory page, click **Add Directory**.

Step 4 Enter the **Hostname / IP Address** and **Port** for the LDAP server or Active Directory domain controller. The system sends an LDAP query to the hostname or IP address you specify. If the host name resolves to the IP address of an LDAP server or Active Directory domain controller, the **Test** succeeds.

Step 5 Select an **Encryption Mode**.

Step 6 Choose an **SSL Certificate** from the list or click **Add** () to add a certificate.

Step 7 To test the connection, click **Test**.

Step 8 Click **OK**.

Step 9 Click **Save**. You are returned to Realms page

Step 10 If you haven't already enabled the realm, on Realms page, slide **State** to enabled.

What to do next

- Download Users and Groups, on page 1960.

Related Topics

Export the Active Directory Server's Root Certificate, on page 1955
Find the Active Directory Server's Name, on page 1954

Download Users and Groups

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	Control	Any	Any	Administrator, Access Admin, Network Admin

This section discusses how to download users and groups from your Active Directory server to the Firepower Management Center. If you do not specify any groups to include, the system retrieves user data for all the groups that match the parameters you provided. For performance reasons, Cisco recommends that you explicitly include only the groups that represent the users you want to use in access control.

The maximum number of users the Firepower Management Center can retrieve from the server depends on your Firepower Management Center model. If the download parameters in your realm are too broad, the Firepower Management Center obtains information on as many users as it can and reports the number of users it failed to retrieve in Task of the Message Center.

Note User names that include Unicode characters do not display in the Firepower Management Center. Before you download users and groups, make sure to replace Unicode characters with alphanumeric characters.

For more information about realm configuration fields, see Realm Fields, on page 1950.

Procedure

Step 1 Log in to the Firepower Management Center.

Step 2 Click **System** > **Integration** > **Realms**.

Step 3 To download users and groups manually, click **Download** ( ) next to the realm to download users and user groups. If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration. You can skip the remainder of this procedure.

Step 4 To configure the realm for automatic user and group download, click **Edit** ( ) next to the realm to configure for automatic user and group download.

Step 5 On User Access Control page, check **Download users and groups (required for user access control)**.

Step 6 Select a time to **Begin automatic download at** from the lists.

Step 7 Select a download interval from the **Repeat Every** list.

Step 8 To include or exclude user groups from the download, choose user groups from the **Available Groups** column and click **Add to Include** or **Add to Exclude**.

Separate multiple users with commas. You can also use an asterisk (*) as a wildcard character in this field.

Note You must **Add to Include** if you want to perform user control on users in that group.

Use the following guidelines:

- If you leave a group in the **Available Groups** box, the group is not downloaded.
- If you move a group to the **Add to Include** box, the group is downloaded and user data is available for user awareness and user control.
- If you move a group to the **Add to Exclude** box, the group is downloaded and user data is available for user awareness, but not for user control.
- To include users from groups that are not included, enter the user name in the field below **Groups to Include** and click **Add**.
- To exclude users from groups that are not excluded, enter the user name in the field below **Groups to Exclude** and click **Add**.

Create an Identity Policy

Before you begin

An identity policy is required to use users and groups in a realm in access control policies. Create and enable one or more realms as described in Create a Realm, on page 1949.

An identity policy is not required if all of the following are true:

- You use the ISE/ISE-PIC identity source.
- You do not use users or groups in access control policies.
- You use Security Group Tags (SGT) in access control policies. For more information, see ISE SGT vs Custom SGT Rule Conditions, on page 389.

Procedure

Step 1 Log in to the Firepower Management Center.

Step 2 Click **Policies** > **Access Control** > **Identity** and click **New Policy**.

Step 3 Enter a **Name** and, optionally, a **Description**.

Step 4 Click **Save**.

Step 5 To add a rule to the policy, click **Add Rule** as described in Create an Identity Rule, on page 1962.

Step 6 To create a rule category, click **Add Category**.

Step 7 To configure captive portal active authentication, click **Active Authentication** as described in Configure the Captive Portal Part 1: Create an Identity Policy, on page 1907.

Step 8 Click **Save** to save the identity policy.

What to do next

- Add rules to your identity policy that specify which users to match and other options; see Create an Identity Rule, on page 1962.
- Associate the identity policy with an access control policy to allow or block selected users from accessing specified resources; see Associating Other Policies with Access Control, on page 1213.
- Deploy configuration changes to managed devices; see Deploy Configuration Changes, on page 346.

If you encounter issues, see Troubleshoot User Control, on page 387.

Create an Identity Rule

For details about configuration options for identity rules, see Identity Rule Fields, on page 1963.

Before you begin

You must create and enable a realm.

- Create a realm as discussed in Create a Realm, on page 1949.
- Create a directory as discussed in Configure a Realm Directory, on page 1959.
- Download users and groups and enable the realm as discussed in Download Users and Groups, on page 1960.

Procedure

Step 1 If you haven't done so already, log in to the Firepower Management Center.

Step 2 Click **Policies** > **Access Control** > **Identity** .

Step 3 Click **Edit** ( ) next to the identity policy to which to add the identity rule.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Click **Add Rule**.

Step 5 Enter a **Name**.

Step 6 Specify whether the rule is **Enabled**.

Step 7 To add the rule to an existing category, indicate where you want to **Insert** the rule. To add a new category, click **Add Category**.

Step 8 Choose a rule **Action** from the list.

Step 9 Click **Realms & Settings**.

Step 10 Choose a realm for the identity rule from the **Realms** list. You must associate a realm with every identity rule.

The only exception to the realm requirement is implementing user control using only the ISE SGT attribute tag. In this case, you do not need to configure a realm for the ISE server. ISE SGT attribute conditions can be configured in policies with or without an associated identity policy.

Step 11 If you're configuring captive portal, see How to Configure the Captive Portal for User Control, on page 1905.

Step 12 (Optional) To add conditions to the identity rule, see Rule Condition Types, on page 363.

Step 13 Click **Add**.

Step 14 In the policy editor, set the rule position. Click and drag or use the right-click menu to cut and paste. Rules are numbered starting at 1. The system matches traffic to rules in top-down order by ascending rule number. The first rule that traffic matches is the rule that handles that traffic. Proper rule order reduces the resources required to process network traffic and prevents rule preemption.

Step 15 Click **Save**.

Related Topics

Snort® Restart Scenarios, on page 348

Identity Rule Fields

Use the following fields to configure identity rules.

Enabled

Choosing this option enables the identity rule in the identity policy. Deselecting this option disables the identity rule.

Action

Specify the type of authentication you want to perform on the users in the specified realm: **Passive Authentication** (default), **Active Authentication**, or **No Authentication**. You must fully configure the authentication method, or *identity source*, before selecting it as the action in an identity rule.

Caution Adding the first or removing the last active authentication rule when SSL decryption is disabled (that is, when the access control policy does not include an SSL policy) restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information.

Note that an active authentication rule has either an **Active Authentication** rule action, or a **Passive Authentication** rule action with **Use active authentication if passive authentication cannot identify user** selected.

For information about which passive and active authentication methods are supported in your version of the Firepower System, see About User Identity Sources, on page 1891.

Realm

The realm containing the users you want to perform the specified **Action** on. You must fully configure a realm before selecting it as the realm in an identity rule.

Note If you select `Kerberos` (or `HTTP Negotiate`, if you want Kerberos as an option) as the **Authentication Type** for the identity rule, the **Realm** you select must be configured with an **AD Join Username** and **AD Join Password** to perform Kerberos captive portal active authentication.

Use active authentication if passive authentication cannot identify user

Selecting this option authenticates users using captive portal active authentication if a passive or a VPN authentication fails to identify them. You must configure captive portal active authentication in your identity policy in order to select this option.

If you disable this option, users that do not have a VPN identity or that passive authentication cannot identify are identified as Unknown.

Identify as Special Identities/Guest if authentication cannot identify user

This field is displayed only if you configure **Active Authentication** (that is, captive portal authentication) as the rule **Action**.

Authentication Type

The method to use to perform captive portal active authentication. The selections vary depending on the type of realm, LDAP or AD:

- Choose **HTTP Basic** if you want to authenticate users using an unencrypted HTTP Basic Authentication (BA) connection. Users log in to the network using their browser's default authentication popup window.

  Most web browsers cache the credentials from **HTTP Basic** logins and use the credentials to seamlessly begin a new session after an old session times out.

- Choose **NTLM** to authenticate users using a NT LAN Manager (NTLM) connection. This selection is available only when you select an AD realm. If transparent authentication is configured in a user's browser, the user is automatically logged in. If transparent authentication is not configured, users log in to the network using their browser's default authentication popup window.
- Choose **Kerberos** to authenticate users using a Kerberos connection. This selection is available only when you select an AD realm for a server with secure LDAP (LDAPS) enabled. If transparent authentication is configured in a user's browser, the user is automatically logged in. If transparent authentication is not configured, users log in to the network using their browser's default authentication popup window.



Note The **Realm** you select must be configured with an **AD Join Username** and **AD Join Password** to perform Kerberos captive portal active authentication.



Note If you are creating an identity rule to perform Kerberos captive portal and you have DNS resolution configured, you must configure your DNS server to resolve the fully qualified domain name (FQDN) of the captive portal device. The FQDN must match the host name you provided when configuring DNS.

For ASA with FirePOWER Services and Firepower Threat Defense devices, the FQDN must resolve to the IP address of the routed interface used for captive portal.

- Choose **HTTP Negotiate** to allow the captive portal server to choose between HTTP Basic, Kerberos, or NTLM for the authentication connection. This type is available only when you select an AD realm.

Note The **Realm** you choose must be configured with an **AD Join Username** and **AD Join Password** for **HTTP Negotiate** to choose Kerberos captive portal active authentication.

Note If you are creating an identity rule to perform **HTTP Negotiate** captive portal and you have DNS resolution configured, you must configure your DNS server to resolve the fully qualified domain name (FQDN) of the captive portal device. The FQDN of the device you are using for captive portal must match the hostname you provided when configuring DNS.

For ASA with FirePOWER Services devices, the FQDN is the FQDN of the ASA FirePOWER module.

Manage a Realm

This section discusses how to perform various maintenance tasks for a realm using controls on the Realms page. Note the following:

- If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Procedure

Step 1 Log in to the Firepower Management Center.

Step 2 Click **System** > **Integration**.

Step 3 Click **Realms**.

Step 4 To delete a realm, click **Delete** ( ).

Step 5 To edit a realm, click **Edit** ( ) next to the realm and make changes as described in Create a Realm, on page 1949.

Step 6 To enable a realm, slide **State** to the right; to disable a realm, slide it to the left.

Step 7 To download users and user groups, click **Download** ( ).

Step 8 To copy a realm, click **Copy** ( ).

Step 9 To compare realms, see Compare Realms, on page 1966.

Compare Realms

You must be an Admin, Access Admin, Network Admin, or Security Approver to perform this task.

Procedure

Step 1	Log in to the Firepower Management Center.
Step 2	Click **System** > **Integration**.
Step 3	Click **Realms**.
Step 4	Click **System** > **Integration**.
Step 5	Click **Realms**.
Step 6	Click **Compare Realms**.
Step 7	Choose **Compare Realm** from the **Compare Against** list.
Step 8	Choose the realms you want to compare from the **Realm A** and **Realm B** lists.
Step 9	Click **OK**.
Step 10	To navigate individually through changes, click **Previous** or **Next** above the title bar.
Step 11	(Optional.) Click **Comparison Report** to generate the realm comparison report.
Step 12	(Optional.) Click **New Comparison** to generate a new realm comparison view.

Manage an Identity Policy

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Procedure

Step 1	If you haven't done so already, log in to the Firepower Management Center.
Step 2	Click **Policies** > **Access Control** > **Identity** .
Step 3	To delete a policy, click **Delete** ( ). If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
Step 4	To edit a policy, click **Edit** ( ) next to the policy and make changes as described in Create an Identity Policy, on page 1961. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
Step 5	To copy a policy, click **Copy** ( ).
Step 6	To generate a report for the policy, click **Report** ( ) as described in Generating Current Policy Reports, on page 356.

Step 7 To compare policies, see Comparing Policies, on page 355.

Manage an Identity Rule

Procedure

Step 1 If you haven't already done so, log in to the Firepower Management Center.

Step 2 Click **Policies** > **Access Control** > **Identity** .

Step 3 Click **Edit** ( ) next to the policy you want to edit. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 To edit an identity rule, click **Edit** ( ) and make changes as described in Create an Identity Policy, on page 1961.

Step 5 To delete an identity rule, click **Delete** ( ).

Step 6 To create a rule category, click **Add Category** and choose the position and the rule.

Step 7 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

History for Realms

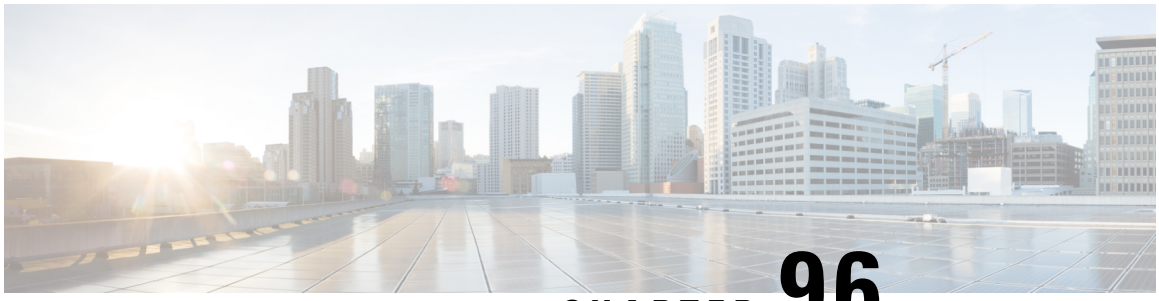
Feature	Version	Details
Realms for user control.	—	Feature introduced before Version 6.0. A realm is a connection between the FMC either an Active Directory or LDAP user repository.



PART XXI

Correlation and Compliance

- Compliance White Lists, on page 1971
- Correlation Policies, on page 1985
- Traffic Profiling, on page 2021
- Remediations, on page 2033



CHAPTER 96

Compliance White Lists

The following topics describe how to configure compliance white lists before you add them to correlation policies.

- Introduction to Compliance White Lists, on page 1971
- Requirements and Prerequisites for Compliance, on page 1976
- Creating a Compliance White List, on page 1976
- Managing Compliance White Lists, on page 1982
- Managing Shared Host Profiles, on page 1984

Introduction to Compliance White Lists

A *compliance white list*, sometimes abbreviated as a *white list*, is a set of criteria that specifies which operating systems, applications (web and client), and protocols are allowed on hosts on your network. The system generates an event (violation) if a host is not on this list.

A compliance white list has two main components:

- *Targets* are the hosts you select for compliance evaluation. You can evaluate all or some monitored hosts, constraining by subnet, VLAN, and host attribute. In a multidomain deployment, you can target domains and subnets within or across domains.
- *Host profiles* specify the compliance criteria for the targets. The global host profile is operating system agnostic. You can also configure operating-system specific host profiles, either unique to one white list or shared across multiple white lists.

The Cisco Talos Intelligence Group (Talos) provides a default white list with recommended settings. You can also create custom white lists. A simple custom list might allow only hosts running a certain operating system. A more complex list might allow all operating systems, but specify which operating system a host must use to run a certain application protocol on a specific port.

Note The system can add hosts to the network map from exported NetFlow records, but the available information for these hosts is limited; see Differences between NetFlow and Managed Device Data, on page 1824. This limitation may affect the way you build compliance white lists.

Implementing Compliance White Lists

To implement white lists, add the list to an active correlation policy. The system evaluates the targets and assigns every host a corresponding attribute:

- Compliant — The host does not violate the list.
- Non-Compliant — The host violates the list.
- Not Evaluated — The host is not a target of the list, the host is currently being evaluated, or the system has insufficient information to determine whether the host is in compliance.



Note To delete the host attribute, delete its corresponding white list. Deactivating, deleting, or removing a white list from a correlation policy does **not** delete the host attribute, nor does it change the attribute's value for each host.

After its initial evaluation, the system generates a *white list event* whenever a monitored host goes out of compliance with an active white list; it also records a *white list violation*.

You can use workflows, dashboards, and network maps to monitor system-wide compliance activity and determine when and how an individual host violates your white lists. You can also automatically respond to such violations with remediations and alerts.

Example: Restricting HTTP to Web Servers

Your security policy states that only web servers may run HTTP. You create a white list that evaluates your entire network, excluding your web farm, to determine which hosts are running HTTP.

Using the network map and the dashboard, you can obtain an at-a-glance summary of the compliance of your network. In just a few seconds, you can determine exactly which hosts in your organization are running HTTP in violation of your policy, and take appropriate action.

Then, using the correlation feature, you can configure the system to alert you whenever a host that is not in your web farm starts running HTTP.

Related Topics

Configuring Correlation Policies, on page 1987

Compliance White List Target Networks

A *target network* specifies the hosts you want to evaluate for compliance. A white list can have more than one target network, and it evaluates hosts that meet the criteria of any of its targets.

Initially, you constrain a target network by IP address or range. In multidomain deployments, the initial constraints also include a domain.

The system-provided default white list targets all monitored hosts: 0.0.0.0/0 and ::/0. In a multidomain deployment, the default white list is constrained to (and only available in) the Global domain.

If you modify a target network or a host so that the host is no longer a valid target for the white list, the host is no longer evaluated by the list and is considered neither compliant nor non-compliant.

Surveying and Refining Target Networks

When you add a target network to a white list, the system prompts you to survey the network map to help you characterize compliant hosts. The survey adds a target to the white list that represents the hosts you surveyed.

You can survey a subnet or individual host. In a multidomain deployment, you can survey an entire domain, or you can survey across domains. Surveying an ancestor domain causes the system to survey that domain's descendants.

In addition to the added target, the survey also populates the white list with one host profile for each operating system detected in the survey. These host profiles allow all the clients, application protocols, web applications, and protocols that the system has detected on the applicable operating systems.

After you survey a target network (or skip the survey), refine the target. You can exclude hosts by IP address, or constrain target networks by host attribute or VLAN.

Targeting Domains with Compliance White Lists

In a multidomain deployment, domains and target networks are closely linked.

- Leaf-domain administrators can create white lists that evaluate hosts within their leaf domains.
- Higher-level domain administrators can create white lists that evaluate hosts across domains. You can target different subnets in different domains in the same white list.

Consider a scenario where you are a Global domain administrator, and you want to apply the same compliance criteria to web servers across the entire deployment. You can create one white list in the Global domain that defines the compliance criteria. Then, constrain the white list with target networks that specify the IP space (or individual IP addresses) of the web servers in each leaf domain.



Note In addition to targeting IP addresses and ranges in leaf domains, you can also constrain a target network using a higher-level domain. Targeting a subnet in a higher-level domain targets the **same** subnet in **each** of the descendant leaf domains. The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.

Compliance White List Host Profiles

In a compliance white list, host profiles specify which operating systems, clients, application protocols, web applications, and protocols are allowed to run on the target hosts. There are three types of host profile you can use in a compliance white list; each type appears differently in the compliance list editor.

Table 225: Compliance White List Host Profile Types

Host Profile Type	Appearance	Description
global	Any Operating System	specifies what is allowed to run on target hosts, regardless of operating system
operating-system specific	is listed in plain text	specifies what is allowed to run on target hosts of a particular operating system

Host Profile Type	Appearance	Description
shared	is listed in italics	specifies operating-system criteria that can be used in multiple white lists

Operating System-Specific Host Profiles

In a compliance white list, *operating-system specific host profiles* indicate not only which operating systems are allowed to run on your network, but also the application protocols, clients, web applications, and protocols that are allowed to run on those operating systems.

For example, you could require that compliant hosts run a particular version of Microsoft Windows. As another example, you could allow SSH to run on Linux hosts on port 22, and further restrict the vendor and version of the SSH client.

Create one host profile for each operating system you want to allow on your network. To disallow an operating system on your network, do not create a host profile for that operating system. For example, to make sure that all the hosts on your network are running Windows, configure the white list to only contain host profiles for that operating system.



Note Unidentified hosts remain in compliance with all white lists until they are identified. You can, however, create a white list host profile for unknown hosts. *Unidentified* hosts are hosts about which the system has not yet gathered enough information to identify their operating systems. *Unknown* hosts are hosts whose operating systems do not match known fingerprints.

Shared Host Profiles

In a compliance white list, *shared host profiles* are tied to specific operating systems, but you can use each shared host profile in more than one white list.

For example, you might have offices worldwide with a separate white list for each location, but you want to use the same profile for all hosts running Apple Mac OS X. You can create a shared profile for that operating system and use it in all your white lists.

The default white list uses a special category of shared host profiles, called *built-in host profiles*. These profiles use built-in application protocols, web applications, protocols, and clients. In the compliance white list editor, the system marks these profiles with the **Built-In Host Profile icon**.

In a multidomain deployment, the system displays shared host profiles created in the current domain, which you can edit. It also displays shared host profiles from ancestor domains, which you cannot edit. To view and edit shared host profiles created in a lower domain, switch to that domain.



Note If you modify a shared host profile (including built-ins), or modify a built-in application protocol, protocol, or client, your change affects every white list that uses it. If you make unintended changes to or delete these built-in elements, you can reset to factory defaults.

White Violation Triggers

The white list compliance of a host can change when the system:

- detects a change in a host's operating system
- detects an identity conflict for a host's operating system or an application protocol on the host
- detects a new TCP server port (for example, a port used by SMTP or web servers) active on a host, or a new UDP server running on a host
- detects a change in a discovered TCP or UDP server running on a host, for example, a version change due to an upgrade
- detects a new client or web application running on a host
- drops a client or web application from its database due to inactivity
- detects that a host is communicating with a new network or transport protocol
- detects a new jailbroken mobile device
- detects that a TCP or UDP port has closed or timed out on a host

In addition, you can trigger a compliance change for a host by using the host input feature or the host profile to:

- add a client, protocol, or server to a host
- delete a client, protocol, or server from a host
- set the operating system definition for a host
- change a host attribute for a host so that the host is no longer a valid target



Note To avoid overwhelming you with events, the system does not generate white list events for non-compliant hosts on its initial evaluation, nor hosts made non-compliant as a result of you modifying an active white list or shared host profile. The violations, however, are still recorded. If you want to generate white list events for all non-compliant targets, purge discovery data. Rediscovering network assets may trigger white list events.

Operating System Compliance

If your white list specifies that only Microsoft Windows hosts are allowed on your network, and the system detects a host running Mac OS X, the system generates a white list event. In addition, the host attribute associated with the white list changes from Compliant to Non-Compliant for that host.

For the host in this example to come back into compliance, one of the following must occur:

- you edit the white list so that the Mac OS X operating system is allowed
- you manually change the operating system definition of the host to Microsoft Windows
- the system detects that the operating system has changed back to Microsoft Windows

Deleting a Non-Compliant Asset from the Network Map

If your white list disallows the use of FTP, and you then delete FTP from the application protocols network map or from an event view, hosts running FTP become compliant. However, if the system detects the application protocol again, the system generates a white list event and the hosts become non-compliant.

Triggering on Complete Information Only

If your white list allows only TCP FTP traffic on port 21, and the system detects indeterminate activity on port 21/TCP, the white list does not trigger. The white list triggers only when the system identifies the traffic as something other than FTP, or you use the host input feature to designate the traffic as non-FTP traffic. The system does not record a violation with only partial information.

Requirements and Prerequisites for Compliance

Model Support

Any

Supported Domains

Any

User Roles

- Admin

Creating a Compliance White List

When you create a compliance white list, the system prompts you to survey your network to create an initial target and to help you characterize compliant hosts.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **White List**.

Step 2 Click **NewWhite List**.

Step 3 Optionally, enter the **IP Address** and **Netmask** for an initial target network. In a multidomain deployment, choose the **Domain** where the target network resides.

Tip To survey the entire monitored network, use the default values of 0.0.0.0/0 and ::/0.

Note After you choose a domain for the target network, you cannot change it. Targeting a subnet in a higher-level domain targets the **same** subnet in **each** of the descendant leaf domains. The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.

Step 4 Add the target network:

- Add—To add the target network without a survey, click **Add**.
- Add and Survey Network—To add and survey the target network, click **Add and Survey Network**.
- Skip—To create a white list without surveying your network, click **Skip**.

Step 5 Optionally, enter a new **Name** and **Description** for the white list.

Step 6 Optionally, **Allow Jailbroken Mobile Devices** on your network. Disabling this option causes jailbroken devices to generate white list violations.

Step 7 Add at least one **Target Network** to the white list, as described in Setting Target Networks for a Compliance White List, on page 1977.

Step 8 Characterize compliant hosts using **Allowed Host Profiles**:

- Global Host Profile—To edit the white list's global host profile, click **Any Operating System** and proceed as described in Building White List Host Profiles, on page 1978.
- Edit Surveyed Profiles—To edit an existing operating system-specific host profile created by a network survey, click its name and proceed as described in Building White List Host Profiles, on page 1978.
- Create New Profiles—To create a new operating system-specific host profile for this white list, click **Add** () next to **Allowed Host Profiles**, and proceed as described in Building White List Host Profiles, on page 1978.
- Add Shared Host Profile—To add an existing shared host profile to the white list, click **Add Shared Host Profile**, select the shared host profile you want to add, then click **OK**. Shared host profiles appear in italics.

Step 9 Click **SaveWhite List**.

What to do next

- Add the white list to an active correlation policy as described in Configuring Correlation Policies, on page 1987. The system immediately starts evaluating the white list and generating violations.

Related Topics

Compliance White List Target Networks, on page 1972
Creating a Compliance White List Based on Selected Hosts, on page 2371
Firepower System IP Address Conventions, on page 15

Setting Target Networks for a Compliance White List

When you add a target network, you can survey it to characterize compliant hosts. This survey populates the white list with one host profile for each operating system detected in the survey. These host profiles allow all the clients, application protocols, web applications, and protocols that the system has detected on the applicable operating systems.

Procedure

Step 1 In the compliance white list editor, click **Add Target Network**.

Step 2 Enter the **IP Address** and **Netmask** for the target network.

Step 3 In a multidomain deployment, choose the **Domain** where the target network resides.

Note After you choose a domain for the target network, you cannot change it. Targeting a subnet in a higher-level domain targets the **same** subnet in **each** of the descendant leaf domains. The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.

Step 4 Add the target network:

- Add — To add the target network without a survey, click **Add**.
- Add and Survey Network — To add and survey the target network, click **Add and Survey Network**.

Step 5 Optionally, click the new target to configure it further:

- Name — Enter a new **Name**.
- Add Networks — To target additional hosts, click **Add** (), then enter the **IP Address** and **Netmask**. To exclude the network from white list compliance, select **Exclude**.
- Add Host Attributes — To target hosts with a specific host attribute, click **Add** (), then specify the **Attribute** and its **Value**.
- Add VLANs — To target a VLAN, click **Add** (), then type a VLAN number (for 802.1q VLANs).
- Delete — To remove a target restriction, click **Delete** ().

Step 6 To immediately implement all changes made since the last time you saved, click **SaveWhite List**.

Related Topics

Compliance White List Target Networks, on page 1972
Firepower System IP Address Conventions, on page 15

Building White List Host Profiles

Host profiles specify the white list's compliance criteria, that is, which operating systems, clients, application protocols, web applications, and protocols are allowed to run on the target hosts.

Every white list has a global host profile which is operating-system agnostic. For example, instead of editing multiple Microsoft Windows and Linux host profiles to allow Mozilla Firefox, you can configure the global host profile to allow Firefox regardless of the operating system where it was detected.

You can also configure operating-system specific host profiles, either unique to one white list or shared across white lists.

Note If you modify a shared host profile (including built-ins), or modify a built-in application protocol, protocol, or client, your change affects every white list that uses it. If you make unintended changes to or delete these built-in elements, you can reset to factory defaults.

Before you begin

- Create or edit a host profile within a white list as described in Editing a Compliance White List, on page 1982, or create or edit a shared host profile as described in Managing Shared Host Profiles, on page 1984.

Procedure

Step 1 In the compliance white list host profile editor, configure a host profile:

- Name — Type a **Name**.
- Operating System — To restrict the host profile to a specific operating system, use the **OS Vendor**, **OS Name**, and **Version** drop-down lists. Because its purpose is to apply to hosts running any operating system, you cannot restrict a global host profile.
- Application Protocol — To allow an application protocol, click **Add** ( ) and proceed as described in Adding an Application Protocol to a Compliance White List, on page 1979.
- Client — To allow a client, click **Add** ( ) and proceed as described in Adding a Client to a Compliance White List, on page 1980.
- Web Application — To allow a web application, click **Add** ( ) and proceed as described in Adding a Web Application to a Compliance White List, on page 1981.
- Protocol — To allow a protocol, click **Add** ( ) and proceed as described in Adding a Protocol to a Compliance White List, on page 1981.
- Delete — To disallow an item you previously allowed, click **Delete** ( ).
- Edit Properties — To edit the properties of an allowed application protocol, client, or protocol, click its name. The changes you make are reflected in every host profile that uses that element.

Tip Select the appropriate **Allow all...** check box to allow all application protocols, clients, or web applications for hosts matching this profile.

Step 2 To immediately implement all changes made since the last time you saved, click **SaveWhite List** (or **Save All Profiles** if you are editing a shared host profile).

Adding an Application Protocol to a Compliance White List

Using white list host profiles, you can allow application protocols either globally or on specific operating systems. Optionally, you can restrict the application protocol by port, vendor, or version. For example, you could allow a particular version of OpenSSH to run on Linux hosts on port 22/TCP.

Procedure

Step 1 While you are creating or modifying a compliance white list host profile, click **Add** ( ) next to **Allowed Application Protocols** (or next to **Globally Allowed Application Protocols** if you are modifying the global host profile).

Step 2 You have two options:

- If the application protocols you want to allow are listed, select them. The web interface lists application protocols that have been allowed or are currently allowed by the white list.
- To allow an application protocol not in the list, select **<New Application Protocol>** and click **OK** to display the application protocol editor. Select the application protocol **Type** and **Protocol** you want to allow. Optionally, restrict the application protocol by **port**, **Vendor**, and **Version**.

  Note You must type the vendor and version exactly as they would appear in a table view of applications. If you do not specify a vendor or version, the white list allows all vendors and versions as long as the type and protocol match.

Step 3 Click **OK**.

Step 4 To immediately implement all changes made since the last time you saved, click **SaveWhite List**.

Adding a Client to a Compliance White List

Using white list host profiles, you can allow clients either globally or on specific operating systems. Optionally, you can require that the client be a specific version. For example, you could allow only Microsoft Internet Explorer 10 to run on Microsoft Windows hosts.

Procedure

Step 1 While you are creating or modifying a compliance white list host profile, click **Add** ( ) next to **Allowed Clients** (or next to **Globally Allowed Clients** if you are modifying the global host profile).

Step 2 You have two options:

- If the clients you want to allow are listed, select them. The web interface lists clients that have been allowed or are currently allowed by the white list.
- To allow a client not in the list, select **<New Client>** and click **OK** to display the client editor. Select the **Client** you want to allow from the drop-down list, and, optionally, restrict the client to an allowed **Version**.

  Note You must type the version exactly as it would appear in a table view of clients. If you do not specify a version, all versions are allowed.

Step 3 Click **OK**.

Step 4 To immediately implement all changes made since the last time you saved, click **SaveWhite List**.

Adding a Web Application to a Compliance White List

Using white list host profiles, you can allow web applications either globally or on specific operating systems.

Procedure

Step 1 While you are creating or modifying a compliance white list host profile, click **Add** ( ) next to **Allowed Web Applications** (or next to **Globally Allowed Web Applications** if you are modifying the global host profile).

Step 2 Select the web applications you want to allow.

Step 3 Click **OK**

Step 4 To immediately implement all changes made since the last time you saved, click **SaveWhite List**.

Adding a Protocol to a Compliance White List

Using white list host profiles, you can allow protocols either globally or on specific operating systems. ARP, IP, TCP, and UDP are always allowed to run on any host; you cannot disallow them.

Procedure

Step 1 While you are creating or modifying a compliance white list host profile, click **Add** ( ) next to **Allowed Protocols** (or next to **Globally Allowed Protocols** if you are modifying the global host profile).

Step 2 You have two options:

- If the protocols you want to allow are listed, select them. The web interface lists protocols that have been allowed or are currently allowed by the white list.
- To allow a protocol not in the list, select **<New Protocol>** and click **OK** to display the protocol editor. From the **Type** drop-down list, select the protocol type (**Network** or **Transport**), then select the **Protocol** from the drop-down list.

  Tip Select **Other (manual entry)** to specify a protocol that is not in the list. For network protocols, type the appropriate number as listed in http://www.iana.org/assignments/ethernet-numbers/. For transport protocols, type the appropriate number as listed in http://www.iana.org/assignments/protocol-numbers/.

Step 3 Click **OK**.

Step 4 To immediately implement all changes made since the last time you saved, click **SaveWhite List**.

Managing Compliance White Lists

You can use the White List page to manage compliance white lists and shared host profiles. The default white list represents recommended settings and uses a special category of shared host profiles, called *built-in host profiles*.

In a multidomain deployment, the system displays compliance white lists created in the current domain, which you can edit. It also displays selected white lists from ancestor domains, which you cannot edit. To view and edit white lists created in a lower domain, switch to that domain.

Note The system does not display configurations from ancestor domains if the configurations expose information about unrelated domains, including names, managed devices, and so on. The default white list is only available in the Global domain.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **White List**.

Step 2 Manage your compliance white lists:

- Create — To create a new white list, click **NewWhite List** and proceed as described in Creating a Compliance White List, on page 1976.
- Delete — To delete a white list that is not in use, click **Delete** (), then confirm you want to delete the white list. Deleting a white list also removes its associated host attribute from all hosts on your network. If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Edit — To modify an existing white list, click **Edit** () and proceed as described in Editing a Compliance White List, on page 1982. If **View** () appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Shared Host Profiles — To manage your white lists' shared host profiles, click **Edit Shared Profiles** and proceed as described in Managing Shared Host Profiles, on page 1984.

Editing a Compliance White List

When you modify and save a compliance white list that is included in an active correlation policy, the system immediately re-evaluates the compliance of the hosts in the white list's target networks. Although this re-evaluation may bring some hosts into or out of compliance, the system does not generate any white list events.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **White List**.

Step 2 Next to the white list you want to modify, click **Edit** ( ).

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Edit your compliance white list:

- Name and Description — To change the name or description, click the white list name in the left panel to display basic white list information, then type the new information.
- Allow Jailbroken Devices — To allow jailbroken mobile devices on your network, click the white list name in the left panel to display basic white list information, then enable **Allow Jailbroken Mobile Devices**. Disabling this option causes jailbroken devices to generate white list violations.
- Add Allowed Host Profile — To create an operating system-specific host profile for this white list, click **Add** ( ) next to Allowed Host Profiles and proceed as described in Building White List Host Profiles, on page 1978.
- Add Shared Host Profile — To add an existing shared host profile to the white list, click **Add Shared Host Profile**, select the shared host profile you want to add, then click **OK**. Shared host profiles appear in italics.
- Add Target Network — To add a new target network without surveying its hosts, click **Add** ( ) next to Target Networks and proceed as described in Setting Target Networks for a Compliance White List, on page 1977.
- Delete Host Profile — To delete a shared or operating-system specific host profile from the white list, click **Delete** ( ) next to the host profile, then confirm your choice. Deleting a shared host profile removes it from the white list, but does not delete the profile or remove it from any other white lists that use it. You cannot delete a white list's global host profile.
- Delete Target Network — To remove a target network from the white list, click **Delete** ( ) next to the network, then confirm your choice.
- Edit Global Host Profile — To edit the white list's global host profile, click **Any Operating System** and proceed as described in Building White List Host Profiles, on page 1978.
- Edit Other Host Profile — To edit a shared or operating-system specific host profile, click the host profile's name and proceed as described in Building White List Host Profiles, on page 1978.
- Edit Target Network — To edit a target network, click the network's name and proceed as directed in Setting Target Networks for a Compliance White List, on page 1977.

Step 4 To immediately implement all changes made since the last time you saved, click **SaveWhite List**.

Managing Shared Host Profiles

In a compliance white list, *shared host profiles* are tied to specific operating systems, but you can use each shared host profile in more than one white list. If you create multiple white lists but want to use the same host profile to evaluate hosts running a particular operating system across the white lists, use a shared host profile.

In a multidomain deployment, the system displays shared host profiles created in the current domain, which you can edit. It also displays shared host profiles from ancestor domains, which you cannot edit. To view and edit shared host profiles created in a lower domain, switch to that domain.

Note If you modify a shared host profile (including built-ins), or modify a built-in application protocol, protocol, or client, your change affects every white list that uses it. If you make unintended changes to or delete these built-in elements, you can reset to factory defaults.

Procedure

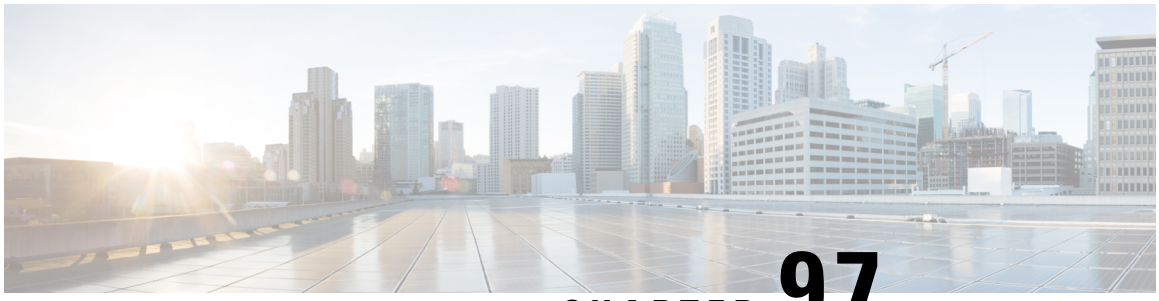
Step 1 Choose **Policies** > **Correlation**, then click **White List**.

Step 2 Click **Edit Shared Profiles**.

Step 3 Manage your shared host profiles:

- Create Shared Host Profile — To create a new shared host profile without surveying hosts, click **Add** ( ) next to Shared Host Profiles and proceed as described in Building White List Host Profiles, on page 1978.
- Create Shared Host Profile by Survey — To create multiple new shared host profiles by surveying a network, click **Add Target Network** and proceed as described in Setting Target Networks for a Compliance White List, on page 1977 .
- Delete — To delete a shared host profile, click **Delete** ( ), then confirm your choice.
- Edit — To modify an existing shared host profile (including a built-in shared host profile), click its name and proceed as described in Building White List Host Profiles, on page 1978.
- Reset Built-In Host Profiles — To reset all built-in host profiles to factory defaults, click **Built-in Host Profiles**, then click **Reset to Factory Defaults** and confirm your choice.

Step 4 To immediately implement all changes made since the last time you saved, click **Save All Profiles**.



CHAPTER 97

Correlation Policies

The following topics describe how to configure correlation policies and rules.

- Introduction to Correlation Policies and Rules, on page 1985
- Requirements and Prerequisites for Compliance, on page 1986
- Configuring Correlation Policies, on page 1987
- Configuring Correlation Rules, on page 1989
- Configuring Correlation Response Groups, on page 2019

Introduction to Correlation Policies and Rules

You can use the *correlation* feature to respond in real time to threats to your network, using *correlation policies*.

A correlation *policy violation* occurs when the activity on your network triggers either a *correlation rule* or *compliance white list* within an active correlation policy.

Correlation Rules

When a correlation rule in an active correlation policy triggers, the system generates a *correlation event*. Correlation rules can trigger when:

- The system generates a specific type of event (connection, intrusion, malware, discovery, user activity, and so on).
- Your network traffic deviates from its normal profile.

You can constrain correlation rules in the following ways:

- Add a *host profile qualification* to constrain the rule using information from the host profile of a host involved in the triggering event.
- Add a *connection tracker* to a correlation rule so that after the rule's initial criteria are met, the system begins tracking certain connections. Then, a correlation event is generated only if the tracked connections meet additional criteria.
- Add a *user qualification* to a correlation rule to track certain users or groups of users. For example, you can constrain a correlation rule so that it triggers only for a particular user's traffic, or traffic from a specific department.

- Add *snooze periods*. When a correlation rule triggers, a snooze period causes that rule not to trigger again for a specified interval. After the snooze period elapses, the rule can trigger again and start a new snooze period.
- Add *inactive periods*. During inactive periods, correlation rules do not trigger.

Although you can configure correlation rules without licensing your deployment, rules that use unlicensed components do not trigger.

Compliance White Lists

A compliance white list specifies which operating systems, applications (web and client), and protocols are allowed on hosts on your network. When a host violates a white list used in an active correlation policy, the system generates a *white list event*.

Correlation Responses

Responses to correlation policy violations include simple alerts and various remediations (such as scanning a host). You can associate each correlation rule or white list with a single response or group of responses.

If network traffic triggers multiple rules or white lists, the system launches all the responses associated with each rule and white list.

Correlation and Multitenancy

In a multidomain deployment, you can create correlation policies at any domain level, using whatever rules, white lists, and responses are available at that level. Higher-level domain administrators can perform correlation within or across domains:

- Constraining a correlation rule by domain matches events reported by that domain's descendants.
- Higher-level domain administrators can create compliance white lists that evaluate hosts across domains. You can target different subnets in different domains in the same white list.



Note The system builds a separate network map for each leaf domain. Using literal configurations (such as IP addresses, VLAN tags, and usernames) to constrain cross-domain correlation rules can have unexpected results.

Related Topics

Introduction to Compliance White Lists, on page 1971
Firepower Management Center Alert Responses, on page 2075
Introduction to Remediations, on page 2033

Requirements and Prerequisites for Compliance

Model Support

Any

Supported Domains

Any

User Roles

- Admin

Configuring Correlation Policies

Use correlation rules, compliance white lists, alert responses, and remediations to build correlation policies.

In a multidomain deployment, you can create correlation policies at any domain level, using whatever constituent configurations are available at that level.

You can assign a priority to each correlation policy, and to each rule and white list used in that policy. Rule and white list priorities override correlation policy priorities. If network traffic violates the correlation policy, the resultant correlation events display the policy priority value, unless the violated rule or white list has its own priority.

Procedure

Step 1 Choose **Policies** > **Correlation**.

Step 2 Click **Create Policy**.

Step 3 Enter a **Policy Name** and **Policy Desription**.

Step 4 From the **Default Priority** drop-down list, choose a priority for the policy. Choose **None** to use rule priorities only.

Step 5 Click **Add Rules**, check the rules and white lists that you want to use in the policy, then click **Add**.

Step 6 From the **Priority** list for each rule or white list, choose a priority:

- A priority value from 1 to 5
- **None**
- **Default** to use the policy's default priority

Step 7 Add responses to rules and white lists as described in Adding Responses to Rules and White Lists, on page 1987.

Step 8 Click **Save**.

What to do next

- Activate the policy by clicking the slider.

Adding Responses to Rules and White Lists

You can associate each correlation rule or white list with a single response or group of responses. If network traffic triggers multiple rules or white lists, the system launches all the responses associated with each rule

and white list. Note that an Nmap remediation does not launch when used as a response to a traffic profile change.

In a multidomain deployment, you can use responses created in the current domain or in ancestor domains.

Procedure

Step 1 In the correlation policy editor, next to a rule or white list where you want to add responses, click **Responses**.

Step 2 Under Unassigned Responses, choose the responses you want to launch when the rule or white list triggers, and click the up arrow (^).

Step 3 Click **Update**.

Related Topics

Firepower Management Center Alert Responses, on page 2075
Introduction to Remediations, on page 2033

Managing Correlation Policies

Changes made to active correlation policies take effect immediately.

When you activate a correlation policy, the system immediately begins processing events and triggering responses. Note that the system does not generate white list events for non-compliant hosts on its initial, post-activation evaluation.

In a multidomain deployment, the system displays correlation policies created in the current domain, which you can edit. It also displays selected correlation policies from ancestor domains, which you cannot edit. To view and edit correlation policies created in a lower domain, switch to that domain.



Note The system does not display configurations from ancestor domains if the configurations expose information about unrelated domains, including names, managed devices, and so on.

Procedure

Step 1 Choose **Policies** > **Correlation**.

Step 2 Manage your correlation policies:

- Activate or Deactivate — Click the slider. If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Create — Click **Create Policy**; see Configuring Correlation Policies, on page 1987.
- Edit — Click **Edit** (); see Configuring Correlation Policies, on page 1987. If **View** () appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

- Delete — Click **Delete** (). If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Configuring Correlation Rules

A simple correlation rule requires only that an event of a certain type occurs. You do not need to provide more specific conditions. For example, correlation rules based on traffic profile changes do not require conditions. You can also create complex correlation rules, with multiple conditions and added constraints.

When you create correlation rule trigger criteria, host profile qualifications, user qualifications, or connection trackers, the syntax varies but the mechanics remain consistent.



Note In a multidomain deployment, constraining a correlation rule by an ancestor domain matches events reported by that domain's descendants.

Before you begin

- Confirm that your deployment is collecting the type of information you want to use to trigger correlation events. For example, the information available for any individual connection or connection summary event depends on several factors, including the detection method, the logging method, and event type. The system can add hosts to the network map from exported NetFlow records, but the available information for these hosts is limited; see Differences between NetFlow and Managed Device Data, on page 1824.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Rule Management**.

Step 2 Click **Create Rule**.

Step 3 Enter a **Rule Name** and **Rule Description**.

Step 4 Optionally, choose a **Rule Group** for the rule.

Step 5 Choose a base event type and, optionally, specify additional trigger criteria for the correlation rule. You can choose the following base event types:

- **an intrusion event occurs**—See Syntax for Intrusion Event Trigger Criteria, on page 1990.
- **a malware event occurs**—See Syntax for Malware Event Trigger Criteria, on page 1993.
- **a discovery event occurs**—See Syntax for Discovery Event Trigger Criteria, on page 1994.
- **user activity is detected** —See Syntax for User Activity Event Trigger Criteria, on page 1997.
- **a host input event occurs**—See Syntax for Host Input Event Trigger Criteria, on page 1998.
- **a connection event occurs**—See Syntax for Connection Event Trigger Criteria, on page 1999.
- **a traffic profile changes**—See Syntax for Traffic Profile Changes, on page 2002.

Step 6 Optionally, further constrain the correlation rule by adding any or all of the following:

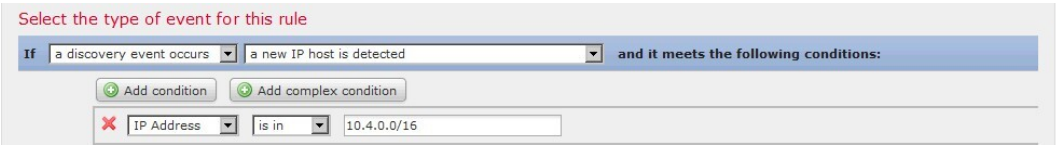
- Host Profile Qualification—Click **Add Host Profile Qualification**; see Syntax for Correlation Host Profile Qualifications, on page 2004.
- Connection Tracker—Click **Add Connection Tracker**; see Connection Trackers, on page 2007.
- User Qualification—Click **Add User Qualification**; see Syntax for User Qualifications, on page 2006.
- Snooze Period—Under Rule Options, use the **Snooze** text field and drop-down list to specify the interval that the system should wait to trigger a correlation rule again, after the rule triggers.
- Inactive Period—Under Rule Options, click **Add Inactive Period**. Using the text field and drop-down lists, specify when and how often you want the system to refrain from evaluating network traffic against the correlation rule.

Tip To remove a snooze period, specify an interval of `0` (seconds, minutes, or hours).

Step 7 Click **Save Rule**.

Example Simple Correlation Rule

The following simple correlation rule triggers if a new host is detected in a specific subnet. Note that when the category represents an IP address, choosing **is in** or **is not in** as the operator allows you to specify whether the IP address *is in* or *is not in* a block of IP addresses, as expressed in special notation such as CIDR.



What to do next

- Use the rule in correlation policies as described in Configuring Correlation Policies, on page 1987.

Related Topics

Managing Correlation Rules, on page 2018
Correlation Rule Building Mechanics, on page 2015
Snooze and Inactive Periods, on page 2015
Differences between NetFlow and Managed Device Data, on page 1824

Syntax for Intrusion Event Trigger Criteria

The following table describes how to build a correlation rule condition when you choose an intrusion event as the base event.

Table 226: Syntax for Intrusion Events

If you specify...	Choose an operator, then...
Access Control Policy	Choose one or more access control policies that use the intrusion policy that generated the intrusion event.

If you specify...	Choose an operator, then...
Access Control Rule Name	Enter all or part of the name of the access control rule that uses the intrusion policy that generated the intrusion event.
Application Protocol	Choose one or more application protocols associated with the intrusion event.
Application Protocol Category	Choose one or more category of application protocol.
Classification	Choose one or more classifications.
Client	Choose one or more clients associated with the intrusion event.
Client Category	Choose one or more category of client.
Destination Country or Source Country	Choose one or more countries associated with the source or destination IP address in the intrusion event.
Destination IP, Source IP, Both Source IP and Destination IP, or Either Source IP or Destination IP	Enter a single IP address or address block.
Destination Port/ICMP Code or Source Port/ICMP Type	Enter the port number or ICMP type for source traffic or the port number or ICMP code for destination traffic.
Device	Choose one or more devices that may have generated the event.
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants. This field is only present if you have ever configured the Firepower Management Center for multitenancy.
Egress Interface or Ingress Interface	Choose one or more interfaces.
Egress Security Zone or Ingress Security Zone	Choose one or more security zones or tunnel zones.
Generator ID	Choose one or more preprocessors.
Impact Flag	Choose the impact level assigned to the intrusion event. Because no operating system information is available for hosts added to the network map from NetFlow data, the system cannot assign Vulnerable (impact level 1: red) impact levels for intrusion events involving those hosts. In such cases, use the host input feature to manually set the operating system identity for the hosts.
Inline Result	Choose whether the system **dropped** or **would have dropped** packets as a result of the intrusion policy violation. The system can drop packets in an inline, switched, or routed deployment. It does not drop packets in a passive deployment, including when an inline set is in tap mode, regardless of intrusion rule state or the drop behavior of the intrusion policy.
Intrusion Policy	Choose one or more intrusion policies that generated the intrusion event.

If you specify...	Choose an operator, then...
IOC Tag	Choose whether an indication of compromise tag was set as a result of the intrusion event.
Priority	Choose the rule priority. For rule-based intrusion events, the priority corresponds to either the value of the `priority` keyword or the value for the `classtype` keyword. For other intrusion events, the priority is determined by the decoder or preprocessor.
Protocol	Enter the name or number of the transport protocol as listed in http://www.iana.org/assignments/protocol-numbers.
Rule Message	Enter all or part of the rule message.
Rule SID	Enter a single Snort ID (SID) or multiple SIDs separated by commas. If you choose **is in** or **is not in** as the operator, you cannot use the multi-selection pop-up window. You must enter a comma-separated list of SIDs.
Rule Type	Specify whether the rule is local. Local rules include custom standard text intrusion rules, standard text rules that you modified, and any new instances of shared object rules created when you saved the rule with modified header information.
SSL Actual Action	Choose the SSL rule action that indicates how the system handled an encrypted connection.
SSL Certificate Fingerprint	Enter the fingerprint of the certificate used to encrypt the traffic, or choose a subject common name associated with the fingerprint.
SSL Certificate Subject Common Name (CN)	Enter all or part of the subject common name of the certificate used to encrypt the session.
SSL Certificate Subject Country (C)	Choose one or more subject country codes of the certificate used to encrypt the session.
SSL Certificate Subject Organization (O)	Enter all or part of the subject organization name of the certificate used to encrypt the session.
SSL Certificate Subject Organizational Unit (OU)	Enter all or part of the subject organizational unit name of the certificate used to encrypt the session.
SSL Flow Status	Choose one or more statuses based on the result of the system's attempt to decrypt the traffic.
Username	Enter the username of the user logged into the source host in the intrusion event.
VLAN ID	Enter the innermost VLAN ID associated with the packet that triggered the intrusion event
Web Application	Choose one or more web applications associated with the intrusion event.
Web Application Category	Choose one or more category of web application.

Related Topics

Intrusion Event Fields, on page 2254
Firepower System IP Address Conventions, on page 15

Syntax for Malware Event Trigger Criteria

To base a correlation rule on a malware event, first specify the type of malware event you want to use. Your choice determines the set of trigger criteria you can use. You can choose:

- **by endpoint-based malware detection** (detection by AMP for Endpoints)
- **by network-based malware detection** (detection by AMP for Networks)
- **by retrospective network-based malware detection** (retroactive detection by AMP for Networks)

The following table describes how to build a correlation rule condition when you choose a malware event as the base event.

Table 227: Syntax for Malware Events

If you specify...	Choose an operator, then...
Application Protocol	Choose one or more application protocols associated with the malware event.
Application Protocol Category	Choose one or more category of application protocol.
Client	Choose one or more clients associated with the malware event.
Client Category	Choose one or more category of client.
Destination Country or Source Country	Choose one or more countries associated with the source or destination IP address in the malware event.
Destination IP, Host IP, or Source IP	Enter a single IP address or address block.
Destination Port/ICMP Code	Enter the port number or ICMP code for destination traffic.
Disposition	Choose either or both **Malware** or **Custom Detection**.
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants. This field is only present if you have ever configured the Firepower Management Center for multitenancy.
Event Type	Choose one or more event types associated with the malware event detected by AMP for Endpoints.
File Name	Enter the name of the file.
File Type	Choose the file type.
File Type Category	Choose one or more file type categories.
IOC Tag	Choose whether an indication of compromise tag **is** or **is not** set as a result of the malware event.
SHA-256	Enter or paste the SHA-256 hash value of the file.
SSL Actual Action	Choose the SSL rule action that indicates how the system handled an encrypted connection.

If you specify...	Choose an operator, then...
SSL Certificate Fingerprint	Enter the fingerprint of the certificate used to encrypt the traffic, or choose a subject common name associated with the fingerprint.
SSL Certificate Subject Common Name (CN)	Enter all or part of the subject common name of the certificate used to encrypt the session.
SSL Certificate Subject Country (C)	Choose one or more subject country codes of the certificate used to encrypt the session.
SSL Certificate Subject Organization (O)	Enter all or part of the subject organization name of the certificate used to encrypt the session.
SSL Certificate Subject Organizational Unit (OU)	Enter all or part of the subject organizational unit name of the certificate used to encrypt the session.
SSL Flow Status	Choose one or more statuses based on the result of the system's attempt to decrypt the traffic.
Source Port/ICMP Type	Enter the port number or ICMP type for source traffic.
Web Application	Choose one or more web applications associated with the malware event.
Web Application Category	Choose one or more category of web application.

Related Topics

File and Malware Event Fields, on page 2300
Firepower System IP Address Conventions, on page 15

Syntax for Discovery Event Trigger Criteria

To base a correlation rule on a discovery event, first specify the type of discovery event you want to use. Your choice determines the set of trigger criteria you can use. The following table lists the discovery event types you can choose.

You cannot trigger a correlation rule on hops changes, or when the system drops a new host due to reaching the host limit. You can, however, choose **there is any type of event** to trigger the rule when any type of discovery event occurs.

Table 228: Correlation Rule Trigger Criteria vs Discovery Event Types

Choose this option...	To use this discovery event type...
a client has changed	Client Update
a client timed out	Client Timeout
a host IP address is reused	DHCP: IP Address Reassigned
a host is deleted because the host limit was reached	Host Deleted: Host Limit Reached
a host is identified as a network device	Host Type Changed to Network Device
a host timed out	Host Timeout

Choose this option...	To use this discovery event type...
a host's IP address has changed	DHCP: IP Address Changed
a NETBIOS name change is detected	NETBIOS Name Change
a new client is detected	New Client
a new IP host is detected	New Host
a new MAC address is detected	Additional MAC Detected for Host
a new MAC host is detected	New Host
a new network protocol is detected	New Network Protocol
a new transport protocol is detected	New Transport Protocol
a TCP port closed	TCP Port Closed
a TCP port timed out	TCP Port Timeout
a UDP port closed	UDP Port Closed
a UDP port timed out	UDP Port Timeout
a VLAN tag was updated	VLAN Tag Information Update
an IOC was set	Indication of Compromise
an open TCP port is detected	New TCP Port
an open UDP port is detected	New UDP Port
the OS information for a host has changed	New OS
the OS or server identity for a host has a conflict	Identity Conflict
the OS or server identity for a host has timed out	Identity Timeout
there is any kind of event	any event type
there is new information about a MAC address	MAC Information Change
there is new information about a TCP server	TCP Server Information Update
there is new information about a UDP server	UDP Server Information Update

The following table describes how to build a correlation rule condition when you choose a discovery event as the base event.

Table 229: Syntax for Discovery Events

If you specify...	Choose an operator, then...
Application Protocol	Choose one or more application protocols.

If you specify...	Choose an operator, then...
Application Protocol Category	Choose one or more category of application protocol.
Application Port	Enter the application protocol port number.
Client	Choose one or more clients.
Client Category	Choose one or more category of client.
Client Version	Enter the version number of the client.
Device	Choose one or more devices that may have generated the discovery event.
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants. This field is only present if you have ever configured the Firepower Management Center for multitenancy.
Hardware	Enter the hardware model for the mobile device. For example, to match all Apple iPhones, enter `iPhone`.
Host Type	Choose one or more host types. You can choose between a host or one of several types of network device.
IP Address or New IP Address	Enter a single IP address or address block.
Jailbroken	Choose **Yes** to indicate that the host in the event is a jailbroken mobile device or **No** to indicate that it is not.
MAC Address	Enter all or part of the MAC address of the host. For example, if you know that devices from a certain hardware manufacturer have MAC addresses that begin with 0A:12:34, you could choose **begins with** as the operator, then enter `0A:12:34` as the value.
MAC Type	Choose whether the MAC address was **ARP/DHCP Detected**. That is, choose whether the system positively identified the MAC address as belonging to the host (**is ARP/DHCP Detected**), or whether the system is seeing many hosts with that MAC address because, for example, there is a router between the managed device and the host (**is not ARP/DHCP Detected**).
MAC Vendor	Enter all or part of the name of the MAC hardware vendor of the NIC used by the network traffic that triggered the discovery event.
Mobile	Choose **Yes** to indicate that the host in the event is a mobile device or **No** to indicate that it is not.
NETBIOS Name	Enter the NetBIOS name of the host.
Network Protocol	Enter the network protocol number as listed in http://www.iana.org/assignments/ethernet-numbers.
OS Name	Choose one or more operating system names.
OS Vendor	Choose one or more operating system vendors.

If you specify...	Choose an operator, then...
OS Version	Choose one or more operating system versions.
Protocol or Transport Protocol	Enter the name or number of the transport protocol as listed in http://www.iana.org/assignments/protocol-numbers.
Source	Choose the source of the host input data (for operating system and server identity changes and timeouts).
Source Type	Choose the type of the source for the host input data (for operating system and server identity changes and timeouts).
VLAN ID	Enter the VLAN ID of the host involved in the event.
Web Application	Choose a web application.

Related Topics

Discovery Event Types, on page 2358
Discovery Event Fields, on page 2364
Firepower System IP Address Conventions, on page 15

Syntax for User Activity Event Trigger Criteria

To base a correlation rule on user activity, first choose the type of user activity you want to use. Your choice determines the set of trigger criteria you can use. You can choose:

- **a new user identity is detected**
- **a user logs into a host**

The following table describes how to build a correlation rule condition when you choose user activity as the base event.

Table 230: Syntax for User Activity

If you specify...	Choose an operator, then...
Device	Choose one or more devices that may have detected the user activity.
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants. This field is only present if you have ever configured the Firepower Management Center for multitenancy.
IP Address	Enter a single IP address or address block.
Username	Enter a username.

Related Topics

User Activity Data Fields
Firepower System IP Address Conventions, on page 15

Syntax for Host Input Event Trigger Criteria

To base a correlation rule on a host input event, first specify the type of host input event you want to use. Your choice determines the set of trigger criteria you can use. The following table lists the host input event types you can choose.

You cannot trigger a correlation rule when you add, delete, or change the definition of a user-defined host attribute, or set a vulnerability impact qualification.

Table 231: Correlation Rule Trigger Criteria vs Host Input Event Types

Choose this option...	To trigger the rule on this event type...
a client is added	Add Client
a client is deleted	Delete Client
a host is added	Add Host
a protocol is added	Add Protocol
a protocol is deleted	Delete Protocol
a scan result is added	Add Scan Result
a server definition is set	Set Server Definition
a server is added	Add Port
a server is deleted	Delete Port
a vulnerability is marked invalid	Vulnerability Set Invalid
a vulnerability is marked valid	Vulnerability Set Valid
an address is deleted	Delete Host/Network
an attribute value is deleted	Host Attribute Delete Value
an attribute value is set	Host Attribute Set Value
an OS definition is set	Set Operating System Definition
host criticality is set	Set Host Criticality

The following table describes how to build a correlation rule condition when you choose a host input event as the base event.

Table 232: Syntax for Host Input Events

If you specify...	Choose an operator, then...
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants. This field is only present if you have ever configured the Firepower Management Center for multitenancy.

If you specify...	Choose an operator, then...
IP Address	Enter a single IP address or address block.
Source	Choose the source for the host input data.
Source Type	Choose the type of the source for the host input data.

Related Topics

Host Input Event Types, on page 2362
Discovery Event Fields, on page 2364
Firepower System IP Address Conventions, on page 15

Syntax for Connection Event Trigger Criteria

To base a correlation rule on a connection event, first specify the type of connection event you want to use. Note that the information available for a connection event can vary depending on how, why, and when the system logged the connection. You can choose:

- **at either the beginning or the end of the connection**
- **at the beginning of the connection**
- **at the end of the connection**

The following table describes how to build a correlation rule condition when you choose a connection event as the base event.

Table 233: Syntax for Connection Events

If you specify...	Choose an operator, then...
Access Control Policy	Choose one or more access control policies that logged the connection.
Access Control Rule Action	Choose one or more actions associated with the access control rule that logged the connection. Choose **Monitor** to trigger correlation events when network traffic matches the conditions of any Monitor rule, regardless of the rule or default action that later handles the connection.
Access Control Rule	Enter all or part of the name of the access control rule that logged the connection. You can enter the name of any Monitor rule whose conditions were matched by a connection, regardless of the rule or default action that later handled the connection.
Application Protocol	Choose one or more application protocols associated with the connection.
Application Protocol Category	Choose one or more categories of application protocol.
Client	Choose one or more clients.
Client Category	Choose one or more categories of client.
Client Version	Enter the version number of the client.
Connection Duration	Enter the duration of the connection event, in seconds.

If you specify...	Choose an operator, then...
Connection Type	Specify whether you want to trigger the correlation rule based on how the connection information was obtained: • Choose **is** and **Netflow** for connection events generated from exported NetFlow data. • Choose **is not** and **Netflow** for connection events detected by a Firepower System managed device.
Destination Country or Source Country	Choose one or more countries associated with the source or destination IP address in the connection event.
Device	Choose one or more devices that either detected the connection, or that processed the connection (for connection data from exported NetFlow records).
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants. This field is only present if you have ever configured the Firepower Management Center for multitenancy.
Egress Interface or Ingress Interface	Choose one or more interfaces.
Egress Security Zone or Ingress Security Zone	Choose one or more security zones or tunnel zones.
Initiator Bytes, Responder Bytes, or Total Bytes	Enter one of: • The number of bytes sent (**Initiator Bytes**). • The number of bytes received (**Responder Bytes**). • The number of bytes both sent and received (**Total Bytes**).
Initiator IP, Responder IP, Both Initiator and Responder IP, or Either Initiator IP or Responder IP	Specify a single IP address or address block.
Initiator Packets, Responder Packets, or Total Packets	Enter one of: • The number of packets sent (**Initiator Packets**). • The number of packets received (**Responder Packets**). • The number of packets both sent and received (**Total Packets**)
Initiator Port/ICMP Type or Responder Port/ICMP Code	Enter the port number or ICMP type for initiator traffic or the port number or ICMP code for responder traffic.
IOC Tag	Specify whether an indication of compromise tag **is** or **is not** set due to the connection event.
NetBIOS Name	Enter the NetBIOS name of the monitored host in the connection.

If you specify...	Choose an operator, then...
NetFlow Device	Choose the IP address of the NetFlow exporter you want to use to trigger the correlation rule. If you did not add any NetFlow exporters to the network discovery policy, the **NetFlow Device** drop-down list is blank.
Prefilter Policy	Choose one or more prefilter policies that handled the connection.
Reason	Choose one or more reasons associated with the connection event.
Security Intelligence Category	Choose one or more Security Intelligence categories associated with the connection event. To use Security Intelligence Category as a condition for end-of-connection events, set that category to **Monitor** instead of **Block** in your access control policy.
SSL Actual Action	Specify the SSL rule action that indicates how the system handled an encrypted connection.
SSL Certificate Fingerprint	Enter the fingerprint of the certificate used to encrypt the traffic, or choose a subject common name associated with the fingerprint.
SSL Certificate Status	Choose one or more statuses associated with the certificate used to encrypt the session.
SSL Certificate Subject Common Name (CN)	Enter all or part of the subject common name of the certificate used to encrypt the session.
SSL Certificate Subject Country (C)	Choose one or more subject country codes of the certificate used to encrypt the session.
SSL Certificate Subject Organization (O)	Enter all or part of the subject organization name of the certificate used to encrypt the session.
SSL Certificate Subject Organizational Unit (OU)	Enter all or part of the subject organizational unit name of the certificate used to encrypt the session.
SSL Cipher Suite	Choose one or more cipher suites used to encrypt the session.
SSL Encrypted Session	Choose **Successfully Decrypted**.
SSL Flow Status	Choose one or more statuses based on the result of the system's attempt to decrypt the traffic.
SSL Policy	Choose one or more SSL policies that logged the encrypted connection.
SSL Rule Name	Enter all or part of the name of the SSL rule that logged the encrypted connection.
SSL Server Name	Enter all or part of the name of the server with which the client established an encrypted connection.
SSL URL Category	Choose one or more URL categories for the URL visited in the encrypted connection.
SSL Version	Choose one or more SSL or TLS versions used to encrypt the session.
TCP Flags	Choose a TCP flag that a connection event must contain in order to trigger the correlation rule. Only connection data generated from NetFlow records contains TCP flags.
Transport Protocol	Enter the transport protocol used by the connection: `TCP` or `UDP`.
Tunnel/Prefilter Rule	Enter all or part of the name of the tunnel or prefilter rule that handled the connection.

If you specify...	Choose an operator, then...
URL	Enter all or part of the URL visited in the connection.
URL Category	Choose one or more URL categories for the URL visited in the connection.
URL Reputation	Choose one or more URL reputation values for the URL visited in the connection.
Username	Enter the username of the user logged in to either host in the connection.
Web Application	Choose one or more web applications associated with the connection.
Web Application Category	Choose one or more categories of web application.

Related Topics

Connection and Security Intelligence Event Fields, on page 2225
Firepower System IP Address Conventions, on page 15

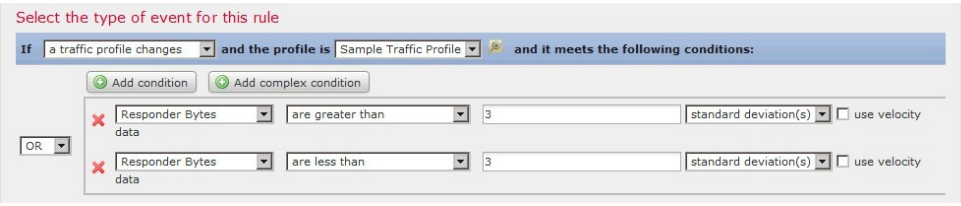
Syntax for Traffic Profile Changes

To base a correlation rule on a traffic profile change, first choose the traffic profile you want to use. The rule triggers when network traffic deviates from the pattern characterized by the profile you choose.

You can trigger the rule based on either raw data or on the statistics calculated from the data. For example, you could write a rule that triggers if the amount of data traversing your network (measured in bytes) suddenly spikes, which could indicate an attack or other security policy violation. You could specify that the rule trigger if either:

- the number of bytes traversing your network spikes above a certain number of bytes
- the number of bytes traversing your network spikes above a certain number of standard deviations above or below the mean amount of traffic

  Note that to create a rule that triggers when the number of bytes traversing your network falls outside a certain number of standard deviations (whether above or below), you must specify upper and lower bounds, as shown in the following graphic.



  To create a rule that triggers when the number of bytes traversing is greater than a certain number of standard deviations *above* the mean, use only the first condition shown in the graphic.

  To create a rule that triggers when the number of bytes traversing is greater than a certain number of standard deviations *below* the mean, use only the second condition.

Check the **use velocity data** check box to trigger the correlation rule based on rates of change between data points. If you wanted to use velocity data in the above example, you could specify that the rule triggers if either:

- the change in the number of bytes traversing your network spikes above or below a certain number of standard deviations above the mean rate of change
- the change in the number of bytes traversing your network spikes above a certain number of bytes

The following table describes how to build a condition in a correlation rule when you choose a traffic profile change as the base event.

Table 234: Syntax for Traffic Profile Changes

If you specify...	Choose an operator, then enter...	Then choose one of...
Number of Connections	the total number of connections detected **or** the number of standard deviations either above or below the mean that the number of connections detected must be in to trigger the rule	connections standard deviation(s)
Total Bytes, Initiator Bytes, or Responder Bytes	one of: • the total bytes transmitted (**Total Bytes**) • the number of bytes transmitted (**Initiator Bytes**) • the number of bytes received (**Responder Bytes**) **or** the number of standard deviations either above or below the mean that one of the above criteria must be in to trigger the rule	bytes standard deviation(s)
Total Packets, Initiator Packets, or Responder Packets	one of: • the total packets transmitted (**Total Packets**) • the number of packets transmitted (**Initiator Packets**) • the number of packets received (**Responder Packets**) **or** the number of standard deviations either above or below the mean that one of the above criteria must be in trigger the rule	packets standard deviation(s)
Unique Initiators	the number of unique hosts that initiated sessions **or** the number of standard deviations either above or below the mean that the number of unique initiators detected must be to trigger the rule	initiators standard deviation(s)

If you specify...	Choose an operator, then enter...	Then choose one of...
Unique Responders	the number of unique hosts that responded to sessions **or** the number of standard deviations either above or below the mean that the number of unique responders detected must be to trigger the rule	responders standard deviation(s)

Syntax for Correlation Host Profile Qualifications

To constrain a correlation rule based on the host profile of a host involved in the event, add a *host profile qualification*. You cannot add a host profile qualification to a correlation rule that triggers on a malware event, traffic profile change, or on the detection of a new IP host.

When you build a host profile qualification, first specify the host you want to use to constrain your correlation rule. The host you can choose depends on the rule's base event type:

- connection event — Choose **Responder Host** or **Initiator Host**.
- intrusion event — Choose **Destination Host** or **Source Host**.
- discovery event, host input event, or user activity — Choose **Host**.

The following table describes how to build a host profile qualification for a correlation rule.

Table 235: Syntax for Host Profile Qualifications

If you specify...	Choose an operator, then...
Application Protocol > Application Protocol	Choose an application protocol.
Application Protocol > Application Port	Enter the application protocol port number.
Application Protocol > Protocol	Choose a protocol.
Application Protocol Category	Choose a category.
Client > Client	Choose a client.
Client > Client Version	Enter the client version.
Client Category	Choose a category.
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants.This field is only present if you have ever configured the Firepower Management Center for multitenancy.
Hardware	Enter the hardware model for the mobile device. For example, to match all Apple iPhones, enter `iPhone`.
Host Criticality	Choose the host criticality.

If you specify...	Choose an operator, then...
Host Type	Choose one or more host types. You can choose between a normal host or one of several types of network device.
IOC Tag	Choose one or more indication of compromise tags.
Jailbroken	Choose **Yes** to indicate that the host in the event is a jailbroken mobile device or **No** to indicate that it is not.
MAC Address > MAC Address	Enter all or part of the MAC address of the host.
MAC Address > MAC Type	Choose whether the MAC type is ARP/DHCP detected: • the system positively identified the MAC address as belonging to the host (**is ARP/DHCP Detected**) • the system is seeing many hosts with that MAC address because, for example, there is a router between the device and the host (**is not ARP/DHCP Detected**) • the MAC type is irrelevant (**is any**)
MAC Vendor	Enter all or part of the MAC vendor of hardware used by the host.
Mobile	Choose **Yes** to indicate that the host in the event is a mobile device or **No** to indicate that it is not.
NetBIOS Name	Enter the NetBIOS name of the host.
Network Protocol	Enter the network protocol number as listed in http://www.iana.org/assignments/ethernet-numbers.
Operating System > OS Vendor	Choose one or more operating system vendor names.
Operating System > OS Name	Choose one or more operating system names.
Operating System > OS Version	Choose one or more operating system versions.
Transport Protocol	Enter the name or number of the transport protocol as listed in http://www.iana.org/assignments/protocol-numbers.
VLAN ID	Enter the VLAN ID number of the host.
Web Application	Choose a web application.
Web Application Category	Choose a category.
any available host attribute, including the default compliance white list host attribute	Enter or choose the appropriate value, depending on the host attribute type.

Using Implied or Generic Clients to Build a Host Profile Qualification

When system reports a detected client using an application protocol name followed by `client` (for example, `HTTPS client`) , that client is an *implied* or *generic* client. In these cases, the system has not detected a particular client, but is inferring the existence of a client based on server response traffic.

To create a host profile qualification using an implied or generic client, constrain using the application protocol running on the responder host, not the client.

Using Event Data to Build a Host Profile Qualification

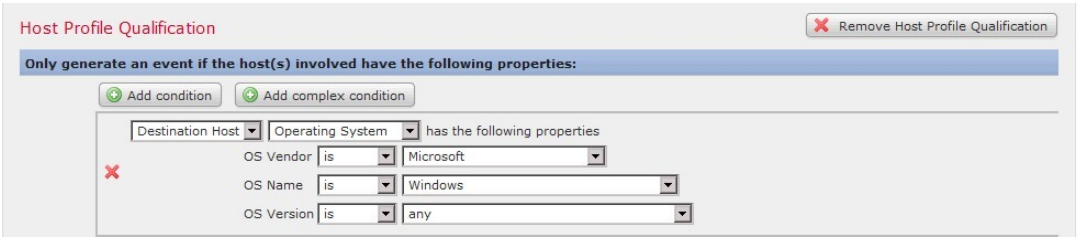
You can often use data from the correlation rule's base event when constructing a host profile qualification.

For example, assume your correlation rule triggers when the system detects the use of a particular browser on one of your monitored hosts. Further assume that when you detect this use, you want to generate an event if the browser version is not the latest.

You could add a host profile qualification to this correlation rule so that the rule triggers only if the **Client** is the **Event Client**, but the **Client Version** is not the latest version.

Example Host Profile Qualification

The following host profile qualification constrains a correlation rule so the rule triggers only if the host involved in the discovery event on which the rule is based is running a version of Microsoft Windows.



Related Topics

Host Data Fields, on page 2366

Syntax for User Qualifications

If you are using a connection, intrusion, discovery, or host input event to trigger your correlation rule, you can constrain the rule based on the identity of a user involved in the event. This constraint is called a *user qualification*. For example, you could constrain a correlation rule so that it triggers only when the identity of the source or destination user is one from the sales department.

You cannot add a user qualification to a correlation rule that triggers on a traffic profile change or on the detection of user activity. Also, the system obtains user details through the Firepower Management Center-server connection established in an identity realm. This information may not be available for all users in the database.

When you build a user qualification, first specify the identity you want to use to constrain your correlation rule. The identity you can choose depends on the rule's base event type:

- connection event — Choose **Identity on Initiator** or **Identity on Responder**.
- intrusion event — Choose **Identity on Destination** or **Identity on Source**.

- discovery event — Choose **Identity on Host**.
- host input event — Choose **Identity on Host**.

The following table describes how to build a user qualification for a correlation rule.

Table 236: Syntax for User Qualifications

If you specify...	Choose an operator, then...
Authentication Protocol	Choose the authentication protocol (or user type) protocol used to detect the user.
Department	Enter a department.
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants. This field is only present if you have ever configured the Firepower Management Center for multitenancy.
Email	Enter an email address.
First Name	Enter a first name.
Last Name	Enter a last name.
Phone	Enter a telephone number.
Username	Enter a username.

Related Topics

User Data Fields

Connection Trackers

A *connection tracker* constrains a correlation rule so that after the rule's initial criteria are met (including host profile and user qualifications), the system begins tracking certain connections. The system generates a correlation event for the rule if the tracked connections meet additional criteria gathered over a time period that you specify.

Tip Connection trackers typically monitor very specific traffic and, when triggered, run only for a finite, specified time. Compare connection trackers with traffic profiles, which typically monitor a broad range of network traffic and run persistently.

There are two ways a connection tracker can generate an event.

Connection Trackers That Fire Immediately When Conditions Are Met

You can configure a connection tracker so that the correlation rule fires as soon as network traffic meets the tracker's conditions. When this happens, the system stops tracking connections for this connection tracker instance, even if the timeout period has not expired. If the same type of policy violation that triggered the correlation rule occurs again, the system creates a new connection tracker.

However, if time expires before network traffic meets the conditions in the connection tracker, the system does not generate a correlation event, and also stops tracking connections for that rule instance.

For example, a connection tracker can serve as a kind of event threshold by generating a correlation event only if a certain type of connection occurs more than a specific number of times within a specific time period. Or, you can generate a correlation event only if the system detects excessive data transfer after an initial connection.

Connection Trackers That Fire at the End of the Timeout Period

You can configure a connection tracker so that it relies on data collected over the entire timeout period, and therefore cannot fire until the end of the timeout period.

For example, if you configure a connection tracker to fire if you detect fewer than a certain number of bytes being transferred during a certain time period, the system waits until that time period passes and then generates an event if network traffic met that condition.

Adding a Connection Tracker

Before you begin

- Create a correlation rule based on a connection, intrusion, discovery, user identity, or host input event. You cannot add a connection tracker to a rule based on a malware event or traffic profile change.

Procedure

Step 1 In the correlation rule editor, click **Add Connection Tracker**.

Step 2 Specify the connections to track; see Syntax for Connection Trackers, on page 2008.

Step 3 Based on the tracked connections, specify when you want to generate a correlation event; see Syntax for Connection Tracker Events, on page 2011.

Step 4 Specify the interval (in seconds, minutes, or hours) during which the tracker's conditions must be met.

Syntax for Connection Trackers

The following table describes how to build a connection tracker condition that specifies the kind of connections you want to track.

Table 237: Syntax for Connection Trackers

If you specify...	Choose an operator, then...
Access Control Policy	Choose one or more access control policies that handled the connections you want to track.
Access Control Rule Action	Choose one or more access control rule actions associated with the access control rule that logged the connections you want to track. Choose **Monitor** to track connections that match the conditions of any Monitor rule, regardless of the rule or default action that later handles the connections.

If you specify...	Choose an operator, then...
Access Control Rule Name	Enter all or part of the name of the access control rule that logged the connections you want to track. To track connections that match a Monitor rule, enter the name of the Monitor rule. The system tracks the connections, regardless of the rule or default action that later handles them.
Application Protocol	Choose one or more application protocols.
Application Protocol Category	Choose one or more application protocol categories.
Client	Choose one or more clients.
Client Category	Choose one or more client categories.
Client Version	Enter the version of the client.
Connection Duration	Enter the connection duration, in seconds.
Connection Type	Specify whether you want to trigger the correlation rule based on how the connection information was obtained: • Choose **is** and **Netflow** for connection events generated from exported NetFlow records. • Choose **is not** and **Netflow** for connection events detected by a Firepower System managed device.
Destination Country or Source Country	Choose one or more countries.
Device	Choose one or more devices whose detected connections you want to track. If you want to track NetFlow connections, choose the devices that process the connection data from exported NetFlow records.
Ingress Interface or Egress Interface	Choose one or more interfaces.
Ingress Security Zone or Egress Security Zone	Choose one or more security zones or tunnel zones.
Initiator IP, Responder IP, or Initiator/Responder IP	Enter a single IP address or address block.
Initiator Bytes, Responder Bytes, or Total Bytes	Enter one of: • the number of bytes transmitted (**Initiator Bytes**) • the number of bytes received (**Responder Bytes**) • the number of bytes both transmitted and received (**Total Bytes**)

If you specify...	Choose an operator, then...
Initiator Packets, Responder Packets, or Total Packets	Enter one of: • the number of packets transmitted (**Initiator Packets**) • the number of packets received (**Responder Packets**) • the number of packets both transmitted and received (**Total Packets**)
Initiator Port/ICMP Type or Responder Port/ICMP Code	Enter the port number or ICMP type for initiator traffic or the port number or ICMP code for responder traffic.
IOC Tag	Choose whether an indication of compromise tag **is** or **is not** set.
NETBIOS Name	Enter the NetBIOS name of the monitored host in the connection.
NetFlow Device	Choose the IP address of the NetFlow exporter you want to track. If you did not add any NetFlow exporters to the network discovery policy, the NetFlow Device drop-down list is blank.
Prefilter Policy	Choose one or more prefilter policies that handled the connections you want to track.
Reason	Choose one or more reasons associated with the connections you want to track.
Security Intelligence Category	Choose one or more Security Intelligence categories associated with the connections you want to track.
TCP Flags	Choose the TCP flag that connections must contain in order to track them. Only connections generated from exported NetFlow records contain TCP flag data.
Transport Protocol	Choose the transport protocol used by the connection.
URL	Enter all or part of the URL visited in the connections you want to track.
URL Category	Choose one or more URL categories for the URL visited in the connections you want to track.
URL Reputation	Choose one or more URL reputation values for the URL visited in the connections you want to track
Username	Enter the username of the user logged into either host in the connections you want to track.
Web Application	Choose one or more web applications.
Web Application Category	Choose one or more web application categories.

Using Event Data to Build a Connection Tracker

You can often use data from the correlation rule's base event when constructing a connection tracker.

For example, assume your correlation rule triggers when the system detects a new client. When you add a connection tracker to this type of correlation rule, the system automatically populates the tracker with constraints that refer to the base event:

- The **Initiator/Responder IP** is set to the **Event IP Address**.
- The **Client** is set to the **Event Client**.

Tip To track connections for a specific IP address or block of IP addresses, click **switch to manual entry** to manually specify the IP. Click **switch to event fields** to go back to using the IP address in the event.

Related Topics

Connection and Security Intelligence Event Fields, on page 2225
Firepower System IP Address Conventions, on page 15

Syntax for Connection Tracker Events

The following table describes how to how to build a connection tracker condition that specifies when you want to generate a correlation event based on the connections you are tracking.

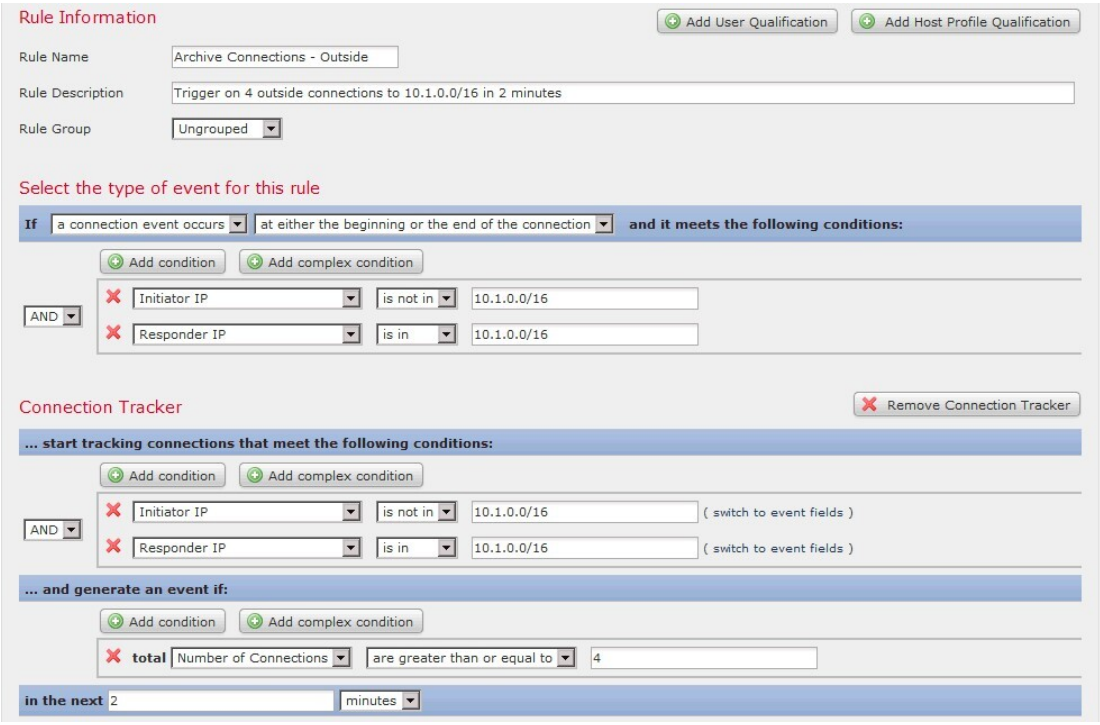
Table 238: Syntax for Connection Tracker Events

If you specify...	Choose an operator, then enter...
Number of Connections	the total number of connections detected
Number of SSL Encrypted Sessions	the total number of SSL- or TLS-encrypted sessions detected
Total Bytes, Initiator Bytes, or Responder Bytes	one of: • the total bytes transmitted (**Total Bytes**) • the number of bytes transmitted (**Initiator Bytes**) • the number of bytes received (**Responder Bytes**)
Total Packets, Initiator Packets, or Responder Packets	one of: • the total packets transmitted (**Total Packets**) • the number of packets transmitted (**Initiator Packets**) • the number of packets received (**Responder Packets**)
Unique Initiators or Unique Responders	one of: • the number of unique hosts that initiated sessions that were detected (**Unique Initiators**) • the number of unique hosts that responded to connections that were detected (**Unique Responders**)

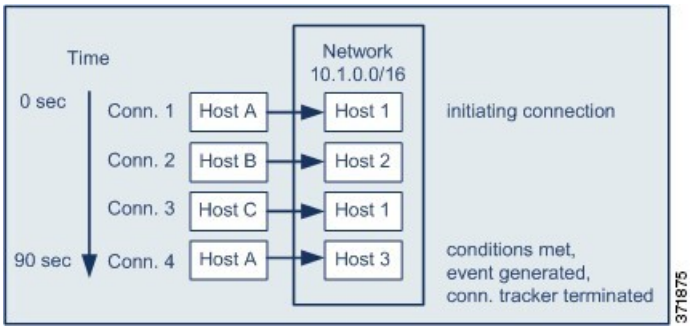
Sample Configuration for Excessive Connections From External Hosts

Consider a scenario where you archive sensitive files on network 10.1.0.0/16, and where hosts outside this network typically do not initiate connections to hosts inside the network. An occasional connection initiated from outside the network might occur, but you have determined that when four or more connections are initiated within two minutes, there is cause for concern.

The rule shown in the following graphic specifies that when a connection occurs from outside the 10.1.0.0/16 network to inside the network, the system begins tracking connections that meet that criterion. The system then generates a correlation event if the system detects four connections (including the original connection) within two minutes that match that signature.



The following diagram shows how network traffic can trigger the above correlation rule.



In this example, the system detected a connection that met the basic conditions of the correlation rule, that is, the system detected a connection from a host outside the 10.1.0.0/16 network to a host inside the network. This created a connection tracker.

The connection tracker is processed in the following stages:

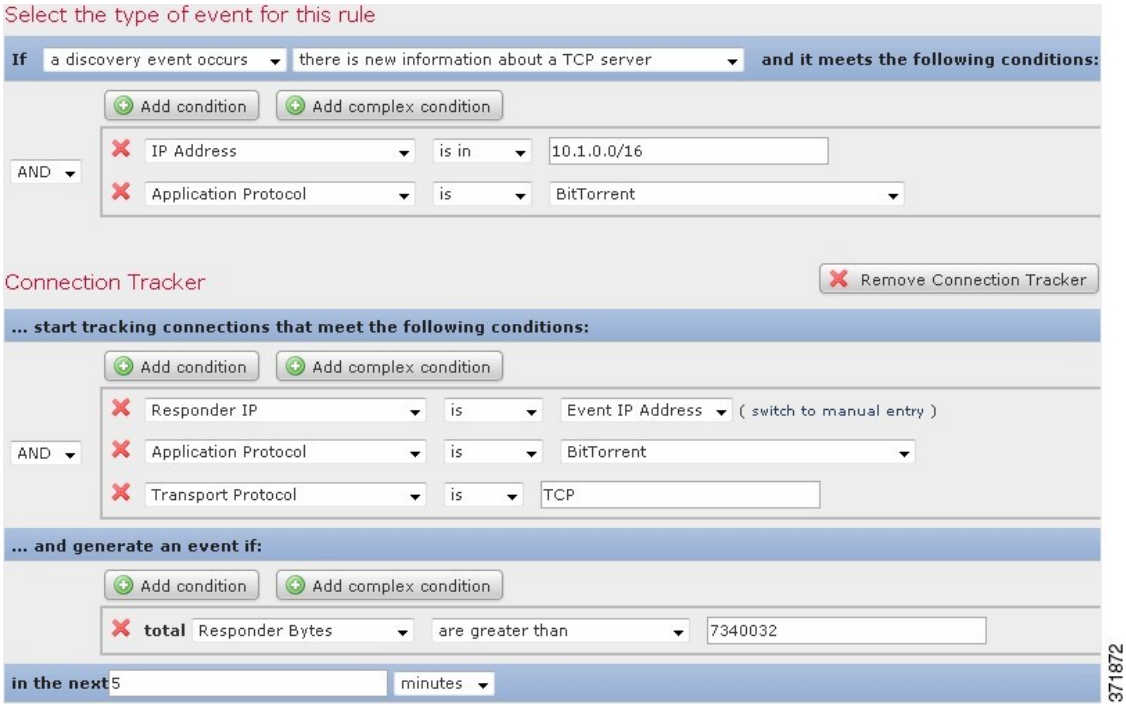
- First, the system starts tracking connections when it detects a connection from Host A outside the network to Host 1 inside the network.
- The system detects two more connections that match the connection tracker signature: Host B to Host 2 and Host C to Host 1.

- The system detects a fourth qualifying connection when Host A connects to Host 3 within the two-minute time limit. The rule conditions are met.
- Finally, the system generates a correlation event and the system stops tracking connections.

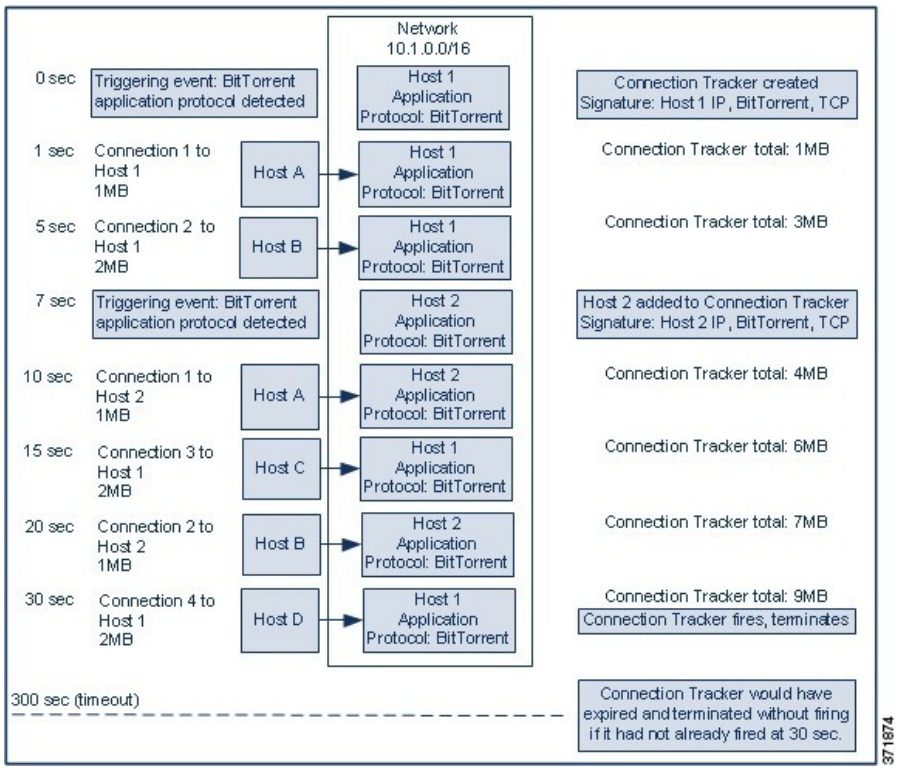
Sample Configuration for Excessive BitTorrent Data Transfers

Consider a scenario where you want to generate a correlation event if the system detects excessive BitTorrent data transfers after an initial connection to any host on your monitored network.

The following graphic shows a correlation rule that triggers when the system detects the BitTorrent application protocol on your monitored network. The rule has a connection tracker that constrains the rule so that the rule triggers only if hosts on your monitored network (in this example, 10.1.0.0/16) collectively transfer more than 7MB of data (7340032 bytes) via BitTorrent in the five minutes following the initial policy violation.



The following diagram shows how network traffic can trigger the above correlation rule.



In this example, the system detected the BitTorrent TCP application protocol on two different hosts: Host 1 and Host 2. These two hosts transmitted data via BitTorrent to four other hosts: Host A, Host B, Host C, and Host D.

This connection tracker is processed in the following stages:

- First, the system starts tracking connections at the 0-second marker when the system detects the BitTorrent application protocol on Host 1. Note that the connection tracker will expire if the system does not detect 7MB of BitTorrent TCP data being transmitted in the next 5 minutes (by the 300-second marker).
- At 5 seconds, Host 1 has transmitted 3MB of data that matches the signature:
  - 1MB from Host 1 to Host A, at the 1-second marker (1MB total BitTorrent traffic counted towards fulfilling the connection tracker)
  - 2MB from Host 1 to Host B, at the 5-second marker (3MB total)
- At 7 seconds, the system detects the BitTorrent application protocol on Host 2 and starts tracking BitTorrent connections for that host as well.
- At 20 seconds, the system has detected additional data matching the signature being transmitted from both Host 1 and Host 2:
  - 1MB from Host 2 to Host A, at the 10-second marker (4MB total)
  - 2MB from Host 1 to Host C, at the 15-second marker (6MB total)
  - 1MB from Host 2 to Host B, at the 20-second marker (7MB total)

- Although Host 1 and Host 2 have now transmitted a combined 7MB of BitTorrent data, the rule does not trigger because the total number of bytes transmitted must be **more** than 7MB (**Responder Bytes are greater than 7340032**). At this point, if the system were to detect no additional BitTorrent transfers for the remaining 280 seconds in the tracker's timeout period, the tracker would expire and the system would not generate a correlation event.
- However, at 30 seconds, the system detects another BitTorrent transfer, and the rule conditions are met:
  - 2MB from Host 1 to Host D at the 30-second marker (9MB total)
- Finally, the system generates a correlation event. The system also stops tracking connections for this connection tracker instance, even though the 5-minute period has not expired. If the system detects a new connection using the BitTorrent TCP application protocol at this point, it will create a new connection tracker. Note that the system generates the correlation event *after* Host 1 transmits the entire 2MB to Host D, because it does not tally connection data until the session terminates.

Snooze and Inactive Periods

You can configure *snooze periods* in correlation rules. When a correlation rule triggers, a snooze period instructs the system to stop firing that rule for a specified interval, even if the rule is violated again during the interval. When the snooze period has elapsed, the rule can trigger again (and start a new snooze period).

For example, you may have a host on your network that should never generate traffic. A simple correlation rule that triggers whenever the system detects a connection involving that host may create multiple correlation events in a short period of time, depending on the network traffic to and from the host. To limit the number of correlation events exposing your policy violation, you can add a snooze period so that the system generates a correlation event only for the first connection (within a time period that you specify) that the system detects involving that host.

You can also set up inactive periods in correlation rules. During inactive periods, the correlation rule will not trigger. You can set up inactive periods to recur daily, weekly, or monthly. For example, you might perform a nightly Nmap scan on your internal network to look for host operating system changes. In that case, you could set a daily inactive period on the affected correlation rules for the time and duration of your scan so that those rules do not trigger erroneously.

Correlation Rule Building Mechanics

You build a correlation rule by specifying the conditions under which it triggers. The syntax you can use within conditions varies depending on the element you are creating, but the mechanics are the same.

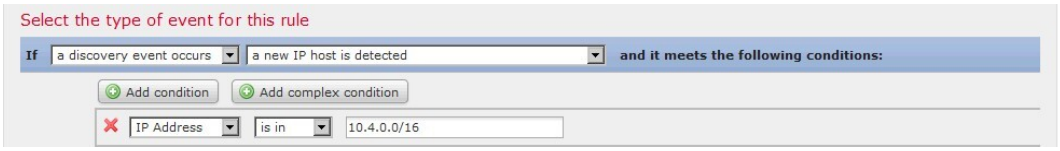
Most conditions have three parts: a *category*, an *operator*, and a *value*:

- The categories you can choose depend on whether you are building correlation rule triggers, a host profile qualification, a connection tracker, or a user qualification. Within correlation rule triggers, the categories further depend on the base event type for the rule. Some conditions may contain several categories, each of which may have their own operators and values.
- A condition's available operators depend on the category.
- The syntax you can use to specify a condition's value depends on the category and operator. Sometimes you type the value in a text field. Other times, you can choose a value (or multiple values) from a drop-down list.

For example, if you want to generate a correlation event every time a new host is detected, you can create a simple rule with no conditions.



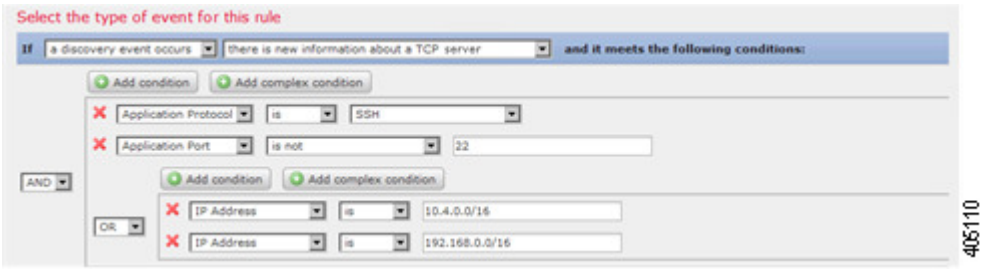
If you want to further constrain the rule and generate an event only if that new host was detected on the 10.4.x.x network, you can add a single condition.



When your construct includes more than one condition, you must link them with an **AND** or an **OR** operator. Conditions on the same level are evaluated together:

- The **AND** operator requires that all conditions on the level it controls must be met.
- The **OR** operator requires that at least one of the conditions on the level it controls must be met.

The following rule, which detects SSH activity on a nonstandard port on the 10.4.x.x network and the 192.168.x.x network, has four conditions, with the bottom two constituting a complex condition.



Logically, the rule is evaluated as follows:

```
(A and B and (C or D))
```

Table 239: Rule Evaluation

Where...	Is the condition that states...
A	Application Protocol is SSH
B	Application Port is not 22
C	IP Address is in 10.0.0.0/8
D	IP Address is in 196.168.0.0/16

Caution Evaluating complex correlation rules that trigger on frequently occurring events can degrade system performance. For example, a multicondition rule that the system must evaluate against every logged connection can cause resource overload.

Adding and Linking Conditions in Correlation Rules

Procedure

Step 1 In the correlation rule editor, add a simple or complex condition:

- Simple — Click **Add condition**.
- Complex — Click **Add complex condition**.

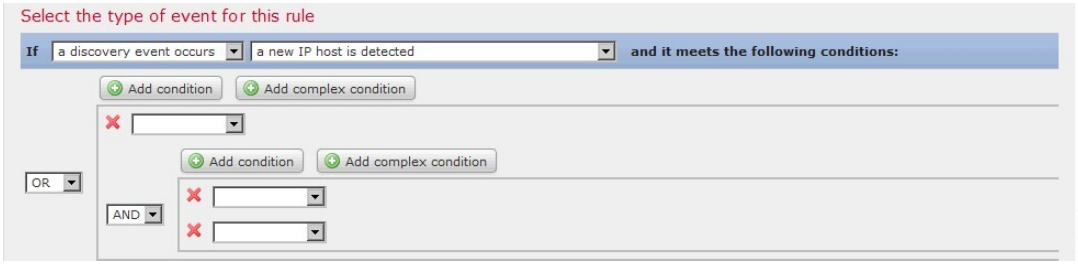
Step 2 Link conditions by choosing the **AND** or **OR** operator from the drop-down list to the left of the conditions.

Example: Simple vs Complex Conditions

The following graphic shows a correlation rule with two simple conditions joined by the **OR** operator.



The following graphic shows a correlation rule with one simple condition and one complex condition, joined by the **OR** operator. The complex condition comprises two simple conditions joined by the **AND** operator.



Using Multiple Values in Correlation Rule Conditions

When you are building a correlation condition, and the condition syntax allows you to pick a value from a drop-down list, you can often use multiple values from the list.

Procedure

Step 1 In the correlation rule editor, build a condition, choosing **is in** or **is not in** as the operator.

Step 2 Click anywhere in the text field or on the **Edit** link.

Step 3 Under **Available**, choose multiple values. You can also click and drag to choose multiple adjacent values.

Step 4 Click the right arrow (>) to move the selected entries to **Selected**.

Step 5 Click **OK**.

Managing Correlation Rules

In a multidomain deployment, the system displays correlation rules and groups created in the current domain, which you can edit. It also displays selected correlation rules and groups from ancestor domains, which you cannot edit. To view and edit correlation rules and groups created in a lower domain, switch to that domain.

Note The system does not display configurations from ancestor domains if the configurations expose information about unrelated domains, including names, managed devices, and so on.

Changes made to rules in active correlation policies take effect immediately.

Before you begin

- If you want to delete a rule, delete it from all correlation policies, as described in Managing Correlation Policies, on page 1988.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Rule Management**.

Step 2 Manage your rules:

- Create — Click **Create Rule**; see Configuring Correlation Rules, on page 1989.
- Create Group — Click **Create Group**, enter a name for the group, and click **Save**. To add a rule to a group, edit the rule.
- Edit — Click **Edit** ( ); see Configuring Correlation Rules, on page 1989. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Delete Rule or Rule Group— Click **Delete** ( ). Deleting a rule group ungroups the rules. If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Configuring Correlation Response Groups

You can create a *correlation response group* of alerts and remediations, then activate and assign the group to a correlation rule within an active correlation policy. The system launches all the grouped responses when network traffic matches the correlation rule.

When used in an active correlation policy, changes to an active group or any of its grouped responses take affect immediately.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Groups**.

Step 2 Click **Create Group**.

Step 3 Enter a **Name**.

Step 4 Check the **Active** check box if you want to activate the group upon creation.

Deactivated groups do not launch responses.

Step 5 Choose the **Available Responses** to group. then click the right arrow (>) to move them to the **Responses in Group**. To move responses the other way, use the left arrow (<).

Step 6 Click **Save**.

What to do next

- If you did not activate the group upon creation and you want to activate it now, click the slider.

Related Topics

Firepower Management Center Alert Responses, on page 2075
Introduction to Remediations, on page 2033

Managing Correlation Response Groups

You can delete a response group if it is not used in a correlation policy. Deleting a response group ungroups its responses. You can also temporarily deactivate a response group without deleting it. This leaves the group on the system but does not launch it when policies are violated.

In a multidomain deployment, the system displays groups created in the current domain, which you can edit. It also displays groups created in ancestor domains, which you cannot edit. To view and edit groups created in a lower domain, switch to that domain.

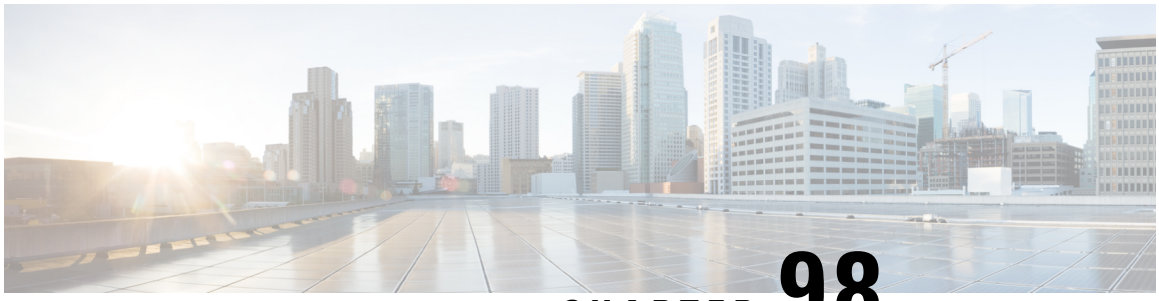
Changes made to active, in-use response groups take effect immediately.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Groups**.

Step 2 Manage response groups:

- Activate or Deactivate — Click the slider. If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Create — Click **Create Group**; see Configuring Correlation Response Groups, on page 2019.
- Edit — Click **Edit** (); see Configuring Correlation Response Groups, on page 2019. If **View** () appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Delete — Click **Delete** (). If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.



CHAPTER 98

Traffic Profiling

The following topics describe how to configure traffic profiles:

- Introduction to Traffic Profiles, on page 2021
- Requirements and Prerequisites for Traffic Profiles, on page 2025
- Managing Traffic Profiles, on page 2025
- Configuring Traffic Profiles, on page 2026

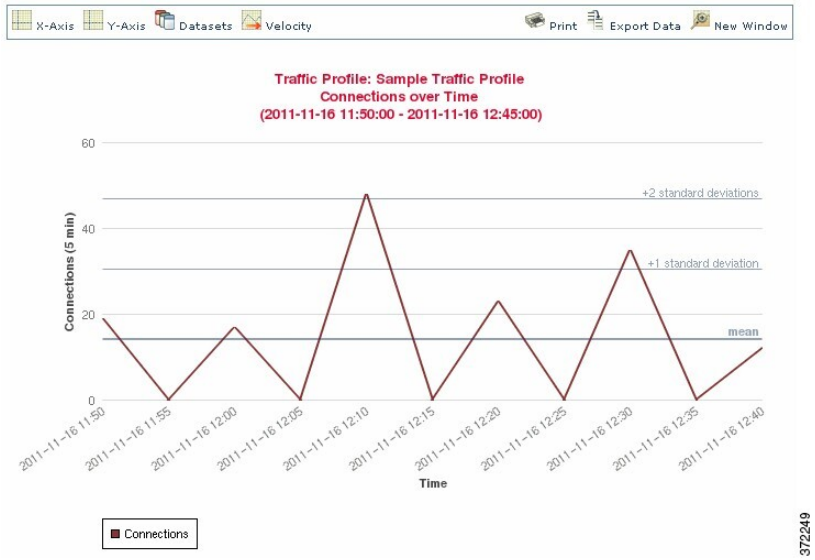
Introduction to Traffic Profiles

A *traffic profile* is a graph of network traffic based on connection data collected over a profiling time window (PTW). This measurement presumably represents normal network traffic. After the learning period, you can detect abnormal network traffic by evaluating new traffic against your profile.

The default PTW is one week, but you can change it to be as short as an hour or as long as several weeks. By default, traffic profiles generate statistics on connection events generated by the system over five-minute intervals. However, you can increase this sampling rate to as long as an hour.

Tip Cisco recommends that the PTW include at least 100 data points. Configure your PTW and sampling rate so that your traffic profiles contain enough data to be statistically meaningful.

The following graphic shows a traffic profile with a PTW of one day and a sampling rate of five minutes.



You can also set up inactive periods in traffic profile. Traffic profiles collect data during inactive periods, but do not use that data when calculating profile statistics. Traffic profile graphs plotted over time show inactive periods as a shaded region.

For example, consider a network infrastructure where all the workstations are backed up at midnight every night. The backup takes about 30 minutes and spikes the network traffic. You could configure recurring inactive period for your traffic profile to coincide with the scheduled backups.

Note The system uses end-of-connection data to create connection graphs and traffic profiles. To use traffic profiles, make sure you log end-of-connection events to the Firepower Management Center database.

Implementing Traffic Profiles

When you activate a traffic profile, the system collects and evaluates connection data for the learning period (PTW) you configured. After the learning period, the system evaluates correlation rules written against the traffic profile.

For example, you could write a rule that triggers if the amount of data traversing your network (measured in packets, KBytes, or number of connections) suddenly spikes to three standard deviations above the mean amount of traffic, which could indicate an attack or other security policy violation. Then, you could include that rule in a correlation policy to alert you of the traffic spike or to perform a remediation in response.

Targeting Traffic Profiles

Profile conditions and *host profile qualifications* constrain traffic profiles.

Using profile conditions, you can profile all network traffic, or you can restrict the traffic profile to monitoring a domain, subnets within or across domains, or individual hosts. In a multidomain deployment:

- Leaf-domain administrators can profile network traffic within their leaf domains.
- Higher-level domain administrators can profile traffic within or across domains.

Profile conditions can also constrain traffic profiles using criteria based on connection data. For example, you could set the profile conditions so that the traffic profile only profiles sessions using a specific port, protocol, or application.

Finally, you can also constrain traffic profiles using information about the tracked hosts. This constraint is called a *host profile qualification*. For example, you could collect connection data only for hosts with high criticality.

Note Constraining a traffic profile to a higher-level domain aggregates and profiles the **same** type of traffic in **each** of the descendant leaf domains. The system builds a separate network map for each leaf domain. In a multidomain deployment, profiling traffic across domains can have unexpected results.

Related Topics

Introduction to Correlation Policies and Rules, on page 1985

Traffic Profile Conditions

You can create simple traffic profile conditions and host profile qualifications, or you can create more elaborate constructs by combining and nesting conditions.

Conditions have three parts: a category, an operator, and a value:

- The categories you can use depend on whether you are building traffic profile conditions or a host profile qualification.
- The operators you can use depend on the category you choose.
- The syntax you can use to specify a condition's value depends on the category and operator. Sometimes you must enter the value in a text field. Other times, you can pick one or more values from a drop-down list.

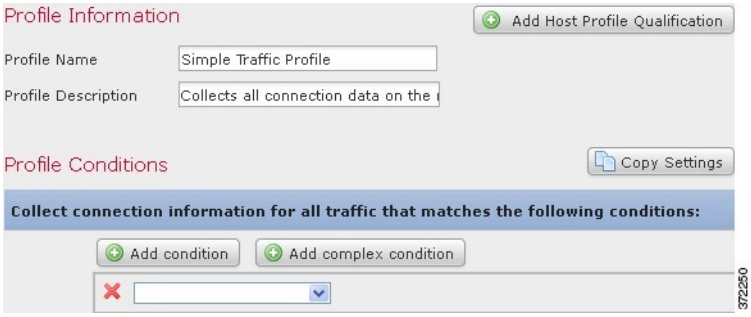
For a host profile qualification, you must also specify whether you are constraining the traffic profile using information data about the initiating or responding hosts.

When your construct includes more than one condition, you must link them with an **AND** or an **OR** operator. Conditions on the same level are evaluated together:

- The **AND** operator requires that all conditions on the level it controls must be met.
- The **OR** operator requires that at least one of the conditions on the level it controls must be met.

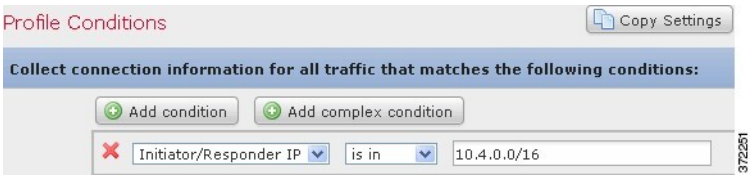
Unconstrained Traffic Profile

If you want to create a traffic profile that collects data for your entire monitored network segment, you can create a very simple profile with no conditions, as shown in the following graphic.



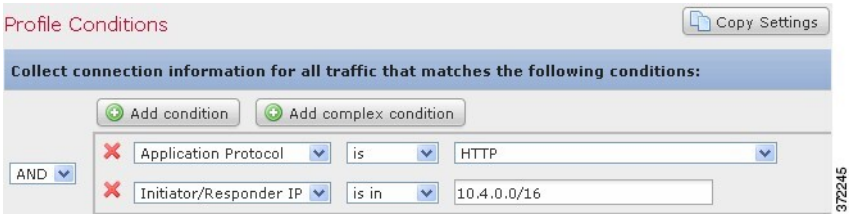
Simple Traffic Profile

If you wanted to constrain the profile and collect data only for a subnet, you can add a single condition, as shown in the following graphic.

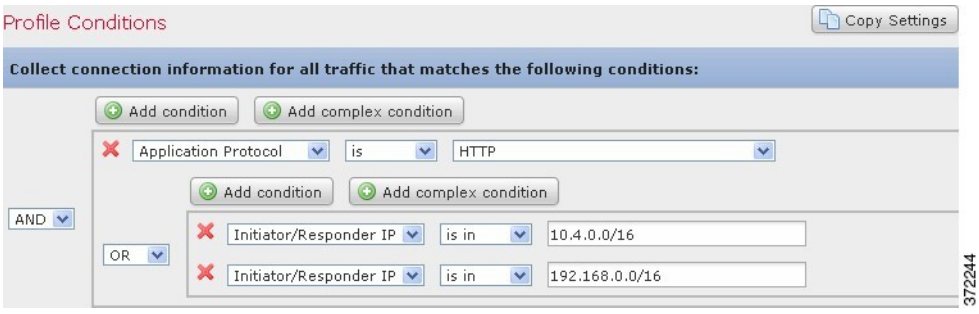


Complex Traffic Profile

The following traffic profile contains two conditions linked by **AND**. This means that the traffic profile collects connection data only if both conditions are true. In this example, it collects HTTP connections for all hosts with IP addresses in a specific subnet.



In contrast, the following traffic profile, which collects connection data for HTTP activity in either of two subnets, has three conditions, with the last constituting a complex condition.



Logically, the above traffic profile is evaluated as follows:

```
(A and (B or C))
```

Where...	Is the condition that states...
A	Application Protocol Name is HTTP
B	IP Address is in 10.4.0.0/16
C	IP Address is in 192.168.0.0/16

Requirements and Prerequisites for Traffic Profiles

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Discovery Admin

Managing Traffic Profiles

Only rules written against active, complete traffic profiles can trigger a correlation policy violation. A slider next to each traffic profile indicates whether the profile is active and collecting data. A progress bar shows the status of the traffic profile's learning period.

In a multidomain deployment, the system displays traffic profiles created in the current domain, which you can edit. It also displays selected traffic profiles from ancestor domains, which you cannot edit. To view and edit traffic profiles created in a lower domain, switch to that domain.

Note The system does not display traffic profiles from ancestor domains if the profiles' conditions expose information about unrelated domains, including names, managed devices, and so on.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Traffic Profiles**.

Step 2 Manage your traffic profiles:

- Activate/Deactivate — To activate or deactivate a traffic profile, click the slider. Deactivating a traffic profile deletes its associated data. If you reactivate the profile, you must wait the length of its PTW before rules written against it will trigger.
- Create — To create a new traffic profile, click **New Profile** and proceed as described in Configuring Traffic Profiles, on page 2026. You can also click **Copy** ( ) to edit a copy of an existing traffic profile.
- Delete — To delete a traffic profile, click **Delete** ( ), then confirm your choice.
- Edit — To modify an existing traffic profile, click **Edit** ( ) and proceed as described in Configuring Traffic Profiles, on page 2026. If a traffic profile is active you can only change its name and description.
- Graph — To view the traffic profile as a graph, click **Graph** ( ). In a multidomain deployment, you cannot view the graph for a traffic profile that belongs to an ancestor domain if the graph exposes information about unrelated domains.

Configuring Traffic Profiles

Constraining a traffic profile to a higher-level domain aggregates and profiles the **same** type of traffic in **each** of the descendant leaf domains. The system builds a separate network map for each leaf domain. In a multidomain deployment, profiling traffic across domains can have unexpected results.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Traffic Profiles**.

Step 2 Click **New Profile**.

Step 3 Enter a **Profile Name**, and optionally, a **Profile Description**.

Step 4 Optionally, constrain the traffic profile:

- Copy Settings — To copy settings from an existing traffic profile, click **Copy Settings**, choose the traffic profile you want to use, and click **Load**.
- Profile Conditions — To constrain the traffic profile using information from tracked connections, proceed as described in Adding Traffic Profile Conditions, on page 2027.
- Host Profile Qualification — To constrain the traffic profile using information from tracked hosts, proceed as described in Adding Host Profile Qualifications to a Traffic Profile, on page 2027.
- Profiling Time Window (PTW) — To change the **Profiling Time Window**, enter a time unit, then choose **hour(s)**, **day(s)**, or **week(s)**.
- Sampling Rate — Choose a **Sampling Rate**, in minutes.
- Inactive Period — Click **Add Inactive Period** and use the drop-down lists to specify when and how often you want the traffic profile remain inactive. Inactive traffic profiles do not trigger correlation rules. Traffic profiles do not include data from inactive periods in profile statistics.

Step 5 Save the traffic profile:

- To save the profile and start collecting data immediately, click **Save & Activate**.
- To save the profile without activating it, click **Save**.

Adding Traffic Profile Conditions

Procedure

Step 1 In the traffic profile editor, under Profile Conditions, click **Add condition** or **Add complex condition** for each condition you want to add. Conditions on the same level are evaluated together.

- To require that all conditions on the level that the operator controls are met, choose **AND**.
- To require that only one of the conditions on the level that the operator controls is met, choose **OR**.

Step 2 Specify a category, operator, and value for each condition as described in Syntax for Traffic Profile Conditions, on page 2028 and Traffic Profile Conditions, on page 2023.

If you choose **is in** or **is not in** as the operator, you can select multiple values in a single condition as described in Using Multiple Values in a Traffic Profile Condition, on page 2031.

When the category represents an IP address, choosing **is in** or **is not in** as the operator allows you to specify whether the IP address *is in* or *is not in* a range of IP addresses.

Example

The following traffic profile collects information on a specific subnet. The category of the condition is **Initiator/Responder IP**, the operator is **is in**, and the value is `10.4.0.0/16`.



Related Topics

Firepower System IP Address Conventions, on page 15

Adding Host Profile Qualifications to a Traffic Profile

Procedure

Step 1 In the traffic profile editor, click **Add Host Profile Qualification**.

Step 2 Under Host Profile Qualification, click **Add condition** or **Add complex condition** for each condition you want to add. Conditions on the same level are evaluated together.

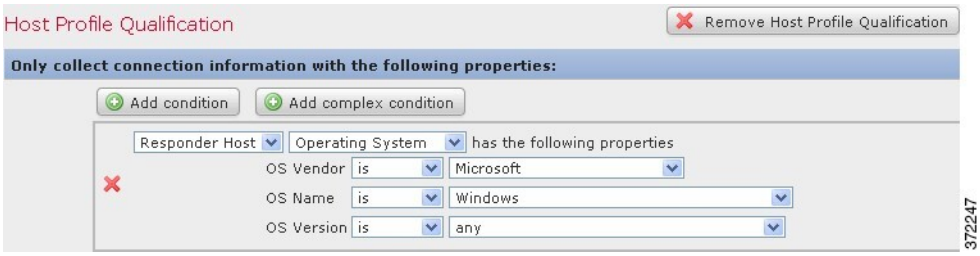
- To require that all conditions on the level that the operator controls are met, choose **AND**.
- To require that only one of the conditions on the level that the operator controls is met, choose **OR**.

Step 3 Specify a host type, category, operator, and value for each condition as described in Syntax for Host Profile Qualifications in a Traffic Profile, on page 2029and Traffic Profile Conditions, on page 2023.

If you choose **is in** or **is not in** as the operator, you can select multiple values in a single condition as described in Using Multiple Values in a Traffic Profile Condition, on page 2031.

Example

The following host profile qualification constrains a traffic profile such that it collects connection data only if the responding host in the detected connection is running a version of Microsoft Windows.



Syntax for Traffic Profile Conditions

The following table describes how to build a traffic profile condition. Keep in mind the connection data available to build a traffic profile depends on several factors, including traffic characteristics and detection method.

Table 240: Syntax for Traffic Profile Conditions

If you choose...	Choose an operator, then...
Application Protocol	Choose one or more application protocols.
Application Protocol Category	Choose one or more application protocol categories.
Client	Choose one or more clients.
Client Category	Choose one or more client categories.
Connection Type	Choose whether the profile uses connection data from traffic monitored by Firepower System managed devices or from exported NetFlow records. If you do not specify a connection type, the traffic profile includes both.
Destination Country or Source Country	Choose one or more countries.
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants.
Initiator IP, Responder IP, or Initiator/Responder IP	Enter an IP address or range of IP addresses. The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.

If you choose...	Choose an operator, then...
NetFlow Device	Choose the NetFlow exporter whose data you want to use to create the traffic profile.
Responder Port/ICMP Code	Enter the port number or ICMP code.
Security Intelligence Category	Choose one or more a Security Intelligence categories. To use a Security Intelligence category for a traffic profile condition, that category must be set to **Monitor** instead of **Block** in your access control policy.
SSL Encrypted Session	Choose **Successfully Decrypted**.
Transport Protocol	Enter **TCP** or **UDP** as the transport protocol.
Web Application	Choose one or more web applications.
Web Application Category	Choose one or more web application categories.

Related Topics

Requirements for Populating Connection Event Fields, on page 2240
Firepower System IP Address Conventions, on page 15

Syntax for Host Profile Qualifications in a Traffic Profile

When you build a host profile qualification condition, you must first choose the host you want to use to constrain your traffic profile. You can choose either **Responder Host** or **Initiator Host**. After you choose the host role, continue building your host profile qualification condition.

Although you can add hosts to the network map using NetFlow records, the available information about these hosts is limited. For example, there is no operating system data available for these hosts, unless you provide it using the host input feature. In addition, if your traffic profile uses connection data from exported NetFlow records, keep in mind that NetFlow records do not contain information about which host in the connection is the initiator and which is the responder. When the system processes NetFlow records, it uses an algorithm to determine this information based on the ports each host is using, and whether those ports are well-known.

To match against *implied* or generic clients, create a host profile qualification based on the application protocol used by the server responding to the client. When the client list on a host that acts as the initiator or source of a connection includes an application protocol name followed by **client**, that client may actually be an implied client. In other words, the system reports that client based on server response traffic that uses the application protocol for that client, not on detected client traffic.

For example, if the system reports **HTTPS client** as a client on a host, create a host profile qualification for **Responder Host** where **Application Protocol** is set to **HTTPS**, because HTTPS client is reported as a generic client based on the HTTPS server response traffic sent by the responder or destination host.

Table 241: Syntax for Host Profile Qualifications

If you choose...	Choose an operator, then...
Application Protocol > Application Protocol	Choose one or more application protocols.

If you choose...	Choose an operator, then...
Application Protocol > Application Port	Enter the application protocol port number.
Application Protocol > Protocol	Choose the protocol.
Application Protocol Category	Choose one or more application protocol categories.
Client > Client	Choose one or more clients.
Client > Client Version	Enter the client version.
Client Category	Choose one or more client categories.
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants.
Hardware	Enter a mobile device hardware model. For example, to match all Apple iPhones, enter `iPhone`.
Host Criticality	Choose a host criticality.
Host Type	Choose one or more host types. You can choose between a normal host or one of several types of network device.
IOC Tag	Choose one or more IOC tags.
Jailbroken	Choose **Yes** to indicate that the host in the event is a jailbroken mobile device or **No** to indicate that it is not.
MAC Address > MAC Address	Enter all or part of the MAC address of the host.
MAC Address > MAC Type	Choose whether the MAC type is **ARP/DHCP Detected**, that is, whether: • The system positively identified the MAC address as belonging to the host (**is ARP/DHCP Detected**) • The system is seeing many hosts with that MAC address because, for example, there is a router between the device and the host (**is not ARP/DHCP Detected**) • The MAC type is irrelevant (**is any**)
MAC Vendor	Enter all or part of the MAC vendor of hardware used by the host.
Mobile	Choose **Yes** to indicate that the host in the event is a mobile device or **No** to indicate that it is not.
NETBIOS Name	Enter the NetBIOS name of the host.
Network Protocol	Enter the network protocol number as listed in http://www.iana.org/assignments/ethernet-numbers.
Operating System > OS Vendor	Choose one or more operating system vendor names.
Operating System > OS Name	Choose one or more operating system names.
Operating System > OS Version	Choose one or more operating system versions.

If you choose...	Choose an operator, then...
Transport Protocol	Enter the name or number of the transport protocol as listed in http://www.iana.org/assignments/protocol-numbers.
VLAN ID	Enter the VLAN ID number of the host. The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal VLAN tags to constrain this configuration can have unexpected results.
Web Application	Choose one or more web applications.
Web Application Category	Choose one or more web application categories.
any available host attribute, including the default compliance white list host attribute	Specify the appropriate value, which depends on the type of host attribute you choose: • If the host attribute type is Integer, enter an integer value in the range defined for the attribute. • If the host attribute type is Text, enter a text value. • If the host attribute type is List, choose a valid list string. • If the host attribute type is URL, enter a URL value.

Using Multiple Values in a Traffic Profile Condition

When you are building a condition, and the condition syntax allows you to pick a value from a drop-down list, you can often use multiple values from the list.

For example, if you want to add a host profile qualification to a traffic profile that requires that a host be running some flavor of UNIX, instead of constructing multiple conditions linked with the OR operator, use the following procedure.

Procedure

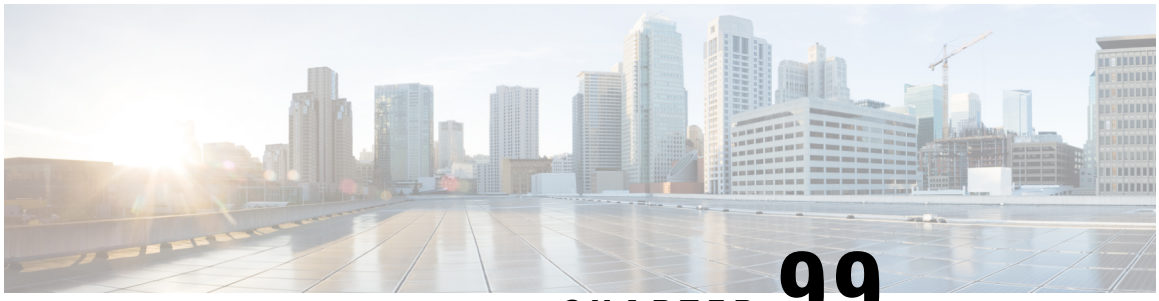
Step 1 While building a traffic profile or host profile qualification condition, choose **is in** or **is not in** as the operator. The drop-down list changes to a text field.

Step 2 Click anywhere in the text field or on the **Edit** link.

Step 3 Under **Available**, choose multiple values.

Step 4 Click the right arrow to move the selected entries to **Selected**.

Step 5 Click **OK**.



CHAPTER 99

Remediations

The following topics contain information on configuring remediations:

- Requirements and Prerequisites for Remediations, on page 2033
- Introduction to Remediations, on page 2033
- Managing Remediation Modules, on page 2043
- Managing Remediation Instances, on page 2044
- Managing Instances for a Single Remediation Module, on page 2044

Requirements and Prerequisites for Remediations

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Discovery Admin

Introduction to Remediations

A *remediation* is a program that the Firepower System launches in response to a correlation policy violation.

When a remediation runs, the system generates a *remediation status event*. Remediation status events include details such as the remediation name, the correlation policy and rule that triggered it, and the exit status message.

The system supports several remediation modules:

- Cisco ISE Endpoint Protection Services (EPS) — quarantines, unquarantines, or shuts down traffic sent to a host or network involved in a correlation policy violation

- Cisco IOS Null Route — blocks traffic sent to a host or network involved in a correlation policy violation (requires Cisco IOS Version 12.0 or higher)
- Nmap Scanning — scans hosts to determine running operating systems and servers
- Set Attribute Value — sets a host attribute on a host involved in a correlation policy violation



Tip You can install custom modules that perform other tasks; see the *Firepower System Remediation API Guide*.

Implementing Remediations

To implement a remediation, first create at least one *instance* for the module you choose. You can create multiple instances per module, where each instance is configured differently. For example, to communicate with multiple routers using the Cisco IOS Null Route remediation module, configure multiples instances of that module.

You can then add multiple *remediations* to each instance that describe the actions you want to perform when a policy is violated.

Finally, associate remediations with rules in correlation policies, so that the system launches the remediations in response to correlation policy violations.

Remediations and Multitenancy

In a multidomain deployment, you can install custom remediation modules at any domain level. The system-provided modules belong to the Global domain.

Though you cannot add a remediation to an instance created in an ancestor domain, you can create a similarly configured instance in the current domain and add remediations to that instance. You can also use remediations created in ancestor domains as correlation responses.

Related Topics

Firepower Management Center Alert Responses, on page 2075
Nmap Scanning, on page 1855
Adding Responses to Rules and White Lists, on page 1987

Cisco ISE EPS Remediations

If you have Endpoint Protection Service (EPS) enabled and configured in your ISE deployment, you can configure your Firepower Management Center to launch remediations using ISE. When fully configured, ISE EPS remediations run the following **Mitigation Actions** on the source or destination host involved in a correlation policy violation:

- **`quarantine`**—Limits or denies an endpoint's access the network
- **`unquarantine`**—Reverses an endpoint's quarantine status and allows full access to the network
- **`shutdown`**—Deactivates an endpoint's network attached system (NAS) port to disconnect it from the network

You can also exempt specific IP addresses from ISE EPS remediation.

Note Your ISE version and configuration impact how you can use ISE in the Firepower System. For more information, see The ISE Identity Source, on page 1895 for more information.

For more information about ISE EPS actions, see the *Cisco Identity Services Engine User Guide*.

Configuring ISE EPS Remediations

You can respond to correlation policy violations by running ISE EPS remediations on the source or destination host.

Before you begin

- Configure EPS operations on your ISE server.
- Configure a connection to ISE as described in Configure ISE for User Control, on page 1897.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Add a pxGrid mitigation instance as described in Adding an ISE EPS Instance, on page 2035.

Step 3 Add one or more ISE EPS remediations as described in Adding ISE EPS Remediations, on page 2036.

What to do next

- Assign remediations as responses to correlation policy violations as described in Adding Responses to Rules and White Lists, on page 1987.

Adding an ISE EPS Instance

Create ISE EPS instances to group individual remediations by logging type.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 From the **Add a New Instance** list, choose **pxGrid Mitigation(v1.0)** as the module type and click **Add**.

Step 3 Enter an **Instance Name** and **Description**.

Step 4 Set **Enable Logging** option to enable or disable system logging.

Step 5 Click **Create**.

What to do next

- Create an ISE EPS remediation as described in Adding Set Attribute Value Remediations, on page 2042.

Related Topics

Firepower System IP Address Conventions, on page 15

Adding ISE EPS Remediations

Create one or more ISE EPS remediations within an instance to run **Mitigation Actions** on the source or destination host involved in a correlation policy violation.

In a multidomain deployment, you cannot add a remediation to an instance created in an ancestor domain.

Before you begin

- Create an ISE EPS instance as described in Adding an ISE EPS Instance, on page 2035.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Next to the instance where you want to add the remediation, click **View** ( ).

Step 3 In the **Configured Remediations** section, choose the **Mitigate Destination** or **Mitigate Source** and click **Add**.

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Enter a **Remediation Name** and **Description**.

Step 5 Choose a **Mitigation Action**: **`quarantine`**, **`unquarantine`**, or **`shutdown`**.

Step 6 (Optional) To exempt IP addresses or ranges from remediation, enter them into the **Whitelist** box.

Step 7 Click **Create**, then click **Done**.

What to do next

- Assign remediations as responses to correlation policy violations; see Adding Responses to Rules and White Lists, on page 1987.

Cisco IOS Null Route Remediations

The Cisco IOS Null Route remediation module allows you to block an IP address or range of addresses using Cisco's "null route" command. This drops all traffic sent to a host or network by routing it to the router's NULL interface. This does not block traffic sent from the violating host or network.

Note Do not use a destination-based remediation as a response to a correlation rule that is based on a discovery or host input event. These events are associated with source hosts.

Caution When a Cisco IOS remediation is activated, there is no timeout period. To unblock the IP address or network, you must manually clear the routing change from the router.

Configuring Remediations for Cisco IOS Routers

Note Do not use a destination-based remediation as a response to a correlation rule that is based on a discovery or host input event. These events are associated with source hosts.

Caution When a Cisco IOS remediation is activated, there is no timeout period. To unblock the IP address or network, you must manually clear the routing change from the router.

Before you begin

- Confirm that your Cisco router is running Cisco IOS 12.0 or higher.
- Confirm that you have level 15 administrative access to the router.

Procedure

Step 1 Enable Telnet on the Cisco router as described in the documentation provided with your Cisco router or IOS software.

Step 2 On the Firepower Management Center, add a Cisco IOS Null Route instance for each Cisco IOS router you plan to use; see Adding a Cisco IOS Instance, on page 2037.

Step 3 Create remediations for each instance, based on the type of response you want to elicit on the router when correlation policies are violated:

- Adding Cisco IOS Block Destination Remediations, on page 2038
- Adding Cisco IOS Block Destination Network Remediations, on page 2039
- Adding Cisco IOS Block Source Remediations, on page 2040
- Adding Cisco IOS Block Source Network Remediations, on page 2040

What to do next

- Assign remediations as responses to correlation policy violations; see Adding Responses to Rules and White Lists, on page 1987.

Adding a Cisco IOS Instance

If you have multiple routers where you want to send remediations, create a separate instance for each router.

Before you begin

- Configure Telnet access on the Cisco IOS router as described in the documentation provided with the router or IOS software.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 From the **Add a New Instance** list, choose **Cisco IOS Null Route** and click **Add**.

Step 3 Enter an **Instance Name** and **Description**.

Step 4 In the **Router IP** field, enter the IP address of the Cisco IOS router you want to use for the remediation.

Step 5 In the **Username** field, enter the Telnet user name for the router. This user must have level 15 administrative access on the router.

Step 6 In the **Connection Password** fields, enter the Telnet user's user password.

Step 7 In the **Enable Password** fields, enter the Telnet user's enable password. This is the password used to enter privileged mode on the router.

Step 8 In the **White List** field, enter IP addresses or ranges that you want to exempt from the remediation, one per line.

Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.

Step 9 Click **Create**.

What to do next

- Add specific remediations to be used by correlation policies as described in Adding Cisco IOS Block Destination Remediations, on page 2038, Adding Cisco IOS Block Destination Network Remediations, on page 2039, Adding Cisco IOS Block Source Remediations, on page 2040, and Adding Cisco IOS Block Source Network Remediations, on page 2040.

Related Topics

Firepower System IP Address Conventions, on page 15

Adding Cisco IOS Block Destination Remediations

The Cisco IOS Block Destination remediation blocks traffic sent from the router to the destination host involved in a correlation policy violation. Do not use this remediation as a response to a correlation rule that is based on a discovery or host input event. These events are associated with source hosts.

In a multidomain deployment, you cannot add a remediation to an instance created in an ancestor domain.

Before you begin

- Add a Cisco IOS instance as described in Adding a Cisco IOS Instance, on page 2037.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Next to the instance where you want to add the remediation, click **View** ( ).

Step 3 In the **Configured Remediations** section, choose **Block Destination** and click **Add**.

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Enter a **Remediation Name** and **Description**.

Step 5 Click **Create**, then click **Done**.

What to do next

- Assign remediations as responses to correlation policy violations; see Adding Responses to Rules and White Lists, on page 1987.

Adding Cisco IOS Block Destination Network Remediations

The Cisco IOS Block Destination Network remediation blocks traffic sent from the router to the network of the destination host involved in a correlation policy violation. Do not use this remediation as a response to a correlation rule that is based on a discovery or host input event. These events are associated with source hosts.

In a multidomain deployment, you cannot add a remediation to an instance created in an ancestor domain.

Before you begin

- Add a Cisco IOS instance as described in Adding a Cisco IOS Instance, on page 2037.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Next to the instance where you want to add the remediation, click **View** ( ).

Step 3 In the **Configured Remediations** section, choose **Block Destination Network** and click **Add**.

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Enter a **Remediation Name** and **Description**.

Step 5 In the **Netmask** field, enter the subnet mask or use CIDR notation to describe the network that you want to block traffic to.

For example, to block traffic to an entire Class C network when a single host triggered a rule (this is not recommended), use `255.255.255.0` or `24` as the netmask.

As another example, to block traffic to 30 addresses that include the triggering IP address, specify `255.255.255.224` or `27` as the netmask. In this case, if the IP address `10.1.1.15` triggers the remediation, all IP addresses between `10.1.1.1` and `10.1.1.30` are blocked. To block only the triggering IP address, leave the field blank, enter `32`, or enter `255.255.255.255`.

Step 6 Click **Create**, then click **Done**.

What to do next

- Assign remediations as responses to correlation policy violations; see Adding Responses to Rules and White Lists, on page 1987.

Related Topics

Firepower System IP Address Conventions, on page 15

Adding Cisco IOS Block Source Remediations

The Cisco IOS Block Source remediation blocks traffic sent from the router to the source host involved in a correlation policy violation.

In a multidomain deployment, you cannot add a remediation to an instance created in an ancestor domain.

Before you begin

- Add a Cisco IOS instance as described in Adding a Cisco IOS Instance, on page 2037.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Next to the instance where you want to add the remediation, click **View** ( ).

Step 3 In the **Configured Remediations** section, choose **Block Source** and click **Add**.

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Enter a **Remediation Name** and **Description**.

Step 5 Click **Create**, then click **Done**.

What to do next

- Assign remediations as responses to correlation policy violations; see Adding Responses to Rules and White Lists, on page 1987.

Adding Cisco IOS Block Source Network Remediations

The Cisco IOS Block Source Network remediation blocks traffic sent from the router to the network of the source host involved in a correlation policy violation.

In a multidomain deployment, you cannot add a remediation to an instance created in an ancestor domain.

Before you begin

- Add a Cisco IOS instance as described in Adding a Cisco IOS Instance, on page 2037.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Next to the instance where you want to add the remediation, click **View** ( ).

Step 3 In the **Configured Remediations** section, choose **Block Source Network** and click **Add**.

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Enter a **Remediation Name** and **Description**.

Step 5 In the **Netmask** field, enter the subnet mask or CIDR notation that describes the network that you want to block traffic to.

For example, to block traffic to an entire Class C network when a single host triggered a rule (this is not recommended), use `255.255.255.0` or `24` as the netmask.

As another example, to block traffic to 30 addresses that include the triggering IP address, specify `255.255.255.224` or `27` as the netmask. In this case, if the IP address `10.1.1.15` triggers the remediation, all IP addresses between `10.1.1.1` and `10.1.1.30` are blocked. To block only the triggering IP address, leave the field blank, enter `32`, or enter `255.255.255.255`.

Step 6 Click **Create**, then click **Done**.

What to do next

- Assign remediations as responses to correlation policy violations; see Adding Responses to Rules and White Lists, on page 1987.

Related Topics

Firepower System IP Address Conventions, on page 15

Nmap Scan Remediations

The Firepower System integrates with Nmap™, an open source active scanner for network exploration and security auditing. You can respond to a correlation policy violation using an Nmap remediation, which triggers an Nmap scan remediation.

For more information about Nmap scanning, see Nmap Scanning, on page 1855.

Set Attribute Value Remediations

You can respond to a correlation policy violation by setting a host attribute value on the host where the triggering event occurred. For text host attributes, you can use the description from the event as the attribute value.

Configuring Set Attribute Remediations

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Create a set attribute instance as described in Adding a Set Attribute Value Instance, on page 2042.

Step 3 Add a set attribute remediation as described in Adding Set Attribute Value Remediations, on page 2042.

What to do next

- Assign remediations as responses to correlation policy violations; see Adding Responses to Rules and White Lists, on page 1987.

Related Topics

Predefined Host Attributes, on page 2341
User-Defined Host Attributes, on page 2342

Adding a Set Attribute Value Instance

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 From the **Add a New Instance** list, choose **Set Attribute Value** and click **Add**.

Step 3 Enter an **Instance Name** and **Description**.

Step 4 Click **Create**.

What to do next

- Create a set attribute remediation as described in Adding Set Attribute Value Remediations, on page 2042.

Adding Set Attribute Value Remediations

The Set Attribute Value remediation sets a host attribute on a host involved in a correlation policy violation. Create a remediation for each attribute value you want set. For text attributes, you can use the description from the triggering event as the attribute value.

In a multidomain deployment, you cannot add a remediation to an instance created in an ancestor domain.

Before you begin

- Create a set attribute instance as described in Adding a Set Attribute Value Instance, on page 2042.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Next to the instance where you want to add the remediation, click **View** ( ).

Step 3 In the **Configured Remediations** section, choose **Set Attribute Value** and click **Add**.

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Enter a **Remediation Name** and **Description**.

Step 5 To use this remediation in response to an event with source and destination data, choose an **Update Which Host(s) From Event** option.

Step 6 For text attributes, specify whether you want to **Use Description From Event For Attribute Value**:

- To use the description from the event as the attribute value, click **On** and enter the **Attribute Value** you want to set.
- To use the **Attribute Value** setting for the remediation as the attribute value, click **Off**.

Step 7 Click **Create**, then click **Done**.

What to do next

- Assign remediations as responses to correlation policy violations; see Adding Responses to Rules and White Lists, on page 1987.

Managing Remediation Modules

In a multidomain deployment, the system displays remediation modules installed in the current domain, which you can delete. It also displays modules installed in ancestor domains, which you cannot delete. To manage remediation modules in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Policies** > **Actions** > **Modules**.

Step 2 Manage your remediation modules:

- Configure — To view the Module Detail page for a module and configure its instances and remediations, click **View** ( ). In a multidomain deployment, you cannot use the Module Detail page to add, delete, or edit instances in the current domain for a module installed in an ancestor domain. Instead, use the Instances page (**Policies** > **Actions** > **Instances**); see Managing Remediation Instances, on page 2044.
- Delete — To delete a custom module that is not in use, click **Delete** ( ). You cannot delete system-provided modules.

- Install — To install a custom module, click **Choose File**, browse to the module, and click **Install**. For more information, see the *Firepower System Remediation API Guide*.

Managing Remediation Instances

The Instances page lists all configured instances for all remediation modules.

In a multidomain deployment, the system displays remediation instances created in the current domain, which you can edit. It also displays instances created in ancestor domains, which you cannot edit. To manage remediation instances in a lower domain, switch to that domain.

Though you cannot add a remediation to an instance created in an ancestor domain, you can create a similarly configured instance in the current domain and add remediations to that instance. You can also use remediations created in ancestor domains as correlation responses.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Manage your remediation instances:

- Add—To add an instance, choose the remediation module for which you want to add an instance and click **Add**. For system-provided modules, see:
  - Adding an ISE EPS Instance, on page 2035
  - Adding a Cisco IOS Instance, on page 2037
  - Adding an Nmap Scan Instance, on page 1863
  - Adding a Set Attribute Value Instance, on page 2042

  For help adding a custom module, see the documentation for that module, if available.
- Configure—To configure instance details and add remediations to the instance, click **View** ( ).
- Delete—To delete an instance that is not in use, click **Delete** ( ).

Managing Instances for a Single Remediation Module

The Module Detail page displays all of the instances and remediations configured for a particular remediation module.

In a multidomain deployment, you can access the Module Detail page for remediation modules installed in the current domain and in ancestor domains. However, you cannot use the Module Detail page to add, delete, or edit instances in the current domain for a module installed in an ancestor domain. Instead, use the Instances page (**Policies** > **Actions** > **Instances**); see Managing Remediation Instances, on page 2044 .

Procedure

Step 1 Choose **Policies** > **Actions** > **Modules**.

Step 2 Next to the remediation module whose instances you want to manage, click **View** ( ).

Step 3 Manage your remediation instances:

- Add — To add an instance, click **Add**. For system-provided modules, see:
  - Adding an ISE EPS Instance, on page 2035
  - Adding a Cisco IOS Instance, on page 2037
  - Adding an Nmap Scan Instance, on page 1863
  - Adding a Set Attribute Value Instance, on page 2042

  For help adding an instance for a custom module, see the documentation for that module, if available.
- Configure — To configure instance details and add remediations to the instance, click **View** ( ).
- Delete — To delete an instance that is not in use, click **Delete** ( ).



PART XXII

Reporting and Alerting

- Working with Reports, on page 2049
- External Alerting with Alert Responses, on page 2075
- External Alerting for Intrusion Events, on page 2085



CHAPTER 100

Working with Reports

The following topics describe how to work with reports in the Firepower System:

- Requirements and Prerequisites for Reports, on page 2049
- Introduction to Reports, on page 2049
- Risk Reports, on page 2050
- Standard Reports, on page 2051
- About Working with Generated Reports, on page 2071

Requirements and Prerequisites for Reports

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Maintenance User (risk reports only)
- Security Analyst

Introduction to Reports

The Firepower System offers two types of reports:

- Risk Reports, on page 2050 — High-level summaries of risks found on your network.
- Standard Reports, on page 2051 — Detailed, customizable reports about all aspects of your Firepower System.

Risk Reports

Risk reports are portable, high-level, easy-to-interpret summaries of risks found in your organization. You can use these reports to share information about areas of risk, and recommendations for addressing these risks, with people who do not have access to your system and who may not be network security experts. These reports are intended to facilitate discussion about areas for investment in the security of your network.

Generating, Viewing, and Printing Risk Reports

Templates for standard reports do not apply to risk reports.

Reports pertain to the current domain.

Each risk report generates as an HTML file.

To schedule risk report generation, see Automating Report Generation, on page 185.

Before you begin

- Make sure your system is configured to detect the risks that you want to summarize.
- If you want to email the report and you have not yet configured a Relay Host, you can do so now. For information, see Configuring a Mail Relay Host and Notification Address, on page 906.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Report Templates**.

Step 3 Click **Generate Report** for the desired report.

Step 4 Enter information.

- Information that you enter in the Input Parameters section will appear on the title page of the report. You can leave these fields blank.

Step 5 Click **Generate**.

Step 6 Click **OK**.

What to do next

- To view, download, move, or delete a risk report, see About Working with Generated Reports, on page 2071.
- You can print to PDF any risk report from most supported browsers. For best results, enable background colors, images, and optionally headers and footers, in the print or print preview settings of your browser. Supported page sizes are A4 and US letter.

Standard Reports

The Firepower System provides a flexible reporting system that allows you to quickly and easily generate multi-section reports with the event views or dashboards that appear on your Firepower Management Center. You can also design your own custom reports from scratch.

A report is a document file formatted in PDF, HTML, or CSV with the content you want to communicate. A report template specifies the data searches and formats for the report and its sections. The Firepower System includes a powerful report designer that automates the design of report templates. You can replicate the content of any event view table or dashboard graphic displayed in the web interface.

You can build as many report templates as you need. Each report template defines the individual sections in the report and specifies the database search that creates the report's content, as well as the presentation format (table, chart, detail view, and so on) and the time frame. Your template also specifies document attributes, such as the cover page and table of contents and whether the document pages have headers and footers (available only for reports in PDF format). You can export a report template in a single configuration package file and import it for reuse on another Firepower Management Center.

You can include input parameters in a template to expand its usefulness. Input parameters allow you to produce tailored variations of the same report. When you generate a report with input parameters, the generation process prompts you to enter a value for each input parameter. The values you type constrain the report contents on a one-time basis. For example, you can place an input parameter in the destination IP field of the search that produces an intrusion event report; at report generation time, you can specify a department's network segment when prompted for the destination IP address. The generated report then contains only information concerning that particular department.

About Designing Reports

Report Templates

You use report templates to define the content and format of the data in each of the report's sections, as well as the document attributes of the report file (cover page, table of contents, and page headers and footers). After you generate a report, the template stays available for reuse until you delete it.

Your reports contain one or more information sections. You choose the format (text, table, or chart) for each section individually. The format you select for a section may constrain the data that can be included. For example, you cannot show time-based information in certain tables using a pie chart format. You can change the data criteria or format of a section at any time to obtain optimum presentation.

You can base a report's initial design on a predefined event view, or you can start your design by importing content from any defined dashboard, workflow, or summary. You can also start with an empty template, adding sections and defining their attributes one by one.



Note In a multidomain deployment, you can view but not edit report templates belonging to ancestor domains. To generate reports from these templates, you must copy them to your current domain.

Report Template Fields

The following table describes the fields you can use to build a section in your report template. Not all fields are used in all types of sections; after you choose the section format, the system displays the appropriate fields.

Field Name	Section Types	Definition
Format	n/a	Choose the format of the section data: **Bar chart** (): Compares quantities of the selected variables. **Line chart** (): Shows trends/changes over time of a selected variable. Available only for time-based tables. **Pie chart** (): Shows each selected variable as a percentage of the whole. Variables with quantities of zero are dropped from the chart. Very small quantities are clustered into a category labeled **Other**. **Table view** (): Shows values of attributes for each record. Not available for summary or statistical data. **Detail view** (): Shows complex object data associated with certain events, such as packets (for intrusion events) and host profiles (for host events). This format is available only for certain event types that involve such objects. Output may degrade performance if large numbers are requested.
Table	All	Choose the table from which the section data is extracted.
Preset	All	Predefined searches. Select an appropriate preset to initialize the search criteria when you define a new search.
Search or Filter	All	For most tables, you can constrain a report using a predefined or saved **Search**. You can also create a new search by clicking **Edit** (). For the Application Statistics table, you use a user-defined application **Filter** to constrain a report.
X-Axis	Bar chart Line chart Pie chart	Available data for the X-axis of the selected chart. For line charts, the X-axis value is always **Time**. For bar and pie charts, you cannot select **Time** as the X-axis value.
Y-Axis	Bar chart Line chart Pie chart	Available data for the Y-axis of the selected chart.
Section Description	All	Descriptive text that precedes the search data in the section. Enter a combination of text and input parameters. The default for a new section is `$<Time Window>` and `$<Constraints>`.

Field Name	Section Types	Definition
Time Window	All	The time window for the data that appears in the section. If the section searches time-based tables, you can select the check box to inherit the report's global time window. Alternatively, you can set a specific time window for the section.
Maximum Results	Table view Detail view	The maximum number of matching records to include.
Results	Bar chart Pie chart	Choose either **Top** or **Bottom** and enter the number of matching records you want to use to build the chart.
Color	Bar chart Line chart	Colors for graphed data in the section.

Report Template Creation

A report template is a framework of sections, each independently built from its own database query.

You can build a new report template by creating a new template, using an existing template, basing a template off an event view, or importing a dashboard or workflow.

If you do not want to copy an existing report template, you can create an entirely new template. The first step in creating a template is to generate the framework that allows you to add and format the sections. Then, in the order you prefer, you design the individual template sections and set attributes for the report document.

Each template section consists of a dataset generated by a search or filter, and has a format specification (table, pie chart, and so on) that determines the mode of presentation. You further determine section content by selecting the fields in the data records you want to include in the output, as well as the time frame and number of records to show.



Note Use the section preview utility to check the column selection and output characteristics such as pie chart colors. It is not a reliable indicator of the correctness of your configured search.

The report you generate from the template has several document attributes that span all sections and control features, such as the cover page, headers and footers, page numbering, and so on.

Note that if you selected CSV as your document format, you have no document attributes to set.

If you identify a good model among your existing templates, you can copy the template and edit its attributes to create a new report template. Cisco also provides a set of predefined report templates, visible on the **Reports Tab** in the list of templates.

From an event view, you can create a report template and modify it to meet your needs. You can add additional sections, modify automatically included sections, and delete sections.

You can quickly create a new report by importing dashboards, workflows, and statistics summaries. The import creates a section for each widget graphic in your dashboard and each event view in your workflow. You can delete any unnecessary sections to focus on the most important information.

Creating a Custom Report Template

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Report Templates**.

Step 3 Click **Create Report Template**.

Step 4 Enter a name for your new template in the **Report Title** field.

Step 5 To add an input parameter to the report title, place your cursor in the title where the parameter value should appear, then click insert **Input Parameter** ( ).

Step 6 Use the set of add under the Report Sections title bar to insert sections as necessary.

Step 7 Configure section content as described in Report Template Configuration, on page 2056.

Tip You can click **Preview** at the bottom of the section window to view the column layout or graphic format you chose.

Step 8 Click **Advanced** to set attributes for PDF and HTML reports as described in Document Attributes in a Report Template, on page 2064.

Step 9 Click **Save**.

If you see an error, look for a yellow triangle beside the results value in each section. If you see any such triangles, do one of the following:

- For each field that displays a yellow triangle, mouse over the triangle and reduce the number of results to the number indicated.
- Click **Generate** and include an output format other than PDF.

Creating a Report Template from an Existing Template

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Report Templates**.

Step 3 Click **Copy** ( ) next to the report template you want to copy.

Step 4 In the **Report Title** field, enter a name.

Step 5 Make changes to the template as needed.

Step 6 Click **Save**.

Creating a Report Template from an Event View

Procedure

Step 1 Populate an event view with the events you want in the report:

- Use an event search to define the events you want to view.
- Drill down through a workflow until you have the appropriate events in your event view.

Step 2 From the event view page, click **Report Designer**.

The Report Sections page displays a section for each view in the captured workflow.

Step 3 Optionally, enter a new name in the **Report Title** field and click **Save**.

Step 4 You can:

- Add a cover page, table of contents, starting page number, or header and footer text — Click **Advanced** settings.
- Add page breaks — Click **Add Page Break** (), and drag the new page break object from the template bottom to the front of the section that should start the new page.
- Add text sections — Click **Add Text Section** (), and drag the new text section from the template bottom to the place where you want it to appear in the report template.
- Change the title of a section — Click the section title in the title bar, enter the section title, and click **OK**.
- Configure the report sections — Adjust the field settings in each section.

  Tip To view the current column layout or chart formatting for a section, click the section's **Preview** link.

- Exclude template sections from the report — Click **Delete** () in the section's title bar, and confirm the deletion.

  Note The last report section in some workflows contains detail views that show packets, host profiles, or vulnerabilities, depending on the workflow. Retrieving large numbers of events with these detail views when generating your report may affect performance of the Firepower Management Center.

Step 5 Click **Save**.

Creating a Report Template by Importing a Dashboard or Workflow

Procedure

Step 1 Identify the dashboard, workflow, or summary you want to replicate in your report.

Step 2 Choose **Overview** > **Reporting**.

Step 3 Click **Report Templates**.

Step 4 Click **Create Report Template**.

Step 5 Enter a name for your new report template in the **Report Title** field.

Step 6 Click **Save**.

Step 7 Click **Import Section** (). You can choose any of the data sources described in Data Source Options on Import Report Sections, on page 2056.

Step 8 Choose a dashboard, workflow, or summary from the drop-down menus.

Step 9 For the data sources you want to add, click **Import**.

For dashboards, each widget graphic will have its own section; for workflows, each event view will have its own section.

Step 10 Make changes to the content of your sections as needed.

Note The last report section in some workflows contains detail views that show packets, host profiles, or vulnerabilities, depending on the workflow. Retrieving large numbers of events with these detail views when generating your report may affect performance of the Firepower Management Center.

Step 11 Click **Save**.

Data Source Options on Import Report Sections

Table 242: Data Source Options on Import Report Sections Window

Select this option...	To import...
Import Dashboard	any custom analysis widget on the selected dashboard.
Import Workflow	any predefined or custom workflow. Selections have the format: `Table - Workflow name` For example, `Connection Events - Traffic by Port` imports the views in the Traffic by Port workflow generated from the Connection Events table.
Import Summary Sections	any of the following generic summaries: • Intrusion Detailed Summary • Intrusion Short Summary • Discovery Detailed Summary • Discovery Short Summary

Report Template Configuration

You can modify and customize a report template once you create it. You can modify a variety of report section attributes to adjust the content of the section and its data presentation.

Each section in a report template queries a database table to generate content for that section. Changing the section's data format uses the same data query, but modifies the fields that appear in the section according to the analytical purpose of the format type. For example, the table view of intrusion events populates the section

with a large number of data fields per event record, while a pie chart section shows the portion of all matching records that each selected attribute represents, with no details about individual events. Bar chart sections compare the total counts of matching records that have specific attributes. Line charts summarize changes in the matching records over time with respect to a single attribute. Line charts are available only for data that is time-based, not for information about hosts, users, third-party vulnerabilities, and so on.

The search or filter in a report section specifies the database query on which the section content is based. For most tables, you can constrain a report using a predefined or saved search, or you can create a new search on the fly:

- Predefined searches serve as examples for searching certain event tables and can provide quick access to important information about your network that you may want to include in reports.
- Saved event searches include all public event searches that you or others have created, plus all your saved private event searches.
- Saved searches for the current report template are accessible only in the report template itself. The search names of saved report template searches end with the string "Custom Search." Users create these searches while designing reports.

For the Application Statistics table, you use a user-defined application filter to constrain a report.

If you include table data in a section, you can choose which fields in the data record to show. All fields in the table are available for inclusion or exclusion. You select fields that accomplish the purpose of the report, then order and sort them accordingly.

You can add text sections to your templates to provide custom text, such as an introduction, for the whole report or for individual sections.

You can add page breaks before or after any section in the template. This feature is particularly helpful for multi-section reports with text pages that introduce the various sections.

A report template's time window defines the template's reporting period.



Note Security Analysts can edit only report templates they created. In multidomain deployments, you cannot edit report templates from ancestor domains, but you can copy them to create descendant versions.

Setting the Table and Data Format for a Report Template Section

Procedure

Step 1 In the report template section, use the **Table** drop-down menu to choose the table to query.

The **Format** field represents each of the output formats available for the table you chose.

Step 2 Choose the applicable output format for the section.

Step 3 To change the search constraints, click **Edit** () next to the **Section description field** or **Filter** field.

Step 4 For graphic output formats (pie chart, bar chart, and so on), adjust the **X-Axis** and **Y-Axis** parameters using the drop-down menus.

When you choose a value for the X-axis, only compatible values appear in the Y-axis drop-down menu, and vice versa.

Step 5 For table output, choose the columns, order of appearance, and sort order in your output.

Step 6 Click **Save**.

Related Topics

Report Template Fields, on page 2052

Specifying the Search or Filter for a Report Template Section

Procedure

Step 1 In the report template section, choose the database table to query from the **Table** drop-down menu:

- For most tables, the **Search** drop-down list appears.
- For the Application Statistics table, the **Filter** drop-down list appears.

Step 2 Choose the search or filter you want to use to constrain the report.

You can view the search criteria or create a new search by clicking **Edit** ().

Related Topics

Application Filters, on page 406

Setting the Search Fields that Appear in Table Format Sections

Procedure

Step 1 For table format report sections, click **Edit** () next to the **Fields** parameter.

Step 2 If you want to modify the section, you must add and delete fields, and drag field into the column order you want.

Step 3 If you want to change the sort order of any column, you must use the drop-down lists on each field to set the sort order and priority.

Step 4 Click **OK**.

Adding a Text Section to a Report Template

Text sections can have rich text with multiple font sizes and styles (bold, italic, and so on) as well as input parameters and imported images.



Tip Text sections are useful for introductions to your report or your report sections.

Procedure

Step 1 In the report template editor, click **Add Text Section** ().

Step 2 Drag the new text section to its intended position in the report template.

Step 3 If you want to position the text section first or last on a page, add page breaks before or after the text section.

Step 4 If you want to change the text section's generic name, click section's name in the title bar, and enter a new name.

Step 5 Add formatted text and images to the body of the text section.

You can include input parameters that dynamically update when you generate the report.

Step 6 Click **Save**.

Related Topics

Input Parameters, on page 2061

Adding a Page Break to a Report Template

Procedure

Step 1 In the report template editor, click **Add Page Break** ().

A page break appears at the bottom of the template.

Step 2 Drag the page break to its intended location, before or after a section.

Step 3 Click **Save**.

Global Time Windows and Report Template Sections

Report templates with time-based data (such as intrusion or discovery events) have a global time window, which the time-based sections in the template inherit by default when created. Changing the global time window changes the local time window for the sections that are configured to inherit the global time window. You can disable time window inheritance for an individual section by clearing its **Inherit Time Window** check box. You can then edit the local time window.



Note Global time window inheritance applies only to report sections with data from time-based tables, such as intrusion events and discovery events. For sections that report on network assets (hosts and devices) and related information (such as vulnerabilities), you must set each time window individually.

Setting the Global Time Window for a Report Template and Its Sections



Tip Your report can have different time ranges per section. For example, your first section could be a summary for the month, and the remaining sections could drill down into details at the week level. In such cases, you set the section-level time windows individually.

Procedure

Step 1 In the report template editor, click **Generate**.

Step 2 To modify the global time window, click **Time Window** ().

Step 3 Modify time settings in **Events Time Window**.

Step 4 Click **Apply**.

Step 5 Click **Generate** to generate the report and **Yes** to confirm.

Setting the Local Time Window for Report Template Sections

Procedure

Step 1 On the Report Sections page of a template, clear the **Inherit Time Window** check box for the section if it is present.

Step 2 To change the section's local time window, click **Time Window** ().

Note Sections with data from statistics tables can have only sliding time windows.

Step 3 Click **Apply** on the Events Time Window.

Step 4 Click **Save**.

Renaming a Report Template Section

Procedure

Step 1 In the report template editor, click the current section name in the section header.

Step 2 Enter a new name for the section.

Step 3 Click **OK**.

Previewing a Report Template Section

The preview function shows the field layout and sort order for table views and important legibility characteristics of graphics, such as pie chart colors.

Procedure

Step 1 At any time while editing a report template section, click **Preview** for the section.

Step 2 Close the preview by clicking **OK**.

Searches in Report Template Sections

The key to generating successful reports is defining the searches that populate the report's sections. The Firepower System provides a search editor to view the searches available in your report templates and to define new custom searches.

Searching in Report Template Sections

Procedure

Step 1 From the relevant section in the report template, click **Edit** ( ) next to the **Search** field.

Step 2 If you want to base a custom search on a predefined search, you must choose a predefined search from the **Saved Searches** drop-down list.

This list includes all available predefined searches for this table, including system-wide and report-specific predefined searches.

Step 3 Edit the search criteria in the appropriate fields.

For certain fields, your constraints can include the same operators (<, <>, and so on) as event searches. If you enter multiple criteria, the search returns only the records that match all the criteria.

Step 4 If you want to insert an input parameter from the drop-down menu instead of entering a constraint value, you must click **Input Parameter** ( ).

Note When you edit the constraints of a reporting search, the system saves your edited search under the following name: *section* `custom search`, where *section* is the name in the section title bar followed by the string `custom search`. To have meaningful names for your saved custom searches, be sure you change the section name before you save the edited search. You cannot rename a saved reporting search.

Step 5 Click **OK.**

Input Parameters

You can use input parameters in a report template that the report can dynamically update at generation time. The **Input Parameter** ( ) indicates the fields that can process them. There are two kinds of input parameters:

- *Predefined input parameters* are resolved by internal system functions or configuration information. For example, at report generation time, the system replaces the `$<Time>` parameter with the current date and time.

- *User-defined input parameters* supply constraints in section searches. Constraining a search with an input parameter instructs the system to collect a value at generation time from the person who requests the report. In this way, you can dynamically tailor a report at generation time to show a particular subset of data without changing the template. For example, you can provide an input parameter for the **Destination IP** field of a report section's search. Then, when you generate the report, you can enter the IP network segment for a particular department to get data for that department only.

  You can also define string-type input parameters to add dynamic text in certain fields of your report, such as in emails (subject or body), report file names, and text sections. You can personalize reports for different departments, with customized report file names, email addresses, and email messages, using the same template for all.

Predefined Input Parameters

Table 243: Predefined Input Parameters

Insert this parameter...	...to include this information in your template:
`$<Logo>`	The selected uploaded logo
`$<Report Title>`	The report title
`$<Time>`	The date and time of day the report ran, with one-second granularity
`$<Month>`	The current month
`$<Year>`	The current year
`$<System Name>`	The name of the Firepower Management Center
`$<Model Number>`	The model number of the Firepower Management Center
`$<Time Window>`	The time window currently applied to the report section
`$<Constraints>`	The search constraints currently applied to the report section

Table 244: Predefined Input Parameter Usage

Parameter	Report Template Cover Page	Report Template Report Title	Report Template Section Description	Report Template Text Section	Generate Report File Name	Generate Report Email Subject, Body
`$<Logo>`	yes	no	no	no	no	no
`$<Report Title>`	yes	no	yes	yes	yes	yes
`$<Time>`	yes	yes	yes	yes	yes	yes
`$<Month>`	yes	yes	yes	yes	yes	yes
`$<Year>`	yes	yes	yes	yes	yes	yes
`$<System Name>`	yes	yes	yes	yes	yes	yes

Parameter	Report Template Cover Page	Report Template Report Title	Report Template Section Description	Report Template Text Section	Generate Report File Name	Generate Report Email Subject, Body
`$<Model Number>`	yes	yes	yes	yes	yes	yes
`$<Time Window>`	no	no	yes	no	no	no
`$<Constraints>`	no	no	yes	no	no	no

User-Defined Input Parameters

You use input parameters to expand the usefulness of your searches. The input parameter instructs the system to collect a value at generation time from the person who requests the report. In this way, you can dynamically constrain a report at generation time to show a particular subset of data without changing the search. For example, you can provide an input parameter for the **Destination IP** field of a report section that drills down on security events at a department level. When you generate the report, you can type the IP network segment for a particular department to get data for that department only.

An input parameter's type determines the search fields where you can use it. You can use a given type only in appropriate fields. For example, a user parameter you define as a string type is available for insertion in text fields but not in fields that take an IP address.

Each input parameter you define has a name and a type.

Table 245: User-Defined Input Parameter Types

Use this parameter type...	With fields with this data...
Network/IP	any IP address or network segment in CIDR format
Application	name of an application protocol, client application, or web application
Event Message	any event view message
Device	a FMC or managed device
Username	user identification such as initiator user and responder user
Number (VLAN ID, Snort ID, Vuln ID)	any VLAN ID, Snort ID, or vulnerability ID
String	text fields such as application or OS version, notes, or descriptions

Creating User-Defined Input Parameters

Procedure

Step 1 In the report template editor, click **Advanced**.

Step 2 Click **Add Input Parameter** ().

Step 3 Enter the parameter **Name**.

Step 4 Choose a value from the **Type** drop-down list.

Step 5 Click **OK** to add the parameter.

Step 6 Click **OK** to return to the editor.

Editing User-Defined Input Parameters

The **Input Parameters** section of the report template lists all available user-defined parameters for the template.

Procedure

Step 1 In the report template editor, click **Advanced**.

Step 2 Click **Edit** ( ) next to the parameter you want to modify.

Step 3 Enter a new **Name**.

Step 4 Use the **Type** drop-down list to change the parameter type.

Step 5 Click **OK** to save your changes.

Step 6 If you want to delete an input parameter, click **Delete** ( ) next to the input parameter and confirm.

Step 7 Click **OK** to return to the report template editor.

Constraining a Search with User-Defined Input Parameters

Input parameters you define are available only for search fields that match their parameter type. For example, a parameter of type **Network/IP** is available only for fields that accept IP addresses or network segments in CIDR format.

Procedure

Step 1 In the report template editor, click **Edit** ( ) next to the **Search** field within the section.

Fields that can take an input parameter are marked with **Input Parameter** ( ).

Step 2 Click **Input Parameter** ( ) next to the field, then choose the input parameter from the drop-down menu.

User-defined input parameters are marked with ( ).

Step 3 Click **OK**.

Document Attributes in a Report Template

Before you generate your report, you can set document attributes that affect the report's appearance. These attributes include the optional cover page and table of contents. Support for some attributes depends on the selected report format: PDF, HTML, or CSV.

Table 246: Document Attribute Support

Attribute	PDF Support?	HTML Support?	CSV Support?
Cover page	yes, with optional logo and custom appearance	yes, with optional logo and custom appearance	no
Table of contents	yes	yes	no
Page headers and footers	yes, with optional text or logo in any field	no	no
Custom starting page number	yes	no	no
Option to suppress numbering of first page	yes	no	no

Editing Document Attributes in a Report Template

Procedure

Step 1 In the report template editor, click **Advanced**.

Step 2 You have the following choices:

- Add cover page —To add a cover page, check the **Include Cover Page** check box.
- Customize cover page —To edit the cover page design, see Customizing a Cover Page, on page 2065.
- Add table of contents — To add a table of contents, check the **Include Table of Contents** check box.
- Manage logos — To manage the logo image associated with the template, see Managing Report Template Logos, on page 2066.
- Configure header and footer —To specify elements for the header and footer of the template, use the drop-down lists in the **Header** and **Footer** fields.
- Set first page number — To specify the page number of the report's first page, enter a **Page Number Start** value.
- Show first page number —To show the page number on the report's first page, check the **Number First Page?** check box. If you choose this option, the cover page is not numbered.

Step 3 Click **OK** to save your changes.

Customizing a Cover Page

You can customize a report template's cover page. Cover pages can have rich text with multiple font sizes and styles (bold, italic, and so on) as well as input parameters and imported images.

Procedure

Step 1 In the report template editor, click **Advanced**.

Step 2 Click **Edit** ( ) next to **Cover Page Design**.

Step 3 Edit the cover page design within the rich text editor.

Step 4 Click **OK**.

Managing Report Template Logos

You can store multiple logos on the Firepower Management Center and associate them with different report templates. You set the logo association when you design the template. If you export the template, the export package contains the logo.

When you upload a logo to the Firepower Management Center, it is available for:

- all report templates on the Firepower Management Center, or
- in a multidomain deployment, all report templates in your current domain

Logo images can be in GIF, JPG, or PNG format.

You can change the logo in a report to any JPG image uploaded to your Firepower Management Center. For example, if you reuse a template, you can associate a logo for a different organization with the report.

You can delete any uploaded logos. Deleting a logo removes it from all templates where it is used. The deletion cannot be undone. Note that you cannot delete the predefined Cisco logo.

Procedure

Step 1 In the report template editor, click **Advanced**.

The logo currently associated with the template appears under **Logo** in **General Settings**.

Step 2 Click **Edit** ( ) next to the logo.

Step 3 You have the following choices:

- Add — Add a new logo as described in Adding a New Logo, on page 2066.
- Change — Change a report template's logo as described in Changing the Logo for a Report Template, on page 2067.
- Delete — Delete a logo as described in Deleting a Logo, on page 2067.

Adding a New Logo

Procedure

Step 1 In the report template editor, click **Advanced**.

Step 2 Click **Edit** ( ) next to the **Logo** field.

Step 3 Click **Upload Logo**.

Step 4 Click **Browse**, browse to the file's location, and click **Open**.

Step 5 Click **Upload**.

Step 6 If you want to associate the new logo with the current template, choose it, and click **OK**.

Changing the Logo for a Report Template

Procedure

Step 1 In the report template editor, click **Advanced**.

Step 2 Click **Edit** ( ) next to the **Logo** field.

Step 3 From the Select Logo dialog, choose the logo to associate with the report template.

Step 4 Click **OK**.

Deleting a Logo

Procedure

Step 1 In the report template editor, click **Advanced**.

Step 2 Click **Edit** ( ) next to the **Logo** field.

Step 3 From the Select Logo dialog, choose the logo you want to delete.

Step 4 Click **Delete Logo**.

Step 5 Click **OK**.

Managing Report Templates

In a multidomain deployment, the system displays report templates created in the current domain, which you can edit. It also displays report templates created in ancestor domains, which you cannot edit. To view and edit report templates in a lower domain, switch to that domain. The system displays reports created in the current domain only.

You must be an Admin user to perform this task.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Report Templates**.

Step 3 You have the following choices:

- Delete — Next to the template you want to delete, click **Delete** ( ) and confirm.

You cannot delete system-provided report templates. Security Analysts can delete only report templates they created. In a multidomain deployment, you can delete report templates belonging to the current domain only.

- Edit — To edit report templates; see Editing Report Templates, on page 2068.
- Export — To export report templates, see Exporting Report Templates, on page 2069.

  Tip You can also export report templates using the standard configuration export process; see Exporting Configurations, on page 175.

- Import — To import report templates, see Importing Configurations, on page 176.

Editing Report Templates

In a multidomain deployment, the system displays report templates created in the current domain, which you can edit. It also displays report templates created in ancestor domains, which you cannot edit. To view and edit report templates in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Report Templates**.

Step 3 Click **Edit** ( ) for the template you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 You have the following choices:

- Add a page break; see Adding a Page Break to a Report Template, on page 2059.
- Add a text section; see Adding a Text Section to a Report Template, on page 2058.
- Configure section content as described in Report Template Configuration, on page 2056.
- Create input parameters; see Creating User-Defined Input Parameters, on page 2063.
- Edit input parameters; see Editing User-Defined Input Parameters, on page 2064.
- Edit document attributes; see Editing Document Attributes in a Report Template, on page 2065.
- Search template sections; see Searching in Report Template Sections, on page 2061.
- Set document attributes described in Document Attributes in a Report Template, on page 2064 by clicking **Advanced**.
- Set the global time window; see Setting the Global Time Window for a Report Template and Its Sections, on page 2060.
- Set the local time window; see Setting the Local Time Window for Report Template Sections, on page 2060.
- Set the search fields; see Setting the Search Fields that Appear in Table Format Sections, on page 2058.
- Set the table and data format; see Setting the Table and Data Format for a Report Template Section, on page 2057.

- Specify searches and filters; see Specifying the Search or Filter for a Report Template Section, on page 2058.

Exporting Report Templates

You must be an Admin user to perform this task.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Choose **Report Templates**.

Step 3 For the template you want to export, click **YouTube EDU** ( ).

Step 4 Click **Save file** and **OK** to save the file to your local computer.

About Generating Reports

Generating Reports

After you create and customize your report template, you are ready to generate the report. The generation process lets you select the report's format (HTML, PDF, or CSV). You can also adjust the report's global time window, which applies a consistent time frame to all sections except those you exempt.

For PDF reports:

- File names using Unicode (UTF-8) characters are not supported.
- Any report sections that include special Unicode file names (such as those appearing in file or malware events) display these file names in transliterated form.

If the report template includes user input parameters in its search specification, the generation process prompts you to enter values, which tailor this run of the report to a subset of the data.

If you have a DNS server configured and IP address resolution enabled, reports contain host names if resolution was successful.

In a multidomain deployment, when you generate a report in an ancestor domain, it can include results from all descendant domains. To generate a report for a specific leaf domain, switch to that domain.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Report Templates**.

Step 3 Click **Report** ( ) next to the template you want to use to generate a report.

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Tip To generate a report from an ancestor's template, copy the template into the current domain.

Step 4 Optionally, configure the report name:

- Enter a new **File Name**. If you do not enter a new name, the system uses the name specified in the report template.
- Use **Input Parameter** ( ) to add one or more input parameters to the file name.

Step 5 Choose the output format for the report by clicking: HTML, PDF, or CSV.

Step 6 If you want to change the global time window, click **Time Window** ( ).

Note Setting the global time window affects the content of individual report sections only if they are configured to inherit the global setting.

Step 7 Enter values for any fields that appear in the **Input Parameters** section.

Tip You can ignore user parameters by typing the * wildcard character in the field. This eliminates the user parameter's constraint on the search.

Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses or VLAN tags to constrain report results can have unexpected results.

Step 8 If you enabled an email relay host in the Firepower Management Center configuration, click **Email** to automate email delivery of the report when it generates.

Step 9 Click **Generate** and confirm when prompted.

Clicking **Generate** saves Generate settings with the report template.

If you click **Cancel**, your selections are saved only for the duration of your session.

Step 10 You have the following choices:

- Click the report link to display the report in a new window.
- Click **OK** to return to the report template editor.

Report Generation Options

You can configure report generation options to:

- Schedule generation of future reports, either once or recurring. See Automating Report Generation, on page 185. You can customize the schedule on a full range of time frames such as daily, weekly, monthly, and so on.
- Distribute email reports using the scheduler. You must configure your report template and a mail relay host **before** scheduling the task.
- Automatically send the report as an email attachment to a list of recipients when you generate a report. You must have a properly configured mail relay host to deliver a report by email.
- Save newly generated report files to your configured remote storage location. To use remote storage, you must first configure a remote storage location.

Note If you store remotely and then switch back to local storage, the reports in remote storage do not appear on the Reports tab list. Similarly, if you switch from one remote storage location to another, the reports in the previous location do not appear in the list.

Distributing Reports by Email at Generation Time

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Report Templates**.

Step 3 Click **Report** ( ) next to the template you want to use to generate a report.

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Tip To generate a report from an ancestor's template, copy the template into the current domain.

Step 4 Expand the **Email** section of the window.

Step 5 In the **Email Options** field, choose **Send Email**.

Step 6 In the **Recipient List**, **CC**, and **BCC** fields, enter recipients' email addresses in comma-separated lists.

Step 7 In the **Subject** field, enter an email subject.

Tip You can provide input parameters in the **Subject** field and the message body to dynamically generate information in the email, such as a timestamp or the name of the Firepower Management Center.

Step 8 Enter a cover letter in the email body as necessary.

Step 9 Click **OK** and confirm.

Related Topics

Configuring a Mail Relay Host and Notification Address, on page 906

Schedule Future Reports

See Automating Report Generation, on page 185.

About Working with Generated Reports

Access and work with previously-generated reports on the Reports tab page.

Viewing Reports

The Reports lists all previously generated reports, with report name, date and time of generation, generating user, and whether the report is stored locally or remotely. A status column indicates whether the report is already generated, is in the generation queue (for example, for scheduled tasks), or failed to generate (for example, due to lack of disk space).

Note that users with Administrator access can view all reports; other users can view only the reports they generated.

In a multidomain deployment, you can view reports generated in the current domain only.

The Reports page shows all locally stored reports. It shows remotely stored reports as well, if remote storage is currently configured. The **Location** column data for remotely-stored reports is `Remote`.



Note If you store remotely and then switch back to local storage, the reports in remote storage do not appear on the Reports tab list. Similarly, if you switch from one remote storage location to another, the reports in the previous location do not appear in the list.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Reports**.

Step 3 Click the report you want to view.

Downloading Reports

You can download any report file to your local computer. From there, you can email it or distribute it electronically by other available means.

In a multidomain deployment, you can download reports generated in the current domain only.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Reports**.

Step 3 Check the check boxes next to the reports you want to download, then click **Download**.

Tip Click the check box at the top left of the page to download all reports on the page. If you have multiple pages of reports, a second check box appears that you can click to download all reports on all pages.

Step 4 Follow your browser's prompts to download the reports. If you chose multiple reports, they are downloaded in a single `.zip` file.

Storing Reports Remotely

The location of your currently configured report storage appears at the bottom of the Overview> Reporting > Reports page, with disk usage for local, NFS, and SMB storage. If you access remote storage using SSH, disk usage data is not available.

Note If you store remotely and then switch back to local storage, the reports in remote storage do not appear on the Reports tab list. Similarly, if you switch from one remote storage location to another, the reports in the previous location do not appear in the list.

Before you begin

- Configure a remote storage location as described in Remote Storage Management, on page 890.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Choose **Reports**.

Step 3 Check the **Enable Remote Storage of Reports** check box at the bottom of the page.

What to do next

- Move reports from local storage to remote storage; see Moving Reports to Remote Storage, on page 2073.

Related Topics

Remote Storage Management, on page 890
Moving Reports to Remote Storage, on page 2073

Moving Reports to Remote Storage

You can move your reports in local storage to a remote storage location in batch mode or singly.

Note If you store remotely and then switch back to local storage, the reports in remote storage do not appear on the Reports tab list. Similarly, if you switch from one remote storage location to another, the reports in the previous location do not appear in the list.

Before you begin

- Configure a remote storage location as described in Remote Storage Management, on page 890.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Choose **Reports**.

Step 3 Choose the check boxes next to the reports you want to move, then click **Move**.

Tip Check the check box at the top left of the page to move all reports on the page. If you have multiple pages of reports, a second check box appears that you can check to move all reports on all pages.

Step 4 Confirm that you want to move the reports.

Deleting Reports

You can delete your report files at any time. The procedure completely removes the files, and no recovery is possible. Although you still have the report template that generated the report, it may be difficult to regenerate a particular report file if the time window was expanding or sliding. Regeneration may also be difficult if your template uses input parameters.

In a multidomain deployment, you can delete reports generated in the current domain only.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Reports**.

Step 3 You have the following choices:

- Delete selected — Check the check boxes next to the reports you want to delete, then click **Delete**.
- Delete all — Check the check box at the top left of the page to delete all reports on the page. If you have multiple pages of reports, a second check box appears that you can check to delete all reports on all pages.

Step 4 Confirm the deletion.



CHAPTER 101

External Alerting with Alert Responses

The following topics describe how to send external event alerts from the Firepower Management Center using alert responses:

- Firepower Management Center Alert Responses, on page 2075
- Requirements and Prerequisites for Alert Responses, on page 2076
- Creating an SNMP Alert Response, on page 2076
- Creating a Syslog Alert Response, on page 2078
- Creating an Email Alert Response, on page 2081
- Configuring Impact Flag Alerting, on page 2081
- Configuring Discovery Event Alerting, on page 2082
- Configuring AMP for Networks Alerting, on page 2082

Firepower Management Center Alert Responses

External event notification via SNMP, syslog, or email can help with critical-system monitoring. The Firepower Management Center uses configurable *alert responses* to interact with external servers. An *alert response* is a configuration that represents a connection to an email, SNMP, or syslog server. They are called *responses* because you can use them to send alerts in response to events detected by Firepower. You can configure multiple alert responses to send different types of alerts to different monitoring servers and/or people.

Note Alerts that use alert responses are sent by the Firepower Management Center. Intrusion email alerts, which do not use alert responses, are also sent by the Firepower Management Center. By contrast, SNMP and syslog alerts that are based on individual intrusion rules triggering are sent directly by managed devices. For more information, see External Alerting for Intrusion Events, on page 2085.

In most cases, the information in an external alert is the same as the information in any associated event you logged to the database. However, for correlation event alerts where the correlation rule contains a connection tracker, the information you receive is the same as for an alert on a traffic profile change, regardless of the base event type.

You create and manage alert responses on the Alerts page (**Policies** > **Actions** > **Alerts**). New alert responses are automatically enabled. To temporarily stop alert generation, you can disable alert responses rather than deleting them.

Changes to alert responses take effect immediately, except when sending connection logs to an SNMP trap or syslog server.

In a multidomain deployment, when you create an alert response it belongs to the current domain. This alert response can also be used by descendant domains.

Configurations Supporting Alert Responses

After you create an alert reponse, you can use it to send the following external alerts from the Firepower Management Center.

Alert/Event Type	For More Information
Intrusion events, by impact flag	Configuring Impact Flag Alerting, on page 2081
Discovery events, by type	Configuring Discovery Event Alerting, on page 2082
Malware and retrospective malware events detected by AMP for Networks ("network-based")	Configuring AMP for Networks Alerting, on page 2082
Correlation events, by correlation policy violation	Adding Responses to Rules and White Lists, on page 1987
Connection events, by the logging rule or default action (email alerts not supported)	Other Connections You Can Log, on page 2209
Health events, by health module and severity level	Creating Health Monitor Alerts, on page 289

Requirements and Prerequisites for Alert Responses

Model Support

Any.

Supported Domains

Any

User Roles

- Admin

Creating an SNMP Alert Response

You can create SNMP alert responses using SNMPv1, SNMPv2, or SNMPv3 for an device type except Firepower Threat Defense.

Note When selecting SNMP versions for the SNMP protocol, note that SNMPv2 only supports read-only communities and SNMPv3 only supports read-only users. SNMPv3 also supports encryption with AES128.

If you want to monitor 64-bit values with SNMP, you must use SNMPv2 or SNMPv3. SNMPv1 does not support 64-bit monitoring.

Before you begin

- If your network management system requires the Firepower Management Center's management information base (MIB) file, obtain it at `/etc/sf/DCEALERT.MIB`.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 From the **Create Alert** drop-down menu, choose **Create SNMP Alert**.

Step 3 Edit the SNMP Alert Configuration fields:

a) **Name—**Enter a name to identify the SNMP response.

b) **Trap Server—**Enter the hostname or IP address of the SNMP trap server.

Note The system does **not** warn you if you enter an invalid IPv4 address (such as 192.169.1.456) in this field. Instead, the invalid address is treated as a hostname.

c) **Version—**Choose the SNMP version you want to use from the drop-down list. SNMPv3 is the default.

Choose from:

- **SNMPv1** or **SNMPv2**: Enter a read-only SNMP community name in the **Community String** field, then skip to the end of the procedure.

Note Do not include special characters (< > / % # & ? ', etc.) in the SNMP community string name.

- For **SNMPv3**: Enter the name of the user that you want to authenticate with the SNMP server in the **User Name** field and continue to the next step.

d) **Authentication Protocol—**Choose the protocol you want to use to encrypt authentication from the drop-down list.

Choose from:

- **MD5**—Message Digest 5 (MD5) hash function.
- **SHA**—Secure Hash Algorithm (SHA) hash function.

e) **Authentication Password—**Enter the password to enable authentication.

f) **Privacy Protocol—**Choose the protocol you want to use to encrypt a private password from the drop-down list.

Choose from:

- **DES**—Data Encryption Standard (DES) using 56-bit keys in a symmetric secret-key block algorithm.
- **AES**—Advanced Encryption Standard (AES) using 56-bit keys in a symmetric cipher algorithm.
- **AES128**—AES using 128-bit keys in a symmetric cipher algorithm. A longer key provides higher security but a reduction in performance.

g) **Privacy Password—**Enter the privacy password required by the SNMP server. If you specify a private password, privacy is enabled, and you must also specify an authentication password.

h) **Engine ID—**Enter an identifier for the SNMP engine, in hexadecimal notation, using an even number of digits.

When you use SNMPv3, the system uses an Engine ID value to encode the message. Your SNMP server requires this value to decode the message.

Cisco recommends that you use the hexadecimal version of the Firepower Management Center's IP address. For example, if the Firepower Management Center has an IP address of 10.1.1.77, use 0a01014D0.

Step 4 Click **Save**.

What to do next

Changes take effect immediately, EXCEPT:

If you are using alert responses to send connection logs, you must deploy configuration changes after you edit those alert responses.

Creating a Syslog Alert Response

When configuring a syslog alert response, you can specify the severity and facility associated with the syslog messages to ensure that they are processed properly by the syslog server. The facility indicates the subsystem that creates the message and the severity defines the severity of the message. Facilities and severities are not displayed in the actual message that appears in the syslog, but are instead used to tell the system that receives the syslog message how to categorize it.



Tip For more detailed information about how syslog works and how to configure it, refer to the documentation for your system. On UNIX systems, the `man` pages for `syslog` and `syslog.conf` provide conceptual information and configuration instructions.

Although you can choose any type of facility when creating a syslog alert response, you should choose one that makes sense based on your syslog server; not all syslog servers support all facilities. For UNIX syslog servers, the `syslog.conf` file should indicate which facilities are saved to which log files on the server.

Before you begin

- Confirm that the syslog server can accept remote messages.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 From the **Create Alert** drop-down menu, choose **Create Syslog Alert**.

Step 3 Enter a **Name** for the alert.

Step 4 In the **Host** field, enter the hostname or IP address of your syslog server.

Note The system does **not** warn you if you enter an invalid IPv4 address (such as 192.168.1.456) in this field. Instead, the invalid address is treated as a hostname.

Step 5 In the **Port** field, enter the port the server uses for syslog messages. By default, this value is 514.

Step 6 From the **Facility** list, choose a facility described in Syslog Alert Facilities, on page 2079.

Step 7 From the **Severity** list, choose a severity described in Syslog Severity Levels, on page 2080.

Step 8 In the **Tag** field, enter the tag name that you want to appear with the syslog message.

For example, if you wanted all messages sent to the syslog to be preceded with `FromMC`, enter `FromMC` in the field.

Step 9 Click **Save**.

What to do next

Changes take effect immediately, EXCEPT:

If you are using alert responses to send connection logs to a syslog server, you must deploy configuration changes after you edit those alert responses.

Syslog Alert Facilities

The following table lists the syslog facilities you can select.

Table 247: Available Syslog Facilities

Facility	Description
ALERT	An alert message.
AUDIT	A message generated by the audit subsystem.
AUTH	A message associated with security and authorization.
AUTHPRIV	A restricted access message associated with security and authorization. On many systems, these messages are forwarded to a secure file.
CLOCK	A message generated by the clock daemon. Note that syslog servers running a Windows operating system will use the `CLOCK` facility.

Facility	Description
CRON	A message generated by the clock daemon. Note that syslog servers running a Linux operating system will use the `CRON` facility.
DAEMON	A message generated by a system daemon.
FTP	A message generated by the FTP daemon.
KERN	A message generated by the kernel. On many systems, these messages are printed to the console when they appear.
LOCAL0-LOCAL7	A message generated by an internal process.
LPR	A message generated by the printing subsystem.
MAIL	A message generated by a mail system.
NEWS	A message generated by the network news subsystem.
NTP	A message generated by the NTP daemon.
SYSLOG	A message generated by the syslog daemon.
USER	A message generated by a user-level process.
UUCP	A message generated by the UUCP subsystem.

Syslog Severity Levels

The following table lists the standard syslog severity levels you can select.

Table 248: Syslog Severity Levels

Level	Description
ALERT	A condition that should be corrected immediately.
CRIT	A critical condition.
DEBUG	Messages that contain debugging information.
EMERG	A panic condition broadcast to all users.
ERR	An error condition.
INFO	Informational messages.
NOTICE	Conditions that are not error conditions, but require attention.
WARNING	Warning messages.

Creating an Email Alert Response

Before you begin

- Confirm that the Firepower Management Center can reverse-resolve its own IP address.
- Configure your mail relay host as described in Configuring a Mail Relay Host and Notification Address, on page 906.

Note You **cannot** use email alerting to log connections.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 From the **Create Alert** drop-down menu, choose **Create Email Alert**.

Step 3 Enter a **Name** for the alert response.

Step 4 In the **To** field, enter the email addresses where you want to send alerts, separated by commas.

Step 5 In the **From** field, enter the email address that you want to appear as the sender of the alert.

Step 6 Next to **Relay Host**, verify the listed mail server is the one that you want to use to send the alert.

Tip To change the email server, click **Edit** ( ).

Step 7 Click **Save**.

Configuring Impact Flag Alerting

You can configure the system to alert you whenever an intrusion event with a specific impact flag occurs. Impact flags help you evaluate the impact an intrusion has on your network by correlating intrusion data, network discovery data, and vulnerability information.

You must have the Threat Smart License or Protection Classic License to configure these alerts.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 Click **Impact Flag Alerts**.

Step 3 In the **Alerts** section, choose the alert response you want to use for each alert type.

Tip To create a new alert response, choose **New** from any drop-down list.

Step 4 In the **Impact Configuration** section, check the appropriate check boxes to specify the alerts you want to receive for each impact flag.

For definitions of the impact flags, see Intrusion Event Impact Levels, on page 2263.

Step 5 Click **Save**.

Configuring Discovery Event Alerting

You can configure the system to alert you whenever a specific type of discovery event occurs.

Before you begin

- Configure your network discovery policy to log the discovery event types you want to configure alerting for as described in Configuring Network Discovery Event Logging, on page 1936.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 Click **Discovery Event Alerts**.

Step 3 In the **Alerts** section, choose the alert response you want to use for each alert type.

Tip To create a new alert response, choose **New** from any drop-down list.

Step 4 In the **Events Configuration** section, check the check boxes that correspond to the alerts you want to receive for each discovery event type.

Step 5 Click **Save**.

Configuring AMP for Networks Alerting

You can configure the system to alert you whenever any malware event, including a retrospective event, is generated by AMP for Networks (that is, a "network-based malware event" is generated.) You cannot alert on malware events generated by AMP for Endpoints ("endpoint-based malware events.")

Before you begin

- Configure a file policy to perform malware cloud lookups and associate that policy with an access control rule as described in Understanding Access Control, on page 1185.
- You must have the Malware license to configure these alerts.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 Click **Advanced Malware Protections Alerts**.

Step 3 In the **Alerts** section, choose the alert response you want to use for each alert type.

Tip To create a new alert response, choose **New** from any drop-down list.

Step 4 In the **Event Configuration** section, check the check boxes that correspond to the alerts you want to receive for each malware event type.

Keep in mind that **All network-based malware events** includes **Retrospective Events**.

(By definition, network-based malware events do not include events generated by AMP for Endpoints.)

Step 5 Click **Save**.



CHAPTER 102

External Alerting for Intrusion Events

The following topics describe how to configure external alerting for intrusion events:

- About External Alerting for Intrusion Events, on page 2085
- License Requirements for External Alerting for Intrusion Events, on page 2086
- Requirements and Prerequisites for External Alerting for Intrusion Events, on page 2086
- Configuring SNMP Alerting for Intrusion Events, on page 2086
- Configuring Syslog Alerting for Intrusion Events, on page 2088
- Configuring Email Alerting for Intrusion Events, on page 2090

About External Alerting for Intrusion Events

External intrusion event notification can help with critical-system monitoring:

- SNMP—Configured per intrusion policy and sent from managed devices. You can enable SNMP alerting per intrusion rule.
- Syslog—Configured per intrusion policy and sent from managed devices. When you enable syslog alerting in an intrusion policy, you turn it on for every rule in the policy.
- Email—Configured across all intrusion policies and sent from the Firepower Management Center. You can enable email alerts per intrusion rule, as well as limit their length and frequency.

Keep in mind that if you configured intrusion event suppression or thresholding, the system may not generate intrusion events (and thus may not send alerts) every time a rule triggers.

In a multidomain deployment, you can configure external alerting in any domain. In ancestor domains, the system generates notifications for intrusion events in descendant domains.

Note The Firepower Management Center also uses SNMP, syslog, and email *alert responses* to send different types of external alerts; see Firepower Management Center Alert Responses, on page 2075. The system does **not** use alert responses to send alerts based on individual intrusion events.

Related Topics

Intrusion Event Notification Filters in an Intrusion Policy, on page 1505

License Requirements for External Alerting for Intrusion Events

FTD License

Threat

Classic License

Protection

Requirements and Prerequisites for External Alerting for Intrusion Events

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Intrusion Admin

Configuring SNMP Alerting for Intrusion Events

After you enable external SNMP alerting in an intrusion policy, you can configure individual rules to send SNMP alerts when they trigger. These alerts are sent from the managed device.

Procedure

Step 1 In the intrusion policy editor's navigation pane, click **Advanced Settings**.

Step 2 Make sure **SNMP Alerting** is **Enabled**, then click **Edit**.
A message at the bottom of the page identifies the intrusion policy layer that contains the configuration.

Step 3 Choose an **SNMP Version**, then specify configuration options as described in Intrusion SNMP Alert Options, on page 2087.

Step 4 In the navigation pane, click **Rules**.

Step 5 In the rules pane, choose the rules where you want to set SNMP alerts, then choose **Alerting** > **Add SNMP Alert**.

Step 6 To save changes you made in this policy since the last policy commit, choose **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Intrusion SNMP Alert Options

If your network management system requires a management information base file (MIB), you can obtain it from the Firepower Management Center at `/etc/sf/DCEALERT.MIB`.

SNMP v2 Options

Option	Description
Trap Type	The trap type to use for IP addresses that appear in the alerts. If your network management system correctly renders the INET_IPV4 address type, choose **as Binary**. Otherwise, choose **as String**. For example, HP OpenView requires **as String**.
Trap Server	The server that will receive SNMP traps notification. You can specify a single IP address or hostname.
Community String	The community name.

SNMP v3 Options

Managed devices encode SNMPv3 alerts with an Engine ID value. To decode the alerts, your SNMP server requires this value, which is the hexadecimal version of the sending device's management interface IP address, appended with "01."

For example, if the device sending the SNMP alert has a management interface IP address of 172.16.1.50, the Engine ID value is 0xAC10013201.

Option	Description
Trap Type	The trap type to use for IP addresses that appear in the alerts. If your network management system correctly renders the INET_IPV4 address type, choose **as Binary**. Otherwise, choose **as String**. For example, HP OpenView requires **as String**.
Trap Server	The server that will receive SNMP traps notification. You can specify a single IP address or hostname.

Option	Description
Authentication Password	The password required for authentication. SNMP v3 uses either the Message Digest 5 (MD5) hash function or the Secure Hash Algorithm (SHA) hash function to encrypt this password, depending on configuration. If you specify an authentication password, authentication is enabled.
Private Password	The SNMP key for privacy. SNMP v3 uses the Data Encryption Standard (DES) block cipher to encrypt this password. When you enter an SNMP v3 password, the password displays in plain text during initial configuration but is saved in encrypted format. If you specify a private password, privacy is enabled, and you must also specify an authentication password.
User Name	Your SNMP user name.

Configuring Syslog Alerting for Intrusion Events

After you enable syslog alerting in an intrusion policy, the system sends all intrusion events to the syslog, either on the managed device itself or to an external host or hosts. If you specify an external host, syslog alerts are sent from the managed device.

Procedure

Step 1 In the intrusion policy editor's navigation pane, click **Advanced Settings**.

Step 2 Make sure **Syslog Alerting** is **Enabled**, then click **Edit**.
A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. The **Syslog Alerting** page is added under **Advanced Settings**.

Step 3 Enter the IP addresses of the **Logging Hosts** where you want to send syslog alerts.

If you leave this field blank, the managed device logs intrusion events using its own syslog facility.

The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Step 4 Choose **Facility** and **Priority** levels as described in Facilities and Priorities for Intrusion Syslog Alerts, on page 2089.

Step 5 To save changes you made in this policy since the last policy commit, choose **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Facilities and Priorities for Intrusion Syslog Alerts

Managed devices can send intrusion events as syslog alerts using a particular facility and **Priority**, so that the logging host can categorize the alerts. The *facility* specifies the subsystem that generated it. The *priority* specifies its severity. These facility and **Priority** values do not appear in the actual syslog messages.

Choose values that make sense based on your environment. Local configuration files (such as syslog.conf on UNIX-based logging hosts) may indicate which facilities are saved to which log files.

Syslog Alert Facilities

Facility	Description
ALERT	An alert message.
AUTH	A message associated with security and authorization.
AUTHPRIV	A restricted access message associated with security and authorization. On many systems, these messages are forwarded to a secure file.
CRON	A message generated by the clock daemon.
DAEMON	A message generated by a system daemon.
FTP	A message generated by the FTP daemon.
KERN	A message generated by the kernel. On many systems, these messages are printed to the console when they appear.
LOCAL0-LOCAL7	A message generated by an internal process.
LPR	A message generated by the printing subsystem.
MAIL	A message generated by a mail system.
NEWS	A message generated by the network news subsystem.
SYSLOG	A message generated by the syslog daemon.
USER	A message generated by a user-level process.
UUCP	A message generated by the UUCP subsystem.

Syslog Alert Priorities

Level	Description
EMERG	A panic condition broadcast to all users
ALERT	A condition that should be corrected immediately
CRIT	A critical condition
ERR	An error condition

Level	Description
WARNING	Warning messages
NOTICE	Conditions that are not error conditions, but require attention
INFO	Informational messages
DEBUG	Messages that contain debug information

Configuring Email Alerting for Intrusion Events

If you enable intrusion email alerting, the system can send email when it generates an intrusion event, regardless of which managed device or intrusion policy detected the intrusion. These alerts are sent from the Firepower Management Center.

Before you begin

- Configure your mail host to receive email alerts; see Configuring a Mail Relay Host and Notification Address, on page 906.
- Ensure that the Firepower Management Center can reverse resolve is own IP address.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 Click **Intrusion Email**.

Step 3 Choose alerting options, including the intrusion rules or rule groups for which you want to alert, as described in Intrusion Email Alert Options, on page 2090.

Step 4 Click **Save**.

Intrusion Email Alert Options

On/Off

Enables or disables intrusion email alerts.

Note Enabling it will enable alerting for all rules unless individual rules are selected.

From/To Addresses

The email sender and recipients. You can specify a comma-separated list of recipients.

Max Alerts and Frequency

The maximum number of email alerts (**Max Alerts**) that the Firepower Management Center will send per time interval (**Frequency**).

Coalesce Alerts

Reduces the number of alerts sent by grouping alerts that have the same source IP and rule ID.

Summary Output

Enables brief alerts, suitable for text-limited devices. Brief alerts contain:

- Timestamp
- Protocol
- Source and destination IPs and ports
- Message
- The number of intrusion events generated against the same source IP

For example: `2011-05-18 10:35:10 10.1.1.100 icmp 10.10.10.1:8 -> 10.2.1.3:0 snort_decoder: Unknown Datagram decoding problem! (116:108)`

If you enable **Summary Output**, also consider enabling **Coalesce Alerts**. You may also want to lower **Max Alerts** to avoid exceeding text-message limits.

Time Zone

The time zone for alert timestamps.

Email Alerting on Specific Rules Configuration

Allows you to choose the rules where you want to set email alerts.



PART XXIII

Event and Asset Analysis Tools

- Using the Context Explorer, on page 2095
- Using the Network Map, on page 2117
- Incidents, on page 2127
- Using Lookups, on page 2135



CHAPTER 103

Using the Context Explorer

The following topics describe how to use the Context Explorer in the Firepower System:

- About the Context Explorer, on page 2095
- Requirements and Prerequisites for the Context Explorer, on page 2108
- Refreshing the Context Explorer, on page 2108
- Setting the Context Explorer Time Range, on page 2109
- Minimizing and Maximizing Context Explorer Sections, on page 2109
- Drilling Down on Context Explorer Data, on page 2110
- Filters in the Context Explorer, on page 2111

About the Context Explorer

The Firepower System Context Explorer displays detailed, interactive graphical information in context about the status of your monitored network, including data on applications, application statistics, connections, geolocation, indications of compromise, intrusion events, hosts, servers, Security Intelligence, users, files (including malware files), and relevant URLs. Distinct sections present this data in the form of vivid line, bar, pie, and donut graphs, accompanied by detailed lists. The first section, a line chart of traffic and event counts over time, provides an at-a-glance picture of recent trends in your network's activity.

You can easily create and apply custom filters to fine-tune your analysis, and you can examine data sections in more detail by simply clicking or hovering your cursor over graph areas. You can also configure the explorer's time range to reflect a period as short as the last hour or as long as the last year. Only users with the Administrator, Security Analyst, or Security Analyst (Read Only) user roles have access to the Context Explorer.

The Firepower System dashboard is highly customizable and compartmentalized and updates in real time. In contrast, the Context Explorer is manually updated, designed to provide broader context for its data, and has a single, consistent layout designed for active user exploration.

You use the dashboard to monitor real-time activity on your network and appliances according to your own specific needs. Conversely, you use the Context Explorer to investigate a predefined set of recent data in granular detail and clear context: for example, if you notice that only 15% of hosts on your network use Linux, but account for almost all YouTube traffic, you can quickly apply filters to view data only for Linux hosts, only for YouTube-associated application data, or both. Unlike the compact, narrowly focused dashboard widgets, the Context Explorer sections are designed to provide striking visual representations of system activity in a format useful to both expert and casual users of the Firepower System.

The data displayed depends on such factors as how you license and deploy your managed devices, and whether you configure features that provide the data. You can also apply filters to constrain the data that appears in all Context Explorer sections.

In a multidomain deployment, the Context Explorer displays aggregated data from all subdomains when you view it in an ancestor domain. In a leaf domain, you can view data specific to that domain only.

Differences Between the Dashboard and the Context Explorer

The following table summarizes some of the key differences between the dashboard and the Context Explorer.

Table 249: Comparison: Dashboard and Context Explorer

Feature	Dashboard	Context Explorer
Displayable data	Anything monitored by the Firepower System	Applications, application statistics, geolocation, host indications of compromise, intrusion events, files (including malware files), hosts, Security Intelligence events, servers, users, and URLs
Customizability	• Selection of widgets for a dashboard is customizable • Individual widgets can be customized to varying degrees	• Cannot change base layout • Applied filters appear in explorer URL and can be bookmarked for later use
Data update frequency	Automatic (default); user-configured	Manual
Data filtering	Possible for some widgets (must edit widget preferences)	Possible for all parts of the explorer, with support for multiple filters
Graphical context	Some widgets (particularly Custom Analysis) can display data in graph form	Extensive graphical context for all data, including uniquely detailed donut graphs
Links to relevant web interface pages	In some widgets	In every section
Time range of displayed data	User-configured	User-configured

Related Topics

About Dashboards, on page 255

The Traffic and Intrusion Event Counts Time Graph

At the top of the Context Explorer is a line chart of traffic and intrusion events over time. The X-axis plots time intervals (which range from five minutes to one month, depending on the selected time window). The Y-axis plots traffic in kilobytes (blue line) and intrusion event count (red line).

Note that the smallest X-axis interval is five minutes. To accommodate this, the system will round the beginning and ending points in your selected time range down to the nearest five-minute interval.

By default, this section shows all network traffic and all generated intrusion events for the selected time range. If you apply filters, the chart changes to display only traffic and intrusion events associated with the criteria

specified in the filters. For example, filtering on the **OS Name** of `Windows` causes the time graph to display only traffic and events associated with hosts using Windows operating systems.

If you filter the Context Explorer on intrusion event data (such as a **Priority** of `High`), the blue Traffic line is hidden to allow greater focus on intrusion events alone.

You can hover your pointer over any point on the graph lines to view exact information about traffic and event counts. Hovering your pointer over one of the colored lines also brings that line to the forefront of the graph, providing clearer context.

This section draws data primarily from the Intrusion Events and Connection Events tables.

The Indications of Compromise Section

The Indications of Compromise (IOC) section of the Context Explorer contains two interactive sections that provide an overall picture of potentially compromised hosts on your monitored network: a proportional view of the most prevalent IOC types triggered, as well as a view of hosts by number of triggered indications.

For more information about IOCs, see Indications of Compromise Data, on page 2373.

The Hosts by Indication Graph

The Hosts by Indication graph, in donut form, displays a proportional view of the Indications of Compromise (IOC) triggered by hosts on your monitored network. The inner ring divides by IOC category (such as `CnC Connected` or `Malware Detected`), while the outer ring further divides that data by specific event type (such as `Impact 2 Intrusion Event — attempted-admin` or `Threat Detected in File Transfer`).

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the Hosts and Host Indications of Compromise tables.

The Indications by Host Graph

The Indications by Host graph, in bar form, displays counts of unique Indications of Compromise (IOC) triggered by the 15 most IOC-active hosts on your monitored network.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the Hosts and Host Indications of Compromise tables.

The Network Information Section

The Network Information section of the Context Explorer contains six interactive graphs that display an overall picture of connection traffic on your monitored network: sources, destinations, users, and security zones associated with traffic, a breakdown of operating systems used by hosts on the network, as well as a proportional view of access control actions your Firepower System has performed on network traffic.

The Operating Systems Graph

The Operating Systems graph, in donut form, displays a proportional representation of operating systems detected on hosts on your monitored network. The inner ring divides by OS name (such as `Windows` or `Linux`), while the outer ring further divides that data by specific operating system version (such as `Windows Server 2008` or `Linux 11.x`). Some closely related operating systems (such as Windows 2000, Windows XP, and

Windows Server 2003) are grouped together. Very scarce or unrecognized operating systems are grouped under **Other**.

Note that this graph reflects all available data regardless of date and time constraints. If you change the explorer time range, the graph does not change.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the Hosts table.

The Traffic by Source IP Graph

The Traffic by Source IP graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the top 15 most active source IP addresses on your monitored network. For each source IP address listed, blue bars represent traffic data and red bars represent connection data.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.



Note If you filter on intrusion event information, the Traffic by Source IP graph is hidden.

This graph draws data primarily from the Connection Events table.

The Traffic by Source User Graph

The Traffic by Source User graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the top 15 most active source users on your monitored network. For each source IP address listed, blue bars represent traffic data and red bars represent connection data.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.



Note If you filter on intrusion event information, the Traffic by Source User graph is hidden.

This graph draws data primarily from the Connection Events table. It displays authoritative user data.

The Connections by Access Control Action Graph

The Connections by Access Control Action graph, in pie form, displays a proportional view of access control actions (such as `Block` or `Allow`) that your Firepower System deployment has taken on monitored traffic.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.



Note If you filter on intrusion event information, the Traffic by Source User graph is hidden.

This graph draws data primarily from the Connection Events table.

The Traffic by Destination IP Graph

The Traffic by Destination IP graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the top 15 most active destination IP addresses on your monitored network. For each destination IP address listed, blue bars represent traffic data and red bars represent connection data.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Note If you filter on intrusion event information, the Traffic by Destination IP graph is hidden.

This graph draws data primarily from the Connection Events table.

The Traffic by Ingress/Egress Security Zone Graph

The Traffic by Ingress/Egress Security Zone graph, in bar form, displays counts of incoming or outgoing network traffic (in kilobytes per second) and unique connections for each security zone configured on your monitored network. You can configure this graph to display either ingress (the default) or egress security zone information, according to your needs.

For each security zone listed, blue bars represent traffic data and red bars represent connection data.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information

Tip To constrain the graph so it displays only traffic by egress security zone, hover your pointer over the graph, then click **Egress** on the toggle button that appears. Click **Ingress** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Ingress view.

Note If you filter on intrusion event information, the Traffic by Ingress/Egress Security Zone graph is hidden.

This graph draws data primarily from the Connection Events table.

The Application Information Section

The Application Information section of the Context Explorer contains three interactive graphs and one table-format list that display an overall picture of application activity on your monitored network: traffic, intrusion events, and hosts associated with applications, further organized by the estimated risk or business relevance assigned to each application. The Application Details list provides an interactive list of each application and its risk, business relevance, category, and host count.

For all instances of "application" in this section, the Application Information graph set, by default, specifically examines application protocols (such as DNS or SSH). You can also configure the Application Information section to specifically examine client applications (such as PuTTY or Firefox) or web applications (such as Facebook or Pandora).

Focusing the Application Information Section

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 Hover your pointer over the **Application Protocol Information** section.

Note If you previously changed this setting in the same Context Explorer session, the section title may appear as **Client Application Information** or **Web Application Information** instead.

Step 3 Click **Application Protocol**, **Client Application**, or **Web Application**.

The Traffic by Risk/Business Relevance and Application Graph

The Traffic by Risk/Business Relevance and Application graph, in donut form, displays a proportional representation of application traffic detected on your monitored network, arranged by the applications' estimated risk (the default) or estimated business relevance. The inner ring divides by estimated risk/business relevance level (such as `Medium` or `High`), while the outer ring further divides that data by specific application (such as `SSH` or `NetBIOS`). Scarcely detected applications are grouped under **Other**.

Note that this graph reflects all available data regardless of date and time constraints. If you change the explorer time range, the graph does not change.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Tip To constrain the graph so it displays traffic by business relevance and application, hover your pointer over the graph, then click **Business Relevance** on the toggle button that appears. Click **Risk** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Risk view.

Note If you filter on intrusion event information, the Traffic by Risk/Business and Application graph is hidden.

This graph draws data primarily from the Connection Events and Application Statistics tables.

The Intrusion Events by Risk/Business Relevance and Application Graph

The Intrusion Events by Risk/Business Relevance and Application graph, in donut form, displays a proportional representation of intrusion events detected on your monitored network and the applications associated with those events, arranged by the applications' estimated risk (the default) or estimated business relevance. The inner ring divides by estimated risk/business relevance level (such as `Medium` or `High`), while the outer ring further divides that data by specific application (such as `SSH` or `NetBIOS`). Scarcely detected applications are grouped under **Other**.

Hover your pointer over any part of the donut graph to view more detailed information. Click any part of the graph to filter or drill down on that information, or (where applicable) to view application information.



Tip To constrain the graph so it displays intrusion events by business relevance and application, hover your pointer over the graph, then click **Business Relevance** on the toggle button that appears. Click **Risk** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Risk view.

This graph draws data primarily from the Intrusion Events and Application Statistics tables.

The Hosts by Risk/Business Relevance and Application Graph

The Hosts by Risk/Business Relevance and Application graph, in donut form, displays a proportional representation of hosts detected on your monitored network and the applications associated with those hosts, arranged by the applications' estimated risk (the default) or estimated business relevance. The inner ring divides by estimated risk/business relevance level (such as `Medium` or `High`), while the outer ring further divides that data by specific application (such as `SSH` or `NetBIOS`). Very scarce applications are grouped under **Other**.

Hover your pointer over any part of the donut graph to view more detailed information. Click any part of the graph to filter or drill down on that information.



Tip To constrain the graph so it displays hosts by business relevance and application, hover your pointer over the graph, then click **Business Relevance** on the toggle button that appears. Click **Risk** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Risk view.

This graph draws data primarily from the Applications table.

The Application Details List

At the bottom of the Application Information section is the Application Details List, a table that provides estimated risk, estimated business relevance, category, and hosts count information for each application detected on your monitored network. The applications are listed in descending order of associated host count.

The Application Details List table is not sortable, but you can click on any table entry to filter or drill down on that information, or (where applicable) to view application information. This table draws data primarily from the Applications table.

Note that this list reflects all available data regardless of date and time constraints. If you change the explorer time range, the list does not change.

The Security Intelligence Section

The Security Intelligence section of the Context Explorer contains three interactive bar graphs that display an overall picture of traffic on your monitored network that is blocked or monitored by Security Intelligence. The graphs sort such traffic by category, source IP address, and destination IP address, respectively; both the amount of traffic (in kilobytes per second) and the number of applicable connections appear.

The Security Intelligence Traffic by Category Graph

The Security Intelligence Traffic by Category graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the top Security Intelligence categories of traffic on your monitored network. For each category listed, blue bars represent traffic data and red bars represent connection data.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to drill down on that information.



Note If you filter on intrusion event information, the Security Intelligence Traffic by Category graph is hidden.

This graph draws data primarily from the Security Intelligence Events table.

The Security Intelligence Traffic by Source IP Graph

The Security Intelligence Traffic by Source IP graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the top source IP addresses of Security Intelligence-monitored traffic on your monitored network. For each category listed, blue bars represent traffic data and red bars represent connection data.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to drill down on that information.



Note If you filter on intrusion event information, the Security Intelligence Traffic by Source IP graph is hidden.

This graph draws data primarily from the Security Intelligence Events table.

The Security Intelligence Traffic by Destination IP Graph

The Security Intelligence Traffic by Destination IP graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the top destination IP addresses of Security Intelligence-monitored traffic on your monitored network. For each category listed, blue bars represent traffic data and red bars represent connection data.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to drill down on that information.



Note If you filter on intrusion event information, the Security Intelligence Traffic by Destination IP graph is hidden.

This graph draws data primarily from the Security Intelligence Events table.

The Intrusion Information Section

The Intrusion Information section of the Context Explorer contains six interactive graphs and one table-format list that display an overall picture of intrusion events on your monitored network: impact levels, attack sources, target destinations, users, priority levels, and security zones associated with intrusion events, as well as a detailed list of intrusion event classifications, priorities, and counts.

The Intrusion Events by Impact Graph

The Intrusion Events by Impact graph, in pie form, displays a proportional view of intrusion events on your monitored network, grouped by estimated impact level (from 0 to 4).

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the intrusion detection (IDS Statistics) and Intrusion Events tables.

The Top Attackers Graph

The Top Attackers graph, in bar form, displays counts of intrusion events for the top attacking host IP addresses (causing those events) on your monitored network.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the Intrusion Events table.

The Top Users Graph

The Top Users graph, in bar form, displays users on your monitored network that are associated with the highest intrusion event counts, by event count.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the intrusion detection (IDS) User Statistics and Intrusion Events tables. It displays authoritative user data.

The Intrusion Events by Priority Graph

The Intrusion Events by Priority graph, in pie form, displays a proportional view of intrusion events on your monitored network, grouped by estimated priority level (such as `High`, `Medium`, or `Low`).

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the Intrusion Events table.

The Top Targets Graph

The Top Targets graph, in bar form, displays counts of intrusion events for the top target host IP addresses (targeted in the connections causing those events) on your monitored network.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the Intrusion Events table.

The Top Ingress/Egress Security Zones Graph

The Top Ingress/Egress Security Zones graph, in bar form, displays counts of intrusion events associated with each security zone (ingress or egress, depending on graph settings) configured on your monitored network.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Tip To constrain the graph so it displays only traffic by egress security zone, hover your pointer over the graph, then click **Egress** on the toggle button that appears. Click **Ingress** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Ingress view.

This graph draws data primarily from the Intrusion Events table.

You can configure this graph to display either ingress (the default) or egress security zone information, according to your needs.

The Intrusion Event Details List

At the bottom of the Intrusion Information section is the Intrusion Event Details List, a table that provides classification, estimated priority, and event count information for each intrusion event detected on your monitored network. The events are listed in descending order of event count.

The Intrusion Event Details List table is not sortable, but you can click on any table entry to filter or drill down on that information. This table draws data primarily from the Intrusion Events table.

The Files Information Section

The Files Information section of the Context Explorer contains six interactive graphs that display an overall picture of file and malware events on your monitored network.

Five of the graphs display data related to AMP for Networks (formerly called AMP for Firepower): the file types, file names, and malware dispositions of the files detected in network traffic, as well as the hosts sending (uploading) and receiving (downloading) those files. The final graph displays all malware threats detected in your organization, whether by AMP for Networks or AMP for Endpoints.

Note If you filter on intrusion information, the entire Files Information Section is hidden.

The Top File Types Graph

The Top File Types graph, in donut form, displays a proportional view of the file types detected in network traffic (outer ring), grouped by file category (inner ring).

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Note that you must have a Malware license to for this graph to display AMP for Networks data.

This graph draws data primarily from the File Events table.

The Top File Names Graph

The Top File Names graph, in bar form, displays counts of the top unique file names detected in network traffic.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Note that you must have a Malware license to for this graph to display AMP for Networks data.

This graph draws data primarily from the File Events table.

The Files by Disposition Graph

The Top File Types graph, in pie form, displays a proportional view of the malware dispositions for files detected by the AMP for Networks feature (formerly called AMP for Firepower). Note that only files for which the Firepower Management Center performed a malware cloud lookup have dispositions. Files that did not trigger a cloud lookup have a disposition of `N/A`. The disposition `Unavailable` indicates that the Firepower Management Center could not perform a malware cloud lookup.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Note that you must have a Malware license to for this graph to display AMP for Networks data.

This graph draws data primarily from the File Events table.

The Top Hosts Sending Files Graph

The Top Hosts Sending Files graph, in bar form, displays counts of the number of files detected in network traffic for the top file-sending host IP addresses.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.



Tip To constrain the graph so it displays only hosts sending malware, hover your pointer over the graph, then click **Malware** on the toggle button that appears. Click **Files** to return to the default files view. Note that navigating away from the Context Explorer also returns the graph to the default files view.

Note that you must have a Malware license to for this graph to display AMP for Networks data.

This graph draws data primarily from the File Events table.

The Top Hosts Receiving Files Graph

The Top Hosts Receiving Files graph, in bar form, displays counts of the number of files detected in network traffic for the top file-receiving host IP addresses.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.



Tip To constrain the graph so it displays only hosts receiving malware, hover your pointer over the graph, then click **Malware** on the toggle button that appears. Click **Files** to return to the default files view. Note that navigating away from the Context Explorer also returns the graph to the default files view.

Note that you must have a Malware license to for this graph to display AMP for Networks data.

This graph draws data primarily from the File Events table.

The Top Malware Detections Graph

The Top Malware Detections graph, in bar form, displays counts of the top malware threats detected in your organization, whether by AMP for Networks or AMP for Endpoints.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Note that you must have a Malware license to for this graph to display AMP for Networks data.

This graph draws data primarily from the File Events and Malware Events tables.

The Geolocation Information Section

The Geolocation Information section of the Context Explorer contains three interactive donut graphs that display an overall picture of countries with which hosts on your monitored network are exchanging data: unique connections by initiator or responder country, intrusion events by source or destination country, and file events by sending or receiving country.

The Connections by Initiator/Responder Country Graph

The Connections by Initiator/Responder Country graph, in donut form, displays a proportional view of the countries involved in connections on your network as either the initiator (the default) or the responder. The inner ring groups these countries together by continent.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.



Tip To constrain the graph so it displays only countries acting as the responder in connections, hover your pointer over the graph, then click **Responder** on the toggle button that appears. Click **Initiator** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Initiator view.

This graph draws data primarily from the Connection Summary Data table.

The Intrusion Events by Source/Destination Country Graph

The Intrusion Events by Source/Destination Country graph, in donut form, displays a proportional view of the countries involved in intrusion events on your network as either the source of the event (the default) or the destination. The inner ring groups these countries together by continent.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.



Tip To constrain the graph so it displays only countries acting as the destinations of intrusion events, hover your pointer over the graph, then click **Destination** on the toggle button that appears. Click **Source** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Source view.

This graph draws data primarily from the Intrusion Events table.

The File Events by Sending/Receiving Country Graph

The File Events by Sending/Receiving Country graph, in donut form, displays a proportional view of the countries detected in file events on your network as either sending (the default) or receiving files. The inner ring groups these countries together by continent.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Tip To constrain the graph so it displays only countries receiving files, hover your pointer over the graph, then click **Receiver** on the toggle button that appears. Click **Sender** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Sender view.

This graph draws data primarily from the File Events table.

The URL Information Section

The URL Information section of the Context Explorer contains three interactive bar graphs that display an overall picture of URLs with which hosts on your monitored network are exchanging data: traffic and unique connections associated with URLs, sorted by individual URL, URL category, and URL reputation. You cannot filter on URL information.

Note If you filter on intrusion event information, the entire URL Information Section is hidden.

Note that you must have a URL Filtering license for this graph to include URL category and reputation data.

The Traffic by URL Graph

The Traffic by URL graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the top 15 most requested URLs on your monitored network. For each URL listed, blue bars represent traffic data and red bars represent connection data.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to drill down on that information.

Note If you filter on intrusion event information, the Traffic by URL graph is hidden.

Note that you must have a URL Filtering license for this graph to include URL category and reputation data.

This graph draws data primarily from the Connection Events table.

The Traffic by URL Category Graph

The Traffic by URL Category graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the most requested URL categories (such as `Search Engines` or `Streaming Media`) on your monitored network. For each URL category listed, blue bars represent traffic data and red bars represent connection data.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to drill down on that information.

Note If you filter on intrusion event information, the Traffic by URL Category graph is hidden.

Note that you must have a URL Filtering license for this graph to include URL category and reputation data.

This graph draws data primarily from the URL Statistics and Connection Events tables.

The Traffic by URL Reputation Graph

The Traffic by URL Reputation graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the most requested URL reputation groups (such as Well Known or Benign sites with security risks) on your monitored network. For each URL reputation listed, blue bars represent traffic data and red bars represent connection data.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to drill down on that information.

Note If you filter on intrusion event information, the Traffic by URL Reputation graph is hidden.

Note that you must have a URL Filtering license for this graph to include URL category and reputation data.

This graph draws data primarily from the URL Statistics and Connection Events tables.

Requirements and Prerequisites for the Context Explorer

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Security Analyst

Refreshing the Context Explorer

The Context Explorer does not automatically update the information it displays. To incorporate new data, you must manually refresh the explorer.

Note that, although reloading the Context Explorer itself (by refreshing the browser program or navigating away from, then back to, the Context Explorer) refreshes all displayed information, this does not preserve any changes you made to section configuration (such as the Ingress/Egress graphs and the Application Information section) and may cause delays in loading.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 Click **Reload** at the upper right.

Reload is dimmed until your refresh is finished.

Setting the Context Explorer Time Range

You can configure the Context Explorer time range to reflect a period as short as the last hour (the default) or as long as the last year. Note that when you change the time range, the Context Explorer does not automatically update to reflect the change. To apply the new time range, you must manually refresh the explorer.

Changes to the time range persist even if you navigate away from the Context Explorer or end your login session.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 From the **Show the last** drop-down list, choose a time range.

Step 3 Optionally, to view data from the new time range, click **Reload**.

Tip Clicking **Apply Filters** also applies any time range updates.

Minimizing and Maximizing Context Explorer Sections

You can minimize and hide one or more sections of the Context Explorer. This is useful if you want to focus on only certain sections, or if you want a simpler view. You cannot minimize the Traffic and Intrusion Event Counts Time Graph.

Context Explorer sections retain the minimized or maximized states that you configure even if you refresh the page or log out of the appliance.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 To minimize a section, click **Minimize** ( ) in a section's title bar.

Step 3 To maximize a section, click maximize **Maximize** ( ) in a minimized section's title bar.

Drilling Down on Context Explorer Data

If you want to examine graph or list data in more detail than the Context Explorer allows, you can drill down to the table views of the relevant data. (Note that you cannot drill down on the Traffic and Intrusion Events over Time graph.) For example, drilling down on an IP address in the Traffic by Source IP graph displays the Connections with Application Details view of the Connection Events table, including only data associated with the source IP address you selected.

Depending on the type of data you examine, additional options can appear in the context menu. Data points that are associated with specific IP addresses offer the option to view host or whois information on the IP address you select. Data points associated with specific applications offer the option to view application information on the application you select. Data points associated with a specific user offer the option to view that user's user profile page. Data points associated with an intrusion event message offer the option to view the rule documentation for that event's associated intrusion rule, and data points associated with a specific IP address offer the option to add that address to a Block or Do Not Block list. For more information about these lists, see Global and Domain Security Intelligence Lists, on page 428 and subtopics.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 In any section but **Traffic and Intrusion Events over Time**, click a data point that you want to investigate.

Step 3 Depending on the data point you selected, you have several options:

- To view more details of this data in a table view, choose **Drill into Analysis**.
- If you chose a data point associated with a specific IP address and want more information about the associated host, choose **View Host Information**.
- If you chose a data point with a specific IP address and want to make a whois search on that address, choose **Whois**.
- If you chose a data point associated with a specific application and want more information about that application, choose **View Application Information**.
- If you chose a data point associated with a specific user and want more information about that user, choose **View User Information**.

- If you chose a data point associated with a specific intrusion event message and want more information about the associated intrusion rule, choose **View Rule Documentation**
- If you chose a data point associated with a specific IP address and want to add that IP address to the Security Intelligence global Block or Do Not Block list, choose the appropriate option.

Filters in the Context Explorer

Beyond the basic, wide-ranging data that the Context Explorer initially displays, you have the option to filter that data for a more granular contextual picture of activity on your network. Filters encompass all types of Firepower System data except URL information, support exclusion as well as inclusion, can be applied quickly by clicking on Context Explorer graph data points, and affect the entire explorer. You can apply up to 20 filters at a time.

You can add filters to Context Explorer data in several ways:

- from the Add Filter dialog
- from the context menu, when you select a data point in the explorer
- from the text links that appear in certain detail view pages (Application Detail, Host Profile, Rule Detail, and User Profile). Clicking these links automatically opens and filters the Context Explorer according to the relevant data on the detail view page. For example, clicking the Context Explorer link on a user detail page for the user `jenkins` constrains the explorer to show only data associated with that user

Some filter types are incompatible with others: for example, filters that relate to intrusion events (such as **Device** and **Inline Result**) cannot be applied at the same time as connection event-related filters (such as **Access Control Action**) because the system cannot sort connection event data by intrusion event data. The system automatically prevents incompatible filters from simultaneously applying; when one filter type is more recently activated, filters of the incompatible type are hidden as long as the incompatibility exists.

When multiple filters are active, values for the same data type are treated as OR search criteria: all data that matches at least one of the values appears. Values for different data types are treated as AND search criteria: to appear, data must match at least one value for each filtered data type. For example, data that appears for the filter set of `Application: 2channel`, `Application: Reddit`, and `User: edickinson` must be associated with the user `edickinson` **AND** either the application `2channel` **OR** the application `Reddit`.

In a multidomain deployment, you can filter by multiple descendant domains when viewing the Context Explorer in an ancestor domain. In such cases, use caution when also adding **IP Address** filters. The system builds a separate network map for each leaf domain. Using literal IP addresses to constrain this configuration can have unexpected results.

Note that the data displayed depends on such factors as how you license and deploy your managed devices and whether you configure features that provide the data.



Note Filters function as a simple, agile tool to get the precise Firepower data context you need at any given time. They are not intended as permanent configuration settings, and disappear when you navigate away from the Context Explorer or end your session. To preserve filter settings for later use, see Saving Filtered Context Explorer Views, on page 2115.

Data Type Field Options

The following table lists the data types available as filters, with examples and brief definitions of each.

Table 250: Filter Data Types

Type	Example Values	Definition
Access Control Action	`Allow`, `Block`	Action taken by your access control policy to allow or block traffic.
Application Category	`web browser`, `email`	General classification of an application's most essential function.
Application Name	`Facebook`, `HTTP`	Name of an application.
Application Risk	`Very High`, `Medium`	Estimated security risk of an application.
Application Tag	`encrypts communications`, `sends mail`	Additional information about an application; applications can have any number of tags, including none.
Application Type	`Client`, `Web Application`	Type of an application: application protocol, client, or web application.
Business Relevance	`Very Low`, `High`	Estimated relevance of an application to business activity (as opposed to recreation).
Continent	`North America`, `Asia`	Continent associated with a routable IP address detected on your monitored network.
Country	`Canada`, `Japan`	Country associated with a routable IP address detected on your monitored network.
Device	`device1.example.com`, `192.168.1.3`	Name or IP address of a device on your monitored network.
Domain	`Asia Division`, `Europe Division`	The domain of the device whose network activity you want to graph. This data type is only present in a multidomain deployment.
Event Classification	`Potential Corporate Policy Violation`, `Attempted Denial of Service`	Capsule description of an intrusion event, determined by the classification of the rule, decoder, or preprocessor that triggered it.
Event Message	`dns response`, `P2P`	Message generated by an event, determined by the rule, decoder, or preprocessor that triggered it.
File Disposition	`Malware`, `Clean`	Disposition of a file for which the Firepower Management Center performed a malware cloud lookup.
File Name	`Packages.bz2`	Name of a file detected in network traffic.
File SHA256	any 32-bit string	SHA-256 hash value of a file for which the Firepower Management Center performed a malware cloud lookup.
File Type	`GZ`, `SWF`, `MOV`	File type detected in network traffic.

Type	Example Values	Definition
File Type Category	`Archive`, `Multimedia`, `Executables`	General category of file type detected in network traffic.
IP Address	`192.168.1.3`, `2001:0db8:85a3::0000/24`	IPv4 or IPv6 addresses, address ranges, or address blocks. Note that searching for an IP address returns events where that address was either the source or the destination for the event.
Impact Level	`Impact Level 1`, `Impact Level 2`	Estimated impact of an event on your monitored network.
Inline Result	`dropped`, `would have dropped`	Whether traffic was dropped, would have been dropped, or was not acted upon by the system.
IOC Category	`High Impact Attack`, `Malware Detected`	Category for a triggered Indication of Compromise (IOC) event.
IOC Event Type	`exploit-kit`, `malware-backdoor`	Identifier associated with a specific Indication of Compromise (IOC), referring to the event that triggers it.
Malware Threat Name	`W32.Trojan.a6b1`	The name of a malware threat.
OS Name	`Windows`, `Linux`	Name of an operating system.
OS Version	`XP`, `2.6`	Specific version of an operating system.
Priority	`high`, `low`	Estimated urgency of an event.
Security Intelligence Category	`Malware`, `Spam`	Category of risky traffic, as determined by Security Intelligence.
Security Zone	`My Security Zone`, `Security Zone X`	A set of interfaces through which traffic is analyzed and, in an inline deployment, passes.
SSL	`yes`, `no`	SSL- or TLS-encrypted traffic.
User	`wsmith`, `mtwain`	Identity of a user logged in to a host on your monitored network.

Creating a Filter from the Add Filter Window

Use this procedure to create filters from scratch with the Add Filter window. (You can also use the context menu to create quick filters.)

The Add Filter window, which you access by clicking **Plus** (+) under **Filters** at the top left of the Context Explorer, contains only two fields:

- The **Data Type** drop-down list contains many different types of Firepower System data you can use to constrain the Context Explorer. After you select a data type, you then enter a specific value for that type in the **Filter** field (for example, a value of `Asia` for the type **Continent**). To assist you, the Filter field presents several grayed-out example values for the data type you select. (These are erased when you enter data in the field.)

- In the **Filter** field, you can input special search parameters such as * and ! essentially as you can in event searches. You can create exclusionary filters by prefixing filter parameters with the ! symbol.

Note Filters that you add are not automatically applied; you must click **Apply Filters** to see the filtering in the Context Explorer.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 Under **Filters** at the top left, click **Plus** ( ).

Step 3 From the **Data Type** drop-down list, choose the data type you want to filter on.

Step 4 In the **Filter** field, enter the data type value you want to filter on.

Step 5 Click **OK**.

Step 6 Optionally, repeat the previous steps to add more filters until you have the filter set you need.

Step 7 Click **Apply Filters**.

Related Topics

Data Type Field Options, on page 2112
Search Constraints, on page 2177

Creating a Quick Filter from the Context Menu

While exploring Context Explorer graph and list data, you can click on data points, then use the context menu to quickly create a filter based on that data, either inclusive or exclusive. If you use the context menu to filter on information of data type Application, User, or Intrusion Event Message, or any individual host, the filter widget includes a widget information that links to the relevant detail page for that data type (such as Application Detail for application data). Note that you cannot filter on URL data.

You can also use the context menu to investigate specific graph or list data in more detail.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 In any explorer section except Traffic and Intrusion Events over Time or sections that contain URL data, click a data point you want to filter on.

Step 3 You have two options:

- To add a filter for this data, click **Add Filter**.

- To add an exclusion filter for this data, click **Add Exclude Filter**. The filter, when applied, displays all data **not** associated with the excluded value. Exclude filters display an exclamation point (!) before the filter value.

Saving Filtered Context Explorer Views

To preserve filter settings in the Context Explorer after you navigate away from the Context Explorer or end your session, create a browser bookmark of the Context Explorer with your preferred filters applied. Because applied filters are incorporated in the Context Explorer page URL, loading a bookmark of that page also loads the corresponding filters.

Procedure

Create a browser bookmark of the Context Explorer with your preferred filters applied.

Viewing Filter Data

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 Click **Information** on any eligible filter widget.

Deleting a Filter

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 Under **Filters** at the top left, click **Clear** (×) on any filter widget.

Tip If you want to delete all filters at once, you can click **Clear**.



CHAPTER 104

Using the Network Map

The following topics describe how to use the network map:

- Requirements and Prerequisites for the Network Map, on page 2117
- The Network Map, on page 2117
- Custom Network Topologies, on page 2123

Requirements and Prerequisites for the Network Map

Model Support

Any.

Supported Domains

Leaf

User Roles

- Admin
- Discovery Admin

The Network Map

The Firepower System monitors traffic traveling over your network, decodes the traffic data, and then compares the data to established operating systems and fingerprints. The system then uses this data to build a detailed representation of your network, called a *network map*. In multidomain deployments, the system creates an individual network map for each leaf domain.

The system gathers data from the managed devices identified for monitoring in the network discovery policy. The managed devices detect network assets directly from monitored traffic and indirectly from processed NetFlow records. If multiple devices detect the same network asset, the system combines the information into a composite representation of the asset.

To augment data from passive detection, you can:

- Actively scan hosts using the open-source scanner, Nmap™, and add the scan results to your network map.
- Manually add host data from a third-party application using the host input feature.

The network map displays your network topology in terms of detected hosts and network devices.

You can use the network map to:

- Obtain a quick, overall view of your network.
- Select different views to suit the analysis you want to perform. Each view of the network map has the same format: a hierarchical tree with expandable categories and sub-categories. When you click a category, it expands to show you the sub-categories beneath it.
- Organize and identify subnets via the custom topology feature. For example, if each department in your organization uses a different subnet, you can assign familiar labels to those subnets using the custom topology feature.
- View detailed information by drilling down to any monitored host's *host profile*.
- Delete an asset if you are no longer interested in investigating it.

Note If the system detects activity associated with a host you deleted from a network map, it re-adds the host to the network map. Similarly, deleted applications are re-added to the network map if the system detects a change in the application (for example, if an Apache web server is upgraded to a new version). Vulnerabilities are reactivated on specific hosts if the system detects a change that makes the host vulnerable.

Tip If you want to permanently exclude a host or subnet from the network map, modify the network discovery policy. You may wish to exclude load balancers and NAT devices from monitoring if you find that they are generating excessive or irrelevant events.

Related Topics

Configuring the Network Discovery Policy, on page 1919

The Hosts Network Map

The network map on the Hosts tab displays a host count and a list of host IP addresses and primary MAC addresses. Each address or partial address is a link to the next level. This network map view provides a count of all unique hosts detected by the system, regardless of whether the hosts have one IP address or multiple IP addresses.

Use the hosts network map to view the hosts on your network, organized by subnet in a hierarchical tree, as well as to drill down to the host profiles for specific hosts.

The system can add hosts to the network map from exported NetFlow records, but the available information for these hosts is limited; see Differences between NetFlow and Managed Device Data, on page 1824.

By creating a custom topology for your network, you can assign meaningful labels to your subnets, such as department names, that appear in the hosts network map. You can also view the hosts network map according to the organization you specified in the custom topology.

You can delete entire networks, subnets, or individual hosts from the hosts network map. For example, if you know that a host is no longer attached to your network, you can delete it to simplify your analysis. If the system afterwards detects activity associated with the deleted host, it re-adds the host to the network map. If you want to permanently exclude a host or subnet from the network map, modify the network discovery policy.

Caution Do not delete network devices from the network map. The system uses them to determine network topology.

On the hosts network map page, you can search only for primary MAC addresses, and the Hosts [MAC] counter includes only primary MAC addresses. For descriptions of primary and secondary MAC addresses, see Basic Host Information in the Host Profile, on page 2327.

The Network Devices Network Map

The network map on the Network Devices tab displays the network devices (bridges, routers, NAT devices, and load balancers) that connect one segment of your network to another. The map contains two sections listing devices identified by an IP address and devices identified by a MAC address.

The map also provides a count of all unique network devices detected by the system, regardless of whether the devices have one IP address or multiple IP addresses.

If you create a custom topology for your network, the labels you assign to your subnets appear in the network devices network map.

The methods the system uses to distinguish network devices include:

- the analysis of Cisco Discovery Protocol (CDP) messages, which can identify network devices and their types (Cisco devices only)
- the detection of the Spanning Tree Protocol (STP), which identifies a device as a switch or bridge
- the detection of multiple hosts using the same MAC address, which identifies the MAC address as belonging to a router
- the detection of TTL value changes from the client side, or TTL values that change more frequently than a typical boot time, which identify NAT devices and load balancers

If a network device communicates using CDP, it may have one or more IP addresses. If it communicates using STP, it may only have a MAC address.

You cannot delete network devices from the network map, because the system uses their locations to determine network topology.

The host profile for a network device has a Systems section rather than an Operating Systems section, which includes a Hardware column that reflects the hardware platform for any mobile devices detected behind the network device. If a value for a hardware platform is listed under Systems, that system represents a mobile device or devices detected behind the network device. Note that mobile devices may or may not have hardware platform information, but hardware platform information is never detected for systems that are not mobile devices.

The Mobile Devices Network Map

The network map on the Mobile Devices tab displays mobile devices attached to your network. This network map also provides a count of all unique mobile devices detected by the system, regardless of whether the devices have one IP address or multiple IP addresses.

Each address or partial address is a link to the next level. You can also delete a subnet or IP address; if the system rediscovers the device, it re-adds the device to the network map.

You can also drill down to view the host profiles for the mobile devices.

To identify mobile devices, the system:

- analyzes User-Agent strings in HTTP traffic from the mobile device's mobile browser
- monitors the HTTP traffic of specific mobile applications

If you create a custom topology for your network, the labels you assign to your subnets appear in the mobile devices network map.

The Indications of Compromise Network Map

The network map on the Indications of Compromise tab displays the compromised hosts on your network, organized by IOC category. Affected hosts are listed beneath each category. Each address or partial address is a link to the next level.

From the indications of compromise network map, you can view the host profile of each host determined to have been compromised in a specific way. You can also delete (mark as resolved) any IOC category or any specific host, which removes the IOC tag from the relevant hosts. For example, you can delete an IOC category from the network map if you have determined that the issue is addressed and unlikely to recur.

Marking a host or IOC category resolved from the network map does not remove it from your network. A resolved host or IOC category reappears in the network map if your system newly detects information that triggers that IOC.

For more information about how the system determines indications of compromise, see Indications of Compromise Data, on page 2373 and subtopics.

The Application Protocols Network Map

The network map on the Application Protocols tab displays the applications running on your network, organized in a hierarchical tree by application name, vendor, version, and finally by the hosts running each application.

The applications that the system detects may change with system software and VDB updates, and if you import any add-on detectors. The release notes or advisory text for each system or VDB update contains information on any new and updated detectors. For a comprehensive up-to-date list of detectors, see the Cisco Support Site (http://www.cisco.com/cisco/web/support/index.html).

From this network map, you can view the host profile of each host that runs a specific application.

You can also delete any application category, any application running on all hosts, or any application running on a specific host. For example, you can delete an application from the network map if you know it is disabled on the host and you want to make sure the system does not use it for impact level qualification.

Deleting an application from the network map does not remove it from your network. A deleted application reappears in the network map if your system detects a change in the application (for example, if an Apache web server is upgraded to a new version) or if you restart your system's discovery function.

Depending on what you delete, the behavior differs:

- Application Category — Deleting removes the application category from the network map. All applications that reside beneath the category are removed from any host profile that contains the applications.

  For example, if you delete **http**, all applications identified as **http** are removed from all host profiles and **http** no longer appears in the applications view of the network map.

- Specific Application, Vendor, or Version — Deleting removes the affected application from the network map and from any host profiles that contain it.

  For example, if you expand the **http** category and delete **Apache**, all applications listed as Apache with any version listed beneath Apache are removed from any host profiles that contain them. Similarly, if instead of deleting **Apache**, you delete a specific version (**1.3.17**, for example), only the version you selected will be deleted from affected host profiles.

- Specific IP Address — Deleting the IP address removes it from the application list and removes the application itself from the host profile of the IP address you selected.

  For example, if you expand **http**, **Apache**, **1.3.17 (Win32)**, and then delete **172.16.1.50:80/tcp**, the Apache 1.3.17 (Win32) application is deleted from the host profile of IP address 172.16.1.50.

The Vulnerabilities Network Map

The network map on the Vulnerabilities tab displays vulnerabilities that the system has detected on your network, organized by legacy vulnerability ID (SVID), Bugtraq ID, CVE ID, or Snort ID.

From this network map, you can view the details of specific vulnerabilities, as well as the host profile of any host subject to a specific vulnerability. This information can help you evaluate the threat posed by that vulnerability to specific affected hosts.

If you determine that a specific vulnerability is not applicable to the hosts on your network (for example, you have applied a patch), you can deactivate the vulnerability. Deactivated vulnerabilities still appear on the network map, but the IP addresses of their previously affected hosts appear in gray italics. The host profiles for those hosts show deactivated vulnerabilities as invalid, though you can manually mark them as valid for individual hosts.

If there is an identity conflict for an application or operating system on a host, the system lists the vulnerabilities for both potential identities. When the identity conflict is resolved, the vulnerabilities remain associated with the current identity.

By default, the network map displays the vulnerabilities of a detected application only if the packet contains the application's vendor and version. However, you can configure the system to list the vulnerabilities for applications lacking vendor and version data by enabling the vulnerability mapping setting for the application in the Firepower Management Center configuration.

The numbers next to a vulnerability ID (or range of vulnerability IDs) represent two counts:

Affected Hosts

The first number is a count of non-unique hosts that are affected by a vulnerability or vulnerabilities. If a host is affected by more than one vulnerability, it is counted multiple times. Therefore, it is possible for the count to be higher than the number of hosts on your network. Deactivating a vulnerability

decrements this count by the number of hosts that are potentially affected by the vulnerability. If you have not deactivated any vulnerabilities for any of the potentially affected hosts for a vulnerability or range of vulnerabilities, this count is not displayed.

Potentially Affected Hosts

The second number is a count of the total number of non-unique hosts that the system has determined are *potentially* affected by a vulnerability or vulnerabilities.

Deactivating a vulnerability renders it inactive only for the hosts you designate. You can deactivate a vulnerability for all hosts that have been judged vulnerable or for a specified individual vulnerable host. After a vulnerability is deactivated, the applicable hosts' IP addresses appear in gray italics in the network map. In addition, host profiles for those hosts show deactivated vulnerabilities as invalid.

If the system subsequently detects the vulnerability on a host where it has not been deactivated (for example, on a new host in the network map), the system activates the vulnerability for that host. You have to explicitly deactivate the newly discovered vulnerability. Also, if the system detects an operating system or application change for a host, it may reactivate associated deactivated vulnerabilities.

The Host Attributes Network Map

The network map on the Host Attributes tab displays the hosts on your network organized by either user-defined or compliance white list host attributes. You cannot organize hosts using predefined host attributes in this display.

When you choose the host attribute you want to use to organize your hosts, the Firepower Management Center lists the possible values for that attribute in the network map and groups hosts based on their assigned values. For example, if you choose to organize your hosts by white list host attributes, the system displays them in categories of Compliant, Non-Compliant, and Not Evaluated.

You can also view the host profile of any host assigned a specific host attribute value.

Related Topics

Host Attributes in the Host Profile, on page 2341

Viewing Network Maps

You must be an Admin or Security Analyst user to view the network map.

Procedure

Step 1 Choose **Analysis** > **Hosts** > **Network Map**.

Step 2 Click the network map you want to view.

Step 3 Continue as appropriate:

- Choose Domain — In multidomain environments, choose a leaf domain from the **Domain** drop-down list.
- Filter Hosts — If you want to filter by IP or MAC addresses, enter an address into the search field. To clear the search, click **Clear** ( ).
- Drill Down — If you want to investigate a category or host profile, drill down through the categories or subnets in the map. If you have defined a custom topology, click **(topology)** from **Hosts** to view it, then click on **(hosts)** if you want to toggle back to the default view.

- Delete — Click **Delete** ( ) next to the appropriate element to:
  - Remove an element from the map on **Hosts**, **Network Devices**, **Mobile Devices**, or **Application Protocols**.
  - Mark an IOC category, compromised host, or group of compromised hosts resolved on **Indications of Compromise**.
  - Deactivate a vulnerability for all hosts or a single host on **Vulnerabilities**.
- Specify Vulnerabilities Class — On **Vulnerabilities**, choose the class of vulnerabilities you want to view from the **Type** drop-down list.
- Specify Organizing Attribute — On **Host Attributes**, choose an attribute from the **Attribute** drop-down list.

Related Topics

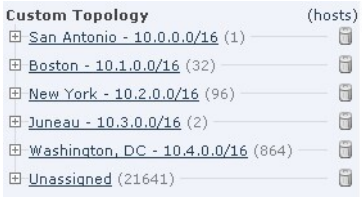
Custom Network Topologies, on page 2123
Host Profiles, on page 2326

Custom Network Topologies

Use the custom topology feature to help you organize and identify subnets in your hosts and network devices network maps.

For example, if each department within your organization uses a different subnet, you can label those subnets using the custom topology feature.

You can also view the hosts network map according to the organization you specified in the custom topology.



You can specify a custom topology's networks using any or all of the following strategies:

- You can import networks from the network discovery policy to add the networks that you configured the system to monitor.
- You can add networks to your topology manually.

The Custom Topology page lists your custom topologies and their status. If the light bulb icon next to the policy name is lit, the topology is active and affects your network map. If it is dimmed, the topology is inactive.

Related Topics

The Hosts Network Map, on page 2118
The Network Devices Network Map, on page 2119

Creating Custom Topologies

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Custom Topology** in the toolbar.

Step 3 Click **Create Topology**.

Step 4 Enter a **Name**.

Step 5 Optionally, enter a **Description**.

Step 6 Add networks to your topology. You can use any or all of the following strategies:

- Import networks from a network discovery policy as described in Importing Networks from the Network Discovery Policy, on page 2124.
- Manually add networks as described in Manually Adding Networks to Your Custom Topology, on page 2125.

Step 7 Click **Save**.

What to do next

- Activate the topology as described in Activating and Deactivating Custom Topologies, on page 2125.

Importing Networks from the Network Discovery Policy

Procedure

Step 1 Access the custom topology to which you want to import the network:

- Create a custom topology; see Creating Custom Topologies, on page 2124.
- Edit an existing custom topology; see Editing Custom Topologies, on page 2126.

Step 2 Click **Import Policy Networks**.

Step 3 Click **Load**. The system displays the topology information for the network discovery policy.

Step 4 Refine your topology:

- Rename a network in the topology by clicking **Edit** () next to the network, typing a name, and clicking **Rename**.
- Remove a network from the topology by clicking **Delete** () and then clicking **OK** to confirm.

Step 5 Click **Save**.

What to do next

- Activate the topology as described in Activating and Deactivating Custom Topologies, on page 2125.

Manually Adding Networks to Your Custom Topology

Procedure

Step 1 Access the custom topology where you want to add the network:

- Create a custom topology; see Creating Custom Topologies, on page 2124.
- Edit an existing custom topology; see Editing Custom Topologies, on page 2126.

Step 2 Click **Add Network**.

Step 3 If you want to add a custom label for the network in the hosts and network devices network maps, type a **Name**.

Step 4 Enter the **IP Address** and **Netmask** (IPv4) that represent the network you want to add.

Step 5 Click **Add**.

Step 6 Click **Save**.

What to do next

- Activate the topology as described in Activating and Deactivating Custom Topologies, on page 2125.

Related Topics

Firepower System IP Address Conventions, on page 15

Activating and Deactivating Custom Topologies



Note Only one custom topology can be active at any time. If you have created multiple topologies, activating one automatically deactivates the currently active topology.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Choose **Custom Topology**.

Step 3 Click the slider next to a topology to activate or deactivate it.

Editing Custom Topologies

Changes you make to an active topology take effect immediately.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Custom Topology**.

Step 3 Click **Edit** ( ) next to the topology you want to edit.

Step 4 Edit the topology as described in Creating Custom Topologies, on page 2124.

Step 5 Click **Save**.



CHAPTER 105

Incidents

The following topics describe how to configure incident handling:

- About Incident Handling, on page 2127
- License Requirements for Incidents, on page 2131
- Requirements and Prerequisites for Incidents, on page 2131
- Creating Custom Incident Types, on page 2131
- Creating an Incident, on page 2132
- Editing an Incident, on page 2132
- Generating Incident Reports, on page 2133

About Incident Handling

Incident handling refers to the response an organization takes when a violation of its security policies is suspected. The Firepower System includes features to support you as you collect and process information that is relevant to your investigation of an incident. You can use these features to gather intrusion events and packet data that may be related to the incident. You can also use the incident as a repository for notes about any activity that you take outside of the Firepower System to mitigate the effects of the attack. For example, if your security policies require that you quarantine compromised hosts from your network, you can note that in the incident.

The Firepower System also supports an incident life cycle, allowing you to change an incident's status as you progress through your response to an attack. When you close an incident, you can note any changes you have made to your security policies as a result of any lessons learned.

Definition of an Incident

Generally, an *incident* is defined as one or more intrusion events that you suspect are involved in a possible violation of your security policies. In the Firepower System, the term also describes the feature you can use to track your response to an incident.

Some intrusion events are more important than others to the availability, confidentiality, and integrity of your network assets. For example, the port scan detection can keep you informed of port scanning activity on your network. Your security policy, however, may not specifically prohibit port scanning or see it as a high priority threat, so rather than take any direct action, you may instead want to keep logs of any port scanning for later forensic study.

On the other hand, if the system generates events that indicate hosts within your network have been compromised and are participating in distributed denial-of-service (DDoS) attacks, this activity is likely a clear violation of your security policy, and you should create an incident in the Firepower System to help you track your investigation of these events.

Common Incident Handling Processes

Preparation

You can prepare for incidents in two ways:

- by having clear and comprehensive security policies in place, as well as the hardware and software resources to enforce them
- by having a clearly defined plan to respond to incidents and a properly trained team that can implement the plan

A key part of incident handling is understanding which parts of your network are at the greatest risk. By deploying Firepower System components on those network segments, you can increase your awareness of when and how incidents occur. Also, by taking the time to carefully tune the intrusion policy for each managed device, you can ensure that the events that are generated are of the highest quality.

Detection and Notification

You cannot respond to an incident unless you can detect it. Your incident handling process should note the kinds of security-related events that you can detect and the mechanisms, both software and hardware, that you use to detect them. You should also note where you can detect violations of your security policies. If your network includes segments that are not actively or passively monitored, you need to note that as well.

The managed devices that you deploy on your network are responsible for analyzing the traffic on the segments where they are installed, for detecting intrusions, and for generating events that describe them. Keep in mind that the access control policy you deploy to each of the managed devices governs what kinds of activity they detect and how it is prioritized. You can also set notification options for certain types of intrusion events so that the incident team does not need to sift through hundreds of events. You can specify that you are notified automatically when certain high priority, high severity events are detected.

Investigation and Qualification

Your incident handling process should specify how, after a security incident is detected, an investigation is conducted. In some organizations, junior members of the team triage all the incidents and handle the less severe or lower priority cases themselves, while more senior members of the team handle high severity and high priority incidents. You should carefully outline the escalation process so that each team member understands the criteria for raising an incident's importance.

Part of the escalation process is tied to understanding how a detected event can affect the security of your network assets. For example, an attack against hosts running Microsoft SQL Server is not a high priority for organizations that use a different database server. Similarly, the attack is less important to you if you use SQL Server on your network, but you are confident that all the servers are patched and are not vulnerable to the attack. However, if someone has recently installed a copy of the vulnerable version of the software (perhaps for testing purposes), you may have a greater problem than a cursory investigation would suggest.

The Firepower System is particularly well suited to supporting the investigation and qualification process. You can create your own event classifications, and then apply them in a way that best describes the

vulnerabilities on your network. When traffic on your network triggers an event, that event is automatically prioritized and qualified for you with special indicators showing which attacks are directed against hosts that are known to be vulnerable.

The incident tracking feature in the Firepower System also includes a status indicator that you can change to show which incidents have been escalated.

Communication

All incident handling processes should specify how an incident is communicated between the incident handling team and both internal and external audiences. For example, you should consider what kinds of incidents require management intervention and at what level. Also, your process should outline how and when you communicate with outside organizations. Consider the following:

- Will some incidents require that you notify law enforcement agencies?
- If your hosts are participating in a distributed denial of service (DDoS) against a remote site, will you inform them?
- Do you want to share information with organizations such as the CERT Coordination Center (CERT/CC) or FIRST?

The Firepower System has features that you can use to gather intrusion data in standard formats such as HTML, PDF, and CSV (comma-separated values) so that you can easily share intrusion data with others.

For example, CERT/CC collects standard information about security incidents on its web site. CERT/CC looks for the kinds of information that you can easily extract from the Firepower System, such as:

- information about the affected machines, including:
  - the host name and IP
  - the time zone
  - the purpose or function of the host
- information about the sources of the attack, including:
  - the host name and IP
  - the time zone
  - whether you had any contact with an attacker
  - the estimated cost of handling the incident
- a description of the incident, including:
  - dates
  - methods of intrusion
  - the intruder tools involved
  - the software versions and patch levels
  - any intruder tool output
  - the details of vulnerabilities exploited

- the source of the attack
- any other relevant information

You can also use the comment section of an incident to record when you communicate issues and with whom.

Containment and Recovery

Your incident handling process should clearly indicate what steps are taken when a host or other network component is compromised. The range of containment and recovery options stretches from applying patches to vulnerable hosts to shutting down the target and removing it from the network. You should also consider the importance, depending upon the nature and severity of the attack, of preserving evidence in case you pursue criminal charges.

You can use the incident feature of Firepower System to maintain a record of the actions you take during the containment and recovery phase of the incident.

Lessons Learned

Each security incident, whether or not it is a successful attack, is an opportunity to review your security policies. Do you need to update your firewall rules? Do you need a more structured approach to patch management? Are unauthorized wireless access points a new security issue? Each lesson learned should feed back into your security policies and help you prepare better for the next incident.

Incident Types in the Firepower System

You can assign an incident type to each incident you create. The following types are supported by default in the Firepower System:

- Intrusion
- Denial of Service
- Unauthorized Admin Access
- Web Site Defacement
- Compromise of System Integrity
- Hoax
- Theft
- Damage
- Unknown

You can also create your own incident types.

License Requirements for Incidents

FTD License

Threat

Classic License

Protection

Requirements and Prerequisites for Incidents

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Intrusion Admin

Creating Custom Incident Types

Procedure

Step 1 Choose **Analysis** > **Intrusions** > **Incidents**.

Step 2 Click **Create Incident**.

Step 3 In the **Type** area, click **Types**.

The default incident types are listed at the bottom of the page.

Step 4 In the **Incident Type Name** field, enter a name for the new incident type.

Step 5 Click **Add**.

Step 6 Click **Done**.

You can use the new incident type the next time you create or edit an incident.

Creating an Incident

In a multidomain deployment, you can view and modify incidents created in the current domain only. In an ancestor domain, you can add events to incidents from any descendant domains.

Procedure

Step 1 Choose **Analysis** > **Intrusions** > **Incidents**.

Step 2 Click **Create Incident**.

Step 3 From the **Type** drop-down menu, choose the option that best describes the incident.

Step 4 In the **Time Spent** field, enter the amount of time you spent on the incident in the #d #h #m #s format, where # represents the number of days, hours, minutes, or seconds.

Step 5 In the **Summary** text box, enter a short description of the incident (up to 255 alphanumeric characters, spaces, and symbols).

Step 6 In the **Add Comment** text box, enter a more complete description for the incident (up to 8191 alphanumeric characters, spaces, and symbols).

Step 7 Add events to the incident:

- To add a selection of events, choose the events on the clipboard, and click **Add to Incident**.
- To add all events from the clipboard, click **Add All to Incident**.

Note If you want to add individual events from more than one page on the clipboard, you must add the events from one page, then add the events from the other pages separately.

Step 8 Click **Save**.

Editing an Incident

In a multidomain deployment, you can view and modify incidents created in the current domain only. In an ancestor domain, you can add events to an incident from any descendant domains.

Procedure

Step 1 Choose **Analysis** > **Intrusions** > **Incidents**.

Step 2 Click **Edit** () next to the incident you want to edit.

Step 3 You can edit any of the following aspects of the incident:

- change the status
- change the type
- add events from the clipboard
- delete events

Step 4 In the **Time Spent** field, enter the amount of additional time you spent on the incident.

Step 5 In the **Add Comment** text box, indicate your changes to the incident (up to 8191 alphanumeric characters, spaces and symbols) for the incident.

Step 6 Optionally, you can add or delete events from the incident:

- To add events from the clipboard, choose the events on the clipboard and click **Add to Incident**.
- To add all the events from the clipboard, click **Add All to Incident**.
- To delete specific events from the incident, choose the events and click **Delete**.
- To delete all events from the incident, click **Delete All**.
- To update the incident without adding or deleting events, click **Save**.

Generating Incident Reports

You can use the Firepower System to generate incident reports. These reports can include the incident summary, incident status, and any comments along with information from the events you add to the incident. You can also specify whether you want to include event summary information in the report.

Procedure

Step 1 Choose **Analysis** > **Intrusions** > **Incidents**.

Step 2 Click **Edit** () next to the incident you want to include in your report.

Step 3 You have two options:

- To include all the events from the incident in the report, click **Generate Report All**.
- To include specific events from the incident in the report, check the check boxes next to the events you want, and click **Generate Report**.

Step 4 Enter a name for the report.

Step 5 In **Incident Report Sections**, check the check boxes for the portions of the incident that you want to include in the report: **status**, **summary**, and **comments**.

Step 6 If you want to include event information in the report, choose the workflow you want to use and then, in **Report Sections**, specify whether you want to include event summary information.

Step 7 Check the check boxes next to the workflow pages you want to include in the report.

Step 8 Check the check boxes next to the output formats you want to use for the report: **PDF**, **HTML**, and **CSV**.

Note CSV-based incident reports include only event information. They do **not** include the status, summary, or comments from the incident.

Step 9 Click **Generate Report** and confirm that you want to update the report profile.



CHAPTER 106

Using Lookups

The following topics explain how to look up information about entities that may or may not be known to the Firepower System:

- Introduction to Lookups, on page 2135
- Performing Whois Lookups, on page 2135
- Finding URL Category and Reputation, on page 2136
- Finding Geolocation Information for an IP Address, on page 2137

Introduction to Lookups

If your Firepower Management Center is connected to the Internet, you can use manual lookup features to find the following information:

- Regional Information Registries (RIR) information (whois) for any IP address.
- URL category and reputation as classified by the URL Filtering feature.
- Geolocation information for any IP address: country name, country code, and continent name. (To ensure that you are using up-to-date geolocation information, Cisco strongly recommends that you regularly update the Geolocation Database (GeoDB) on your Firepower Management Center.)

Related Topics

Update the Geolocation Database, on page 141

Performing Whois Lookups

Before you begin

- Ensure that the Firepower Management Center has Internet access; see Security, Internet Access, and Communication Ports, on page 2425.

Procedure

Step 1 Choose **Analysis** > **Lookup** > **Whois**.

Step 2 Enter an IP address and click **Search**.

Related Topics

The Context Menu, on page 11

Finding URL Category and Reputation

You can manually look up category and reputation of URLs. Use this feature to see how particular URLs are evaluated in order to plan, adjust, or troubleshoot policy processing, or to investigate potentially problematic URLs that come to your attention via sources outside your Cisco solution. The categories and reputations in these results are the same as those that are used by the URL Filtering feature.

Before you begin

- The Firepower Management Center must have Internet access; see Security, Internet Access, and Communication Ports, on page 2425.
- URL Filtering and the **Query Cisco CSI for Unknown URLs** option must be enabled. See Enable URL Filtering Using Category and Reputation, on page 1235 and URL Filtering Options, on page 1235.
- You must be an Admin or Security Analyst user to perform this task.

Procedure

Step 1 Select **Analysis** > **Lookup** > **URL**.

Step 2 Enter up to 250 URLs and public, routable IP addresses, in any common format (for example, URLs may be with or without "http", "www", or a subdomain, or may be shortened). Separate each entity with a space or a return.

Wildcards such as asterisks (*) are not supported.

Step 3 Click **Search**.

If you enter many URLs and your network is slow, processing may take several minutes.

If you see an error message that the URL is not valid, check your spelling or try a different variation of the URL. For example, add or omit the "www" or "http" or "https" prefix.

A URL may belong to up to six categories but has only one reputation.

Step 4 (Optional) Sort the results by clicking a column heading.

Step 5 (Optional) To save the results as a CSV file, click **Export CSV**.

An additional column for reputation level is included in the CSV file so you can sort by risk. Zero (0) represents an unknown risk, for a URL for which the system has insufficient risk data.

What to do next

If you want to view lists of possible categories and reputations, go to **Policies > Access Control > Access Control**, click a policy or add a new one, click **Add Rule**, then click **URLs**.

Finding Geolocation Information for an IP Address

You can use the geolocation lookup feature to find the country name, ISO 3166-1 three-digit country code, and continent name associated with any IP address.

Procedure

Step 1 Choose **Analysis** > **Lookup** > **Geolocation**.

Step 2 To view the geolocation information for one or more IP addresses, enter the address or addresses and click **Search**. You may specify IPv4 addresses, IPv6 addresses, or both. Use a comma, semicolon, return, or any white space character to separate multiple addresses.

Tip Click **Clear** to clear the text box.

Step 3 Optionally, click the column titles to sort the data. You can sort by any field except IP Address.

Step 4 (Optional) To save the results as a CSV file, click **Export CSV**.

Related Topics

Update the Geolocation Database, on page 141



PART XXIV

Workflows

- Workflows, on page 2141
- Searching for Events, on page 2177
- Custom Workflows, on page 2187
- Custom Tables, on page 2195



CHAPTER 107

Workflows

The following topics describe how to use workflows:

- Overview: Workflows, on page 2141
- Predefined Workflows, on page 2141
- Custom Table Workflows, on page 2150
- Using Workflows, on page 2150
- Bookmarks, on page 2174

Overview: Workflows

A workflow is a tailored series of data pages on the Firepower Management Center web interface that analysts can use to evaluate events generated by the system.

The following types of workflows are available on the Firepower Management Center:

Predefined Workflows

Preset workflows delivered with the system. You cannot edit or delete a predefined workflow. You can, however, copy a predefined workflow and use it as the basis for a custom workflow.

Saved Custom Workflows

Custom workflows based on saved custom tables delivered with the Firepower Management Center. You can edit, delete, and copy these workflows.

Custom Workflows

Workflows that you create and customize for your specific needs, or that the system generates automatically when you create custom tables. You can edit, delete, and copy these workflows.

The data displayed in a workflow often depends on such factors as how you license and deploy your managed devices, and whether you configure features that provide the data.

Predefined Workflows

The predefined workflows described in the following sections are delivered with the system. You cannot edit or delete a predefined workflow, but you can copy a predefined workflow and use it as the basis for a custom workflow.

Predefined Intrusion Event Workflows

The following table describes the predefined intrusion event workflows included with the Firepower System.

Table 251: Predefined Intrusion Event Workflows

Workflow Name	Description
Destination Port	Because destination ports are usually tied to an application, this workflow can help you detect applications that are experiencing an uncommonly high volume of alerts. The Destination Port column can also help you identify applications that should not be present on your network.
Event-Specific	This workflow provides two useful features. Events that occur frequently may indicate: • false positives • a worm • a badly misconfigured network Events that occur infrequently are most likely evidence of a targeted attack and warrant special attention.
Events by Priority and Classification	This workflow lists events and their type in order of event priority, along with a count showing how many times each event has occurred.
Events to Destinations	This workflow provides a high-level view of which host IP addresses are being attacked and the nature of the attack; where available, you can also see information about the countries involved in attacks.
IP-Specific	This workflow shows which host IP addresses are generating the most alerts. Hosts with the greatest number of events are either public-facing and receiving worm-type traffic (indicating a good place to look for tuning) or require further investigation to determine the cause of the alerts. Hosts with the lowest counts also warrant investigation as they could be the subject of a targeted attack. Low counts may also indicate that a host may not belong on the network.
Impact and Priority	This workflow lets you find high-impact recurring events quickly. The reported impact level is shown with the number of times the event has occurred. Using this information, you can identify the high-impact events that recur most often, which might be an indicator of a widespread attack on your network.
Impact and Source	This workflow can help you identify the source of an attack in progress. The reported impact level is shown with the associated source IP address for the event. If, for example, events with a level 1 impact are coming from the same source IP address repeatedly, they may indicate an attacker who has identified vulnerable systems and is targeting them.
Impact to Destination	You can use this workflow to identify events repeatedly occurring on vulnerable computers, so you can address the vulnerabilities on those systems and stop any attacks in progress.
Source Port	This workflow indicates which servers are generating the most alerts. You can use this information to identify areas that require tuning, and to decide which servers require attention.
Source and Destination	This workflow identifies host IP addresses sharing high levels of alerts. Pairs at the top of the list could be false positives, and may identify areas that require tuning. You can check pairs at the bottom of the list for targeted attacks, for users accessing resources they should not be accessing, or for hosts that do not belong on the network.

Predefined Malware Workflows

The following table describes the predefined malware workflows included on the Firepower Management Center. All predefined malware workflows use the table view of malware events.

Table 252: Predefined Malware Workflows

Workflow Name	Description
Malware Summary	This workflow provides a list of the malware detected in network traffic or by AMP for Endpoints Connectors, grouped by individual threat.
Malware Event Summary	This workflow provides a quick breakdown of the different malware event types and subtypes.
Hosts Receiving Malware	This workflow provides a list of host IP addresses that have received malware, grouped by the malware files' associated dispositions.
Hosts Sending Malware	This workflow provides a list of host IP addresses that have sent malware, grouped by the malware files' associated dispositions.
Applications Introducing Malware	This workflow provides a list of host IP addresses that have received files, grouped by the associated malware dispositions for those files.

Predefined File Workflows

The following table describes the predefined file event workflows included on the Firepower Management Center. All the predefined file event workflows use the table view of file events.

Table 253: Predefined File Workflows

Workflow Name	Description
File Summary	This workflow provides a quick breakdown of the different file event categories and types, along with any associated malware dispositions.
Hosts Receiving Files	This workflow provides a list of host IP addresses that have received files, grouped by the associated malware dispositions for those files.
Hosts Sending Files	This workflow provides a list of host IP addresses that have sent files, grouped by the associated malware dispositions for those files.

Predefined Captured File Workflows

The following table describes the predefined captured file workflows included on the Firepower Management Center. All predefined captured file workflows use the table view of captured files.

Table 254: Predefined Captured File Workflows

Workflow Name	Description
Captured File Summary	This workflow provides a breakdown of captured files based on type, category, and threat score.
Dynamic Analysis Status	This workflow provides a count of captured files based on whether they have been submitted for dynamic analysis.

Predefined Connection Data Workflows

The following table describes the predefined connection data workflows included on the Firepower Management Center. All the predefined connection data workflows use the table view of connection data.

Table 255: Predefined Connection Data Workflows

Workflow Name	Description
Connection Events	This workflow provides a summary view of basic connection and detected application information, which you can then use to drill down to the table view of events.
Connections by Application	This workflow contains a graph of the 10 most active applications on the monitored network segment, based on the number of detected connections.
Connections by Initiator	This workflow contains a graph of the 10 most active host IP addresses on the monitored network segment, based on the number of connections where the host initiated the connection transaction.
Connections by Port	This workflow contains a graph of the 10 most active ports on the monitored network segment, based on the number of detected connections.
Connections by Responder	This workflow contains a graph of the 10 most active host IP addresses on the monitored network segment, based on the number of connections where the host IP was the responder in the connection transaction.
Connections over Time	This workflow contains a graph of the total number of connections on the monitored network segment over time.
Traffic by Application	This workflow contains a graph of the 10 most active applications on the monitored network segment, based on the number of kilobytes transmitted. Application counts reflect each detector that matched against an application connection. The same application session may be represented more than once in the list depending on whether an application protocol, web application, client detector, or internal detector matched the traffic, as well as whether the traffic originated from a mobile device or was part of an encrypted session. If the application was seen in a client flow and no specific client detector exists, a generic client may be reported. For example, you may see the same session of YouTube traffic reported as **YouTube** (because it matched a YouTube web application detector) and as **YouTube client** (because an internal YouTube detector matched against characteristics typically seen in a client session). Use the information in the connection events and network map for your network to determine more context for specific application connections.

Workflow Name	Description
Traffic by Initiator	This workflow contains a graph of the 10 most active host IP addresses on the monitored network segment, based on the total number of kilobytes transmitted from each address.
Traffic by Port	This workflow contains a graph of the 10 most active ports on the monitored network segment, based on the number of kilobytes transmitted.
Traffic by Responder	This workflow contains a graph of the 10 most active host IP addresses on the monitored network segment, based on the total number of kilobytes received by each address.
Traffic over Time	This workflow contains a graph of the total kilobytes transmitted on the monitored network segment over time.
Unique Initiators by Responder	This workflow contains a graph of the 10 most active responding host IP addresses on the monitored network segment, based on the number of unique initiators that contacted each address.
Unique Responders by Initiator	This workflow contains a graph of the 10 most active initiating host IP addresses on the monitored network segment, based on the number of unique responders that the addresses contacted.

Predefined Security Intelligence Workflows

The following table describes the predefined Security Intelligence workflows included on the Firepower Management Center. All the predefined Security Intelligence workflows use the table view of Security Intelligence events.

Table 256: Predefined Security Intelligence Workflows

Workflow Name	Description
Security Intelligence Events	This workflow provides a summary view of basic Security Intelligence and detected application information, which you can then use to drill down to the table view of events.
Security Intelligence Summary	This workflow is identical to the Security Intelligence Events workflow, but begins with the Security Intelligence Summary page, which lists security intelligence events by category and count only.
Security Intelligence with DNS Details	This workflow is identical to the Security Intelligence Events workflow, but begins with the Security Intelligence with DNS Details page, which lists Security Intelligence events by category and DNS-related characteristics.

Predefined Host Workflows

The following table describes the predefined workflows that you can use with host data.

Table 257: Predefined Host Workflows

Workflow Name	Description
Hosts	This workflow contains a table view of hosts followed by the host view. Workflow views based on the Hosts table allow you to easily view data on all IP addresses associated with a host.

Workflow Name	Description
Operating System Summary	You can use this workflow to analyze the operating systems in use on your network.

Predefined Indications of Compromise Workflows

The following table describes the predefined workflows that you can use with IOC (Indications of Compromise) data.

Table 258: Predefined Indications of Compromise Workflows

Workflow Name	Description
Host Indications of Compromise	This workflow begins with a summary view of IOC data grouped by count and category, and provides a detail view that further subdivides the summary data by event type. Access this workflow via the **Analysis > Hosts** menu.
Indications of Compromise by Host	You can use this workflow to gauge which hosts on your network are most likely to be compromised (based on IOC data). Access this workflow via the **Analysis > Hosts** menu.
User Indications of Compromise	This workflow begins with a summary view of IOC data grouped by count and category, and provides a detail view that further subdivides the summary data by event type. Access this workflow via the **Analysis > Users** menu.
Indications of Compromise by User	Use this workflow to gauge which users on your network are most likely to be involved in potential compromises (based on IOC data.) Access this workflow via the **Analysis > Users** menu.

Predefined Applications Workflows

The following table describes the predefined workflows that you can use with application data.

Table 259: Predefined Applications Workflows

Workflow Name	Description
Application Business Relevance	You can use this workflow to analyze running applications of each estimated business relevance level on your network, so you can monitor appropriate use of your network resources.
Application Category	You can use this workflow to analyze running applications of each category (such as email, search engine, or social networking) on your network, so you can monitor appropriate use of your network resources.
Application Risk	You can use this workflow to analyze running applications of each estimated security risk level on your network, so you can estimate the potential risk of users' activity and take appropriate action.

Workflow Name	Description
Application Summary	You can use this workflow to obtain detailed information about the applications and associated hosts on your network, so you can closely examine host application activity.
Applications	You can use this workflow to analyze running applications on your network, so you can gain an overview of how the network is being used.

Predefined Application Details Workflows

The following table describes the predefined workflows that you can use with application detail and client data.

Table 260: Predefined Application Details Workflows

Workflow Name	Description
Application Details	You can use this workflow to analyze the client applications on your network in more detail. The workflow then provides a table view of client applications, followed by the host view.
Clients	This workflow contains a table view of client applications, followed by the host view.

Predefined Servers Workflows

The following table describes the predefined workflows that you can use with server data.

Table 261: Predefined Servers Workflows

Workflow Name	Description
Network Applications by Count	You can use this workflow to analyze the most frequently used applications on your network.
Network Applications by Hit	You can use this workflow to analyze the most active applications on your network.
Server Details	You can use this workflow to analyze the vendors and versions of detected server application protocols in detail.
Servers	This workflow contains a table view of applications followed by the host view.

Predefined Host Attributes Workflows

The following table describes the predefined workflow that you can use with host attribute data.

Table 262: Predefined Host Attributes Workflows

Workflow Name	Description
Attributes	You can use this workflow to monitor IP addresses of hosts on your network and the hosts' status.

The Predefined Discovery Events Workflow

The following table describes the predefined workflow that you can use to view discovery and identity data.

Table 263: Predefined Discovery Event Workflows

Workflow Name	Description
Discovery Events	This workflow provides a detailed list, in table view form, of discovery events, followed by the host view.

Predefined User Workflows

The following table describes the predefined workflow that you can use to view user discovery and user identity data.

Table 264: Predefined User Workflows

Workflow Name	Description
Users	This workflow provides a list of user information collected by user identity sources.

Predefined Vulnerabilities Workflows

The following table describes the predefined vulnerabilities workflow included on the Firepower Management Center.

Table 265: Predefined Vulnerabilities Workflows

Workflow Name	Description
Vulnerabilities	You can use this workflow to review vulnerabilities in the database, including a table view of only those active vulnerabilities that apply to the detected hosts on your network.The workflow provides a vulnerability detail view, which contains a detailed description for every vulnerability that meets your constraints.

Predefined Third-Party Vulnerabilities Workflows

The following table describes the predefined third-party vulnerabilities workflows included on the Firepower Management Center.

Table 266: Predefined Third-Party Vulnerabilities Workflows

Workflow Name	Description
Vulnerabilities by IP Address	You can use this workflow to quickly see how many third-party vulnerabilities you have detected per host IP address on your monitored network.

Workflow Name	Description
Vulnerabilities by Source	You can use this workflow to quickly see how many third-party vulnerabilities you have detected per third-party vulnerability source, such as the QualysGuard Scanner.

Predefined Correlation and White List Workflows

There is a predefined workflow for each type of correlation data, white list events, white list violations, and remediation status events.

Table 267: Predefined Correlation Workflows

Workflow Name	Description
Correlation Events	This workflow contains a table view of correlation events.
White List Events	This workflow contains a table view of white list events.
Host Violation Count	This workflow provides a series of pages that list all the host IP addresses that violate at least one white list.
White List Violations	This workflow includes a table view of white list violations that lists all violations with the most recently detected violation at the top of the list. Each row in the table contains a single detected violation.
Status	This workflow contains a table view of remediation status, which includes the name of the policy that was violated and the name and status of the remediation that was applied.

Predefined System Workflows

The Firepower System is delivered with some additional workflows, including system events such as audit events and health events, as well as workflows that list results from rule update imports and active scans.

Table 268: Additional Predefined Workflows

Workflow Name	Description
Audit Log	This workflow contains a table view of the audit log that lists audit events.
Health Events	This workflow displays events triggered by the health monitoring policy.
Rule Update Import Log	This workflow contains a table view listing information about both successful and failed rule update imports.
Scan Results	This workflow contains a table view listing each completed scan.

Custom Table Workflows

You can use the custom tables feature to create tables that use the data from two or more types of events. This is useful because you can, for example, create tables and workflows that correlate intrusion event data with discovery data to allow simple searches for events that affect critical systems.

When you create a custom table, the system automatically creates a workflow that you can use to view the events associated with the table. The features in the workflow differ depending on which type of table you use. For example, custom table workflows based on the intrusion event table always end with the packet view. However, custom table workflows based on discovery events end with the host view.

Unlike workflows based on the predefined event tables, workflows based on custom tables do not have links to other types of workflows.

Using Workflows

Procedure

Step 1 Choose the appropriate menu path and option as described in Workflow Selection, on page 2152.

Step 2 Navigate within the current workflow:

- To view all of the columns available in your chosen event data type, use table view pages; see Using Table View Pages, on page 2158.
- To view a subset of the columns available in your chosen event data type, use drill-down pages; see Using Drill-Down Pages, on page 2158.
- To display the corresponding row in the next page of the workflow, click **Down-Arrow** ( ).
- To move among the pages of a multipage workflow, use the tools at the bottom of each page; see Workflow Page Traversal Tools, on page 2155.
- To view the same constraints applied within a workflow for a different type of event, click **Jump to** and choose the event view from the drop-down list.

Step 3 Modify the display of the current workflow:

- Check the check boxes by one or more rows on a page to indicate which row(s) you want to affect, then click one of the buttons at the bottom of the page (for example, **View**) to perform that action for all selected rows.
- Check the check box at the top of the row to select all the rows on the page, then click one of the buttons at the bottom of the page (for example, **View**) to perform that action for all rows on the page.
- Constrain the columns in the display by clicking **Close** ( ) in the column heading that you want to hide. In the pop-up window that appears, click **Apply**

  Tip To hide or show other columns, check or clear the appropriate check boxes before you click **Apply**. To add a disabled column back to the view, click the expand arrow to expand the search constraints, then click the column name under Disabled Columns.

- Constrain the data view by selected values for selected fields. For information, see Event View Constraints, on page 2171 and Compound Event View Constraints, on page 2172.
- Change the time constraints on the event view. The date range located in the upper right corner of the page sets a time range for events to include in the workflow; for information, see Event Time Constraints, on page 2165.

  Note Events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.

- To sort data by columns, click the name of a column. To reverse the sort order, click the column name again. The direction indicates which column the data is sorted by, and whether the sort is **Ascending** or **Descending**.
- Click a workflow page link to display that page using any active constraints. Workflow page links appear in the upper left corner of predefined workflow table views and drill-down pages, above events and below the workflow name.

Step 4 View additional data within the current workflow:

- To view the file's trajectory map in a new window, click network file trajectory in file name and SHA-256 hash value columns. The icon is different depending on the file status; see File Trajectory Icons, on page 2155.
- To display a pop-up window of the host profile associated with an IP address, click host profile in any IP address column. The icon is different depending on the file status; see Host Profile Icons, on page 2156.
- To view the Dynamic Analysis Summary report for the highest threat score associated with a file, click threat score in any threat score column. The icon is different depending on the file's highest threat score; see Threat Score Icons, on page 2156.
- To view user profile information, click **User** or, for users associated with an indication of compromise, **Red User** in any user identity column. The user icon is dimmed if that user cannot be in the database (that is, is an AMP for Endpoints Connector user).
- To view vulnerability details for third-party vulnerabilities, click **Vulnerability** in any third-party vulnerability ID column.
- When viewing aggregated data points, hover your pointer over the flag to view the country name.
- When viewing individual data points, click flag to view further geolocation details described in Geolocation, on page 2159.

Step 5 Navigate to a different workflow:

To view the same event type using a different workflow, click **(switch workflow)** next to the workflow title, then choose the workflow you want to use. Note that you **cannot** use a different workflow for scan results.

Workflow Access by User Role

Access to a workflow is determined by the user's role. See the table below for more information.

User Role	Accessible Workflows
Administrator	Can access any workflow, and are the only users who can access the audit log, scan results, and the rule update import log.
Maintenance User	Can access health events.
Security Analyst and Security Analyst (Read Only)	Can access intrusion, malware, file, connection, discovery, vulnerability, correlation, and health workflows.

Workflow Selection

The Firepower System provides predefined workflows for the types of data listed in the following table.

Table 269: Features Using Workflows

Feature	Menu Path	Option
Intrusion events	Analysis > Intrusions	Events Reviewed Events Clipboard Incidents
Malware events	Analysis > Files	Malware Events
File events	Analysis > Files	File Events
Captured files	Analysis > Files	Captured Files
Connection events	Analysis > Connections	Events
Security Intelligence events	Analysis > Connections	Security Intelligence Events
Host events	Analysis > Hosts	Network Map Hosts Indications of Compromise Applications Application Details Servers Host Attributes Discovery Events

Feature	Menu Path	Option
User events	Analysis > Users	User Activity Users Indications of Compromise
Vulnerability events	Analysis > Vulnerabilities	Vulnerabilities Third-Party Vulnerabilities
Correlation events	Analysis > Correlation	Correlation Events White List Events White List Violations Status
Audit events	System > Monitoring	Audit
Health events	System > Health > Events	n/a
Rule Update Import Log	System > Updates	n/a
Scan Results	Policies > Actions > Scanners	n/a

When you view any of the kinds of data described in the above table, events appear on the first page of the default workflow for that data. You can specify a different default workflow by configuring your event view settings. Note that workflow access depends on your user role.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Related Topics

Configuring Event View Settings, on page 33

Workflow Pages

Although the data in each type of workflow is different, all workflows share a common set of features. Workflows can include several types of pages. The actions you can perform on a workflow page depend on the type of page.

Drill-down and table view pages in workflows allow you to quickly narrow your view of the data so you can zero in on events that are significant to your analysis. Table view pages and drill-down pages both support many features you can use to constrain the set of events you want to view or to navigate the workflow. When viewing data on drill-down pages or in the table view in a workflow, you can sort the data in ascending or descending order based on any available column. If the database contains more events than can be displayed on a single workflow page, you can click the links at the bottom of the page to display more events. When you click one of these links, the time window automatically pauses so that you do not see the same events twice; you can unpause the time window when you are ready.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Table Views

Table views include a column for each of the fields in the database on which your workflow is based if the page is enabled by default.

For best performance, display only the columns you need. The more columns are displayed, the more resources are required to display the data.

Note that when you disable a column on a table view, the Firepower System adds the Count column to the event view if disabling the column could create two or more identical rows.



Important To avoid a significant performance hit, do not disable any of the following columns: First Packet, Last Packet, Device, Initiator IP, Responder IP, Source Port/ICMP Type, Destination Port/ICMP Code.

These fields uniquely identify each event, and removing any of them from the view automatically adds the count field to the table, which is a resource-intensive operation, particularly when there are a large number of fields in the view.

When you click on a value in a table view page, you constrain by that value.

When you create a custom workflow, you add a table view to it by clicking **Add Table View**.

Drill-Down Pages

Generally, drill-down pages are intermediate pages that you use to narrow your investigation to a few events before moving to a table view page. Drill-down pages contain a subset of columns that are available in the database.

For example, a drill-down page for discovery events might include only the IP Address, MAC Address, and Time columns. A drill-down page for intrusion events, on the other hand, might include the Priority, Impact Flag, Inline Result, and Message columns.

Drill-down pages allow you to narrow the scope of events you are viewing and to move forward in the workflow. If you click on a value in a drill-down page, for example, you constrain by that value and move to the next page in the workflow, focusing more closely on events that match your selected values. Clicking a value in a drill-down page does not disable the column where the value is, even if the page you advance to is a table view. Note that drill-down pages for predefined workflows always have a Count column. When you create a custom workflow, you add a drill-down page to it by clicking **Add Page**.

Graphs

Workflows based on connection data can include graph pages, also called *connection graphs*.

For example, a connection graph might display a line graph that shows the number of connections detected by the system over time. Generally, connection graphs are, like drill-down pages, intermediate pages that you use to narrow your investigation.

Final Pages

The final page of a workflow depends on the type of event on which the workflow is based:

- The host view is the final page for workflows based on applications, application details, discovery events, hosts, indications of compromise (IOC), servers, white list violations, host attributes, or third-party vulnerabilities. Viewing host profiles from this page allows you to easily view data on all IP addresses associated with hosts that have multiple addresses.

- The user detail view is the final page for workflows based on users, user activity, and user indications of compromise.
- The vulnerability detail view is the final page for workflows based on Cisco vulnerabilities.
- The packet view is the final page for workflows based on intrusion events.

Workflows based on other kinds of events (for example, audit log events or malware events) do not have final pages.

On the final page of a workflow, you can expand detail sections to view specific information about each object in the set you focused on over the course of the workflow. Although the web interface does not list the constraints on the final page of a workflow, previously set constraints are retained and applied to the set of data.

Workflow Page Navigation Tools

Workflow pages provide visual cues to facilitate navigating among them and choosing what information to display during event analysis.

Workflow Page Traversal Tools

If a workflow contains multiple pages of data, the bottom of each page displays the number of pages in the workflow, as well as the tools listed in the table below which you may use to navigate among the pages:

Table 270: Workflow Page Traversal Tools

Page Traversal Tool	Action
page number (To view a different page, enter the number you wish to view, then press Enter.)	view a different page
>	view the next page
<	view the previous page
>|	jump to the last page
|<	jump to the first page

File Trajectory Icons

When a workflow page provides the opportunity to view the trajectory map for a file in a new window, a network trajectory icon appears. This icon differs depending upon the file status.

Table 271: File Trajectory Icons

File Trajectory Icon	File Status
Clean	Clean

File Trajectory Icon	File Status
Malware	Malware
Custom detection	Custom detection
Unknown	Unknown
Unavailable	Unavailable

Host Profile Icons

When a workflow page provides the opportunity to view the host profile associated with an IP address in a pop-up window, a host profile icon appears. If the host profile icon is dimmed, you cannot view the host profile because that host cannot be in the network map (for example, `0.0.0.0`). This icon appears different depending on the status of the host.

Table 272: Host Profile Icons

Host Profile Icon	Host Status
	Host is not tagged as potentially compromised.
	Host is tagged as potentially compromised by triggered indications of compromise (IOC) rules.
	Added to Block List (Appears only if you are performing traffic filtering based on Security Intelligence data.)
	Added to Block List, set to monitor (Appears only if you are performing traffic filtering based on Security Intelligence data.)

Threat Score Icons

When a workflow page provides the opportunity to view a Dynamic Analysis Summary report for the highest threat score associate with a file, a threat score icon appears. The icon differs depending on the file's highest threat score.

Table 273: Threat Score Icons

Threat Score Icon	Threat Score Level
Low	Low
Medium	Medium
High	High
Very High	Very high

User Icons

When a workflow page provides the opportunity to view the user identity associated with a username in a pop-up window, a user icon appears.

Table 274: User Icons

User Icon	User Status
User	User is not associated with any indications of compromise.
Red User	User is associated with one or more indications of compromise.

The Workflow Toolbar

Each page in a workflow includes a toolbar that offers quick access to related features. The following table describes each of the links on the toolbar.

Table 275: Workflow Toolbar Links

Feature	Description
Bookmark This Page	Bookmarks the current page so you can return to it later. Bookmarking captures the constraints in effect on the page you are viewing so you can return to the same data (assuming the data still exists) at a later time.
Report Designer	Opens the report designer with the currently constrained workflow as the selection criteria.
Dashboard	Opens a dashboard relevant to your current workflow. For example, Connection Events workflows link to the Connection Summary dashboard.
View Bookmarks	Displays a list of saved bookmarks from which you can select.
Search	Displays a Search page where you can perform advanced searches on data in the workflow. You can also click the down arrow icon to select and use a saved search.

Related Topics

Creating a Report Template from an Event View, on page 2055
About Dashboards, on page 255
Event Searches, on page 2177
Bookmarks, on page 2174
Creating Bookmarks, on page 2175
Viewing Bookmarks, on page 2175

Using Drill-Down Pages

Procedure

Step 1 Access a workflow by choosing the appropriate menu path and option as described in Features Using Workflows.

Step 2 In any workflow, you have the following options:

- To drill down to the next workflow page constraining on a specific value, click a value within a row. Note that this works only on drill-down pages. Clicking a value within a row in a table view only constrains the table view and does not drill down to the next page.
- To drill down to the next workflow page constraining on some events, check the check boxes next to the events you want to view on the next workflow page, then click **View**.
- To drill down to the next workflow page keeping the current constraints, click **View All**.

Tip Table views always include "Table View" in the page name.

Using Table View Pages

Table view pages provide some features not available on drill-down, host view, packet view, or vulnerability detail pages. Use these features as described below:

Procedure

Step 1 Access a workflow by choosing the appropriate menu path and option as described in Workflow Selection, on page 2152.

Step 2 Choose a table view from the workflow path displayed beneath the workflow name.

Step 3 Use the features listed below to arrange and navigate within the table view as needed:

- To display the list of disabled columns, click the Search Constraints **Expand Arrow** (▸).
- To hide the list of disabled columns, click the Search Constraints **Collapse Arrow** (▾).
- To add a disabled column back to the event view, click the Search Constraints **Expand Arrow** (▸) to expand the search constraints, then click the column name under Disabled Columns.
- To show or hide (disable) a column, click **Clear** (✖) next to any column name. In the pop-up window that appears, check or clear the appropriate check boxes to indicate which columns you want to display, then click **Apply**.

Geolocation

You can view and filter traffic based on country and continent by leveraging a geolocation database (GeoDB). Note that for mobile devices and other hosts detected moving from country to country, the system may report a continent instead of a specific country.

The system comes with an initial GeoDB that maps IP addresses to countries/continents, so that information should always be available. If you update the GeoDB, the system also downloads contextual data. This can include:

- Region (state, province, or other country subregion), city, and postal code.
- Latitude/longitude, time zone, and clickable maps.
- Autonomous System Number (ASN) and additional information about the ASN.
- Internet service provider (ISP), connection type, and proxy type.
- Home/business, organization, and domain name information.

To view this information, click the small country flag icons and ISO country codes wherever they appear: in events, asset profiles, the Context Explorer, dashboard, and other analysis tools. You cannot view geolocation details for aggregate geolocation information, such as on the Connection Summary dashboard.



Note We issue periodic updates to the GeoDB. You must regularly update the GeoDB to have accurate geolocation information; see Update the Geolocation Database, on page 141.

In May 2022 we split the GeoDB into two packages: a country code package that maps IP addresses to countries/continents, and an IP package that contains contextual data. The new country code package has the same file name as the old all-in-one package. This allows FMCs running Version 7.1 and earlier to continue to obtain GeoDB updates. However, because this package now contains only country code mappings, the contextual data is no longer updated and will grow stale. To obtain fresh data, upgrade or reimage to Version 7.2+ and update the GeoDB. Note that this split does not affect geolocation rules or traffic handling in any way—those rules rely only on the data in the country code package.

Related Topics

Network Conditions, on page 368
Geolocation Objects, on page 409
Introduction to Correlation Policies and Rules, on page 1985
Traffic Profile Conditions, on page 2023
Update the Geolocation Database, on page 141

Connection Event Graphs

In addition to workflows that use tabular drill-down pages and a final table view of events, the system can present certain connection data graphically, using data aggregated over five-minute intervals. Note that you can graph only the information used to aggregate data: source and destination IP addresses (and those hosts' associated users), destination port, transport protocol, and application protocol.

Tip You cannot graph Security Intelligence events separately from their associated connection events. For a graphical overview of Security Intelligence filtering activity, use dashboards and the Context Explorer.

There are three different types of connection graphs:

- *Pie charts* display data from one dataset grouped into discrete categories.
- *Bar graphs* display data from one or more datasets grouped into discrete categories.
- *Line graphs* plot data from one or more datasets over time, using either a standard or a velocity (rate of change) view.

Note The system displays traffic profiles as line graphs, which you can manipulate in the same way as you would any other connection graph, with some restrictions. To view traffic profiles, you must have Administrator access.

Like workflow tables, you can drill down and constrain workflow graphs to focus your analysis.

Both bar graphs and line graphs can display multiple datasets; that is, they can display several values on the y-axis for each x-axis data point. For example, you could display the total number of unique initiators and responders. Pie charts can only display one dataset.

You can display different data and datasets on a connection graph by changing either the x-axis, the y-axis, or both. On a pie chart, changing the x-axis changes the independent variable and changing the y-axis changes the dependent variable.

Related Topics

Connection Summaries (Aggregated Data for Graphs), on page 2224

Using Connection Event Graphs

On the Firepower Management Center, you can view connection event graphs and manipulate them depending on the information you are looking for.

The page you see when you access connection graphs differs depending on the workflow you use. You can use a predefined workflow, which terminates in a table view of connection events. You can also create a custom workflow that displays only the information that matches your specific needs.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Connections** > **Events**.

Note If a connection event table appears instead of a graph, or to view a different graph, click **(switch workflow)** by the workflow title and choose a predefined workflow that includes graphs, or a custom workflow. Note that all predefined connection event workflows—including connection graphs—terminate in a table view of connections.

Step 2 You have the following options:

- Time Range — To adjust the time range, which is useful if the graph is blank, see Changing the Time Window, on page 2168.
- Field Names — To learn more about the data you can graph, see Connection and Security Intelligence Event Fields, on page 2225.
- Host Profile — To view the host profile for an IP address, on a graph displaying connection data by initiator or responder, click either a bar on a bar graph or a wedge on a pie chart and choose **View Host Profile**.
- User Profile — To view user profile information, on a graph displaying connection data by initiator user, click either a bar on a bar graph or a wedge on a pie chart and choose **View User Profile**.
- Other Information — To learn more information about the graphed data, position your cursor over a point on a line graph, a bar in a bar graph, or a wedge in a pie chart.
- Constrain — To constrain a connection graph by any x-axis (independent variable) criterion without advancing the workflow to the next page, click a point on a line graph, a bar on a bar graph, or a wedge on a pie chart, and choose a **View by...** option.
- Data Selection — To change the data displayed on the graph, click **X-Axis** or **Y-Axis** and choose the new data to graph. Note that changing the x-axis to or from **Time** also changes the graph type; changing the y-axis affects the displayed datasets.
- Datasets — To change the graph's dataset, click **Datasets** and choose a new dataset.
- Detach — To detach a connection graph so you can perform further analysis without affecting the default time range, click **Detach**.

  Tip Click **New Window** in a detached graph to create a copy. You can then perform different analyses on each of the detached graphs. Note that traffic profiles are detached graphs.

- Drill Down — To drill down to the next page in the workflow, click a point on a line graph, a bar on a bar graph, or a wedge on a pie chart, then choose **Drill-down**. Clicking a point on a line graph changes the time range on the next page to a 10-minute span, centered on the point you clicked. Clicking a bar on a bar graph or a wedge on a pie chart constrains the next page based on the criterion represented by the bar or wedge.
- Export — To export the connection data for a graph as a CSV (comma-separated values) file, **Export Data**. Then, click **Download CSV File** and save the file.
- Graph Type: Line — To switch between a standard and velocity (rate of change) line graph, click **Velocity**, then choose **Standard** or **Velocity**.
- Graph Type: Bar and Pie — To switch between a bar graph and pie chart, click **Switch to Bar** or **Switch to Pie**. Because you cannot display multiple datasets on a pie char, if you switch to a pie chart from a bar graph that has multiple datasets, the pie chart shows only one dataset, which is selected automatically. When choosing which dataset to display, the Firepower Management Center favors total statistics over initiator and responder statistics, and favors initiator statistics over responder statistics.
- Navigate Between Pages — To navigate between pages in the current workflow, keeping the current constraints, click the appropriate page link at the top left of the workflow page.
- Navigate Between Event Views — To navigate to other event views to view associated events, click **Jump to** and choose the event view from the drop-down list.

- Recenter — To recenter a line graph around a point in time without changing the length of the time range, click that point, then choose **Recenter**.
- Zoom -— To recenter a line graph around a point in time while zooming in or out, click that point, choose **Zoom**, then choose a new time span.

Note Unless you are working with a detached graph, constraining, recentering, and zooming changes the default time range for the Firepower Management Center.

Example: Constraining a Connection Graph

Example: Changing X-Axis and Y-Axis on a Pie Chart

Consider a graph of connections over time. If you constrain a point on the graph by port, a bar graph appears, showing the 10 most active ports based on the number of detected connection events, but constrained by the ten-minute time span that is centered on the point you clicked.

If you further constrain the graph by clicking on one of the bars and choosing **View by Initiator IP**, a new bar graph appears, constrained by not only the same ten-minute time span as before, but also by the port represented by the bar you clicked.

Consider a pie chart that graphs kilobytes per port. In this case, the x-axis is **Responder Port** and the y-axis is **KBytes**. This pie chart represents the total kilobytes of data transmitted over a monitored network during a certain interval. The wedges of the pie represent the percent of the data that was detected on each port.

- If you change the x-axis of the chart to **Application Protocol**, the pie chart still represents the total kilobytes of data transmitted, but the wedges of the pie represent the percentage of the data transmitted for each detected application protocol.
- If you change the y-axis of the chart to **Packets**, the pie chart represents the total number of packets transmitted over the monitored network during a certain interval, and the wedges of the pie represent the percentage of the total number of packets that was detected on each port.

Related Topics

Using Workflows, on page 2150
Configuring Event View Settings, on page 33

Connection Graph Data Options

You can display different data on a connection graph by changing either the x-axis, the y-axis, or both. On a pie chart, changing the x-axis changes the independent variable and changing the y-axis changes the dependent variable.

Table 276: X-Axis Options

X-Axis Option	Graph Type	Graphs This Data
Application Protocol	bar or pie	by the 10 most active application protocols
Device	bar or pie	by the 10 most active managed devices

X-Axis Option	Graph Type	Graphs This Data
Initiator IP	bar or pie	by the 10 most active initiator host IP addresses
Initiator User	bar or pie	by the 10 most active initiator users
Responder IP	bar or pie	by the 10 most active responder host IP addresses
Responder Port	bar or pie	by the 10 most active responder ports
Source Device	bar or pie	by the 10 most active NetFlow data exporters, plus a source device named `Firepower` for all connections detected by Firepower System managed devices.
Time	line	over time Changing the y-axis to and from **Time** also changes the graph type and may change the datasets.

Table 277: Y-Axis Options

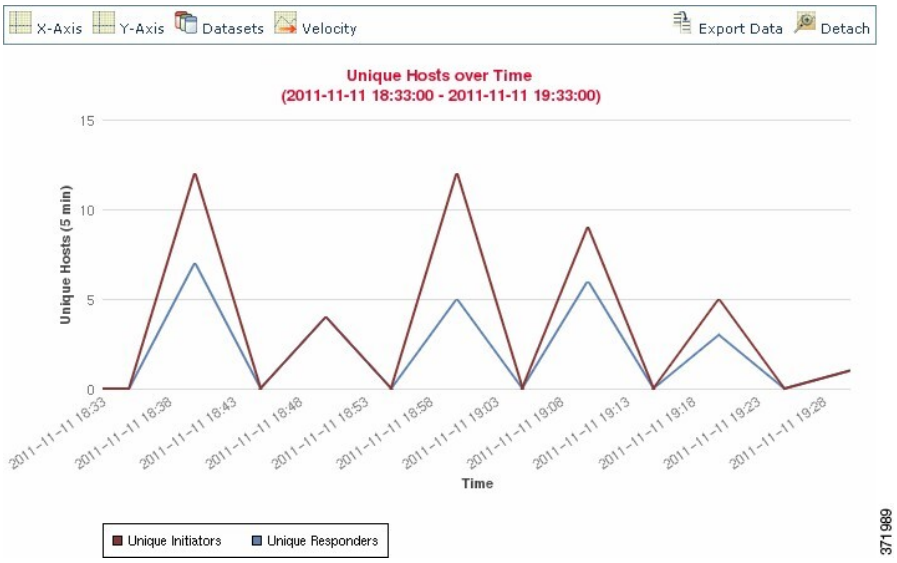
Y-Axis Option	Graphs This Data Using The X-Axis Criterion
Bytes	bytes transmitted
Connections	number of connections
KBytes	kilobytes transmitted
KBytes Per Second	kilobytes per second
Packets	number of packets transmitted
Unique Hosts	number of unique hosts detected
Unique Application Protocols	number of unique application protocols
Unique Users	number of unique users

Connection Graphs with Multiple Datasets

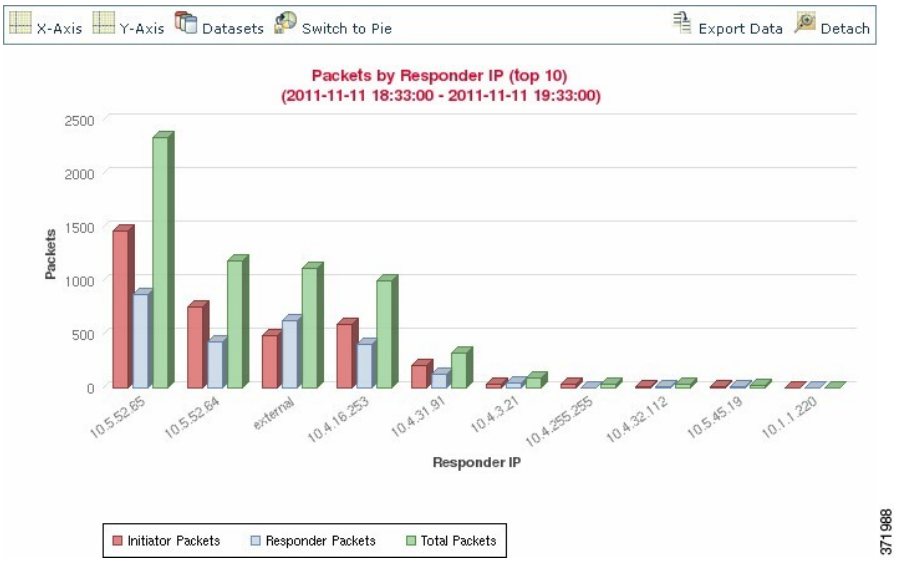
Both bar graphs and line graphs can display multiple datasets; that is, they can display several values on the y-axis for each x-axis data point. For example, you could display the total number of unique initiators and responders.

Note You **cannot** display multiple datasets on a pie chart. If you switch to a pie chart from a bar graph that has multiple datasets, the pie chart shows only one dataset, which is selected automatically. When selecting which dataset to display, the Firepower Management Center favors total statistics over initiator and responder statistics, and favors initiator statistics over responder statistics.

On line graphs, multiple datasets appear as multiple lines, each with a different color. For example, the following graphic displays the total number of unique initiators and the total number of unique responders detected on a monitored network over a one hour interval.



On bar graphs, multiple datasets appear as a set of colored bars for each x-axis data point. For example, the following bar graph displays the total packets transmitted on a monitored network, packets transmitted by initiators, and packets transmitted by responders.



Connection Graph Dataset Options

The following table describes the datasets you can display on the x-axis of a connection graph.

Table 278: Dataset Options

If the y-axis displays...	You can select as datasets...
Connections	the default only, which is the number of connections detected on the monitored network (**Connections**). This is the only option for traffic profile graphs.
KBytes	combinations of: • the total kilobytes transmitted on the monitored network (**Total KBytes**) • the number of kilobytes transmitted from host IP addresses on the monitored network (**Initiator KBytes**) • the number of kilobytes received by host IP addresses on the monitored network (**Responder KBytes**)
KBytes Per Second	the default only, which is the total kilobytes per second transmitted on the monitored network (**Total KBytes Per Second**)
Packets	combinations of: • the total packets transmitted on the monitored network (**Total Packets**) • the number of packets transmitted from host IP addresses on the monitored network (**Initiator Packets**) • the number of packets received by host IP addresses on the monitored network (**Responder Packets**)
Unique Hosts	combinations of: • the number of unique session initiators on the monitored network (**Unique Initiators**) • the number of unique session responders on the monitored network (**Unique Responders**)
Unique Application Protocols	the default only, which is the number of unique application protocols on the monitored network (**Unique Application Protocols**)
Unique Users	the default only, which is the number of unique users logged into session initiators on the monitored network (**Unique Initiator Users**)

Event Time Constraints

Each event has a time stamp that indicates when the event occurred. You can constrain the information that appears in some workflows by setting the time window, sometimes called the time range.

Workflows based on events that can be constrained by time include a time range line at the top of the page.

By default, workflows use an expanding time window set to the past hour. For example, if you log in at 11:30 AM, you will see events that occurred between 10:30 AM and 11:30 AM. As time moves forward, the time window expands. At 12:30 PM, you will see events that occurred between 10:30 AM and 12:30 PM.

You can change this behavior by setting your own default time window in the event view settings. This governs three properties:

- time window type (static, expanding, or sliding)
- time window length
- the number of time windows (either multiple time windows or a single global time window)

Regardless of the default time window setting, you can manually change the time window during your event analysis by clicking the time range at the top of the page, which displays the Date/Time pop-up window. Depending on the number of time windows you configured and the type of appliance you are using, you can also use the Date/Time window to change the default time window for the type of event you are viewing.

Finally, you can pause the time window while looking at a sliding or expanding workflow. See Pause the Time Window to Temporarily Freeze the Data Set, on page 2168.

Related Topics

Configuring Event View Settings, on page 33
Using Connection and Security Intelligence Event Tables, on page 2245

Per-Session Time Window Customization for Events

Regardless of the default time window, you can manually change the time window during your event analysis.



Note Manual time window settings are valid for only the current session. When you log out and then log back in, time windows are reset to the default.

Depending on the number of time windows you configured, changing the time window for one workflow may affect other workflows on the appliance. For example, if you have a single, global time window, changing the time window for one workflow changes it for all other workflows on the appliance. On the other hand, if you are using multiple time windows, changing the audit log or health event workflow time windows has no effect on any other time window, while changing the time window for other kinds of events affects all events that can be constrained by time (with the exception of audit events and health events).

Note that because not all workflows can be constrained by time, time window settings have no effect on workflows based on hosts, host attributes, applications, application details, vulnerabilities, users, or white list violations.

Use the Time Window tab on the Date/Time window to manually configure a time window. Depending on the number of time windows you configured in your default time window settings, the tab's title is one of the following:

- **Events Time Window**, if you configured multiple time windows and are setting the time window for a workflow other than the audit log or health events workflow
- **Health Monitoring Time Window**, if you configured multiple time windows and are setting the time window for the health events workflow
- **Audit Log Time Window**, if you configured multiple time windows and are setting the time window for the audit log
- **Global Time Window**, if you configured a single time window

The first decision you must make when configuring a time window is the type of time window you want to use:

- A *static* time window displays all the events generated from a specific start time to a specific end time.
- An *expanding* time window displays all the events generated from a specific start time to the present; as time moves forward, the time window expands and new events are added to the event view.
- A *sliding* time window displays all the events generated from a specific start time (for example, one week ago) to the present; when you refresh the page, the time window "slides" so that you see only the events in the time range you configured (in this example, for the last week). To temporarily prevent the data set from updating while you are examining it, see Pause the Time Window to Temporarily Freeze the Data Set, on page 2168.

Depending on what type you select, the Date/Time window changes to give you different configuration options.

Note The Firepower System uses a 24-hour clock based on the time you specified in your time zone preferences.

Time Window Settings

The following table explains the various settings you can configure on the Time Window tab.

Table 279: Time Window Settings

Setting	Time Window Type	Description
time window type drop-down list	n/a	Select the type of time window you want to use: static, expanding, or sliding. Note that events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.
Start Time calendar	static and expanding	Specify a start date and time for your time window. The maximum time range for all time windows is from midnight on January 1, 1970 (UTC) to 3:14:07 AM on January 19, 2038 (UTC). Instead of using the calendar, you can use the Presets options, described below.
End Time calendar	static	Specify an end date and time for your time window. The maximum time range for all time windows is from midnight on January 1, 1970 (UTC) to 3:14:07 AM on January 19, 2038 (UTC). Note that if you are using an expanding time window, the End Time calendar is grayed out and specifies that the end time is "Now." Instead of using the calendar, you can use the Presets options, described below.
Show the Last field and drop-down list	sliding	Configure the length of the sliding time window.

Setting	Time Window Type	Description
Presets: Last	all	Click one of the time ranges in the list to change the time window, based on the local time of the appliance. For example, clicking **1 week** changes the time window to reflect the last week. Clicking a preset changes the calendars to reflect the preset you choose.
Presets: Current	static and expanding	Click one of the time ranges in the list to change the time window, based on the local time and date of the appliance. Clicking a preset changes the calendars to reflect the preset you choose. Note that: • the current day begins at midnight • the current week begins at midnight Sunday • the current month begins at midnight on the first of the month
Presets: Synchronize with	all (not available if you are using a global time window)	Click one of: • **Events Time Window** to synchronize the current time window with the events time window • **Health Monitoring Time Window** to synchronize the current time window with the health monitoring time window • **Audit Log Time Window** to synchronize the current time window with the audit log time window

Changing the Time Window

Procedure

Step 1 On a workflow constrained by time, click **Time Range** ( ) to go to the Date/Time window.

Step 2 On **Events Time Window**, set the time window as described in Time Window Settings, on page 2167.

Tip Click **Reset** to change the time window back to the default settings.

Step 3 Click **Apply**.

Pause the Time Window to Temporarily Freeze the Data Set

If you are using a sliding or expanding time window, you can pause the time window to examine a snapshot of the data provided by the workflow. This is useful because when an unpaused workflow updates, it may remove events that you want to examine or add events that you are not interested in.

The time window automatically pauses when you click a link at the bottom of the page to display another page of events; you can unpause the time window when you are ready.

When you are finished with your analysis, you can unpause the time window. Unpausing the time window updates it according to your preferences, and also updates the event view to reflect the unpaused time window.

Pausing an event time window has no effect on dashboards, nor does pausing a dashboard have any effect on pausing an event time window.

Procedure

On a workflow constrained by time, choose the desired time range control:

- To pause the time window, click time range control **Pause** ( ).
- To unpause the time window, click time range control **Play** ( ).

The Default Time Window for Events

During your event analysis, you can use the Preferences tab on the Date/Time window to change the default time window for the type of event you are viewing without having to use the event view settings.

Keep in mind that changing the default time window in this way changes the default time window for only the type of event you are viewing. For example, if you configured multiple time windows, changing the default time window on the Preferences tab changes the settings for either the events, health monitoring, or audit log window, in other words, whichever time window is indicated by the first tab. If you configured a single time window, changing the default time window on the Preferences tab changes the default time window for all types of events.

Related Topics

Default Time Windows, on page 35

Default Time Window Options for Event Types

The following table explains the various settings you can configure on the Preferences tab.

Table 280: Time Window Preferences

Preference	Description
Refresh Interval	Sets the refresh interval for event views, in minutes. Entering zero disables the refresh option.
Number of Time Windows	Specify how many time windows you want to use: • Select **Multiple** to configure separate default time windows for the audit log, for health events, and for workflows based on events that can be constrained by time. • Select **Single** to use a global time window that applies to all events.
Default Time Window: Show the Last - Sliding	This setting allows you to configure a sliding default time window of the length you specify. The appliance displays all the events generated from a specific start time (for example, 1 hour ago) to the present. As you change event views, the time window “slides” so that you always see events from the last hour.

Preference	Description
Default Time Window: Show the Last - Static/Expanding	This setting allows you to configure either a static or expanding default time window of the length you specify. For **static** time windows (enable the **Use End Time** check box), the appliance displays all the events generated from a specific start time (for example, 1 hour ago), to the time when you first viewed the events. As you change event views, the time window stays fixed so that you see only the events that occurred during the static time window. For **expanding** time windows (disable the **Use End Time** check box), the appliance displays all the events generated from a specific start time (for example, 1 hour ago), to the present. As you change event views, the time window expands to the present time.
Default Time Window: Current Day - Static/Expanding	This setting allows you to configure either a static or expanding default time window for the current day. The current day begins at midnight, based on the time zone setting for your current session. For **static** time windows (enable the **Use End Time** check box), the appliance displays all the events generated from midnight to the time when you first viewed the events. As you change event views, the time window stays fixed so that you see only the events that occurred during the static time window. For **expanding** time windows (disable the **Use End Time** check box), the appliance displays all the events generated from midnight to the present. As you change event views, the time window expands to the present time. Note that if your analysis continues for over 24 hours before you log out, this time window can be more than 24 hours.
Default Time Window: Current Week - Static/Expanding	This setting allows you to configure either a static or expanding default time window for the current week. The current week begins at midnight on the previous Sunday, based on the time zone setting for your current session. For **static** time windows (enable the **Use End Time** check box), the appliance displays all the events generated from midnight to the time when you first viewed the events. As you change event views, the time window stays fixed so that you see only the events that occurred during the static time window. For **expanding** time windows (disable the **Use End Time** check box), the appliance displays all the events generated from midnight Sunday to the present. As you change event views, the time window expands to the present time. Note that if your analysis continues for over 1 week before you log out, this time window can be more than 1 week.

Changing the Default Time Window for Your Event Type

Procedure

Step 1 On a workflow constrained by time, click **Time Range** () to go to the Date/Time window.

Step 2 Click **Preferences** and change your preferences, as described in Default Time Window Options for Event Types, on page 2169.

Step 3 Click **Save Preferences**.

Step 4 You have two options:

- To apply your new default time window settings to the event view you are using, click **Apply** to close the Date/Time window and refresh the event view.
- To continue with your analysis without applying the default time window settings, close the Date/Time window without clicking **Apply**.

Event View Constraints

The information that you see on a workflow page is determined by the constraints that you impose. For example, when you initially open an event workflow, the information is constrained to events that were generated in the previous hour.

To advance to the next page in the workflow and constrain the data you are viewing by specific values, select the rows with those values on the page and click **View**. To advance to the next page in the workflow retaining the current constraints and carrying forward all events, select **View All.**



Note If you select a row with multiple non-count values and click **View**, you create a compound constraint.

There is a third method for constraining data in a workflow. To constrain the page to the rows with values that you selected and also add the selected value to the list of constraints at the top of the page, click a value within a row on the page. For example, if you are viewing a list of logged connections and want to constrain the list to only those you allowed using access control, click **Allow** in the **Action** column. As another example, if you are viewing intrusion events and want to constrain the list to only events where the destination port is 80, click **80 (http)/tcp** in the **Destination Port/ICMP Code** column.



Tip The procedure for constraining connection events based on Monitor rule criteria is slightly different and you may need to take some extra steps. Additionally, you cannot constrain connection events by associated file or intrusion information.

You can also use searches to constrain the information in a workflow. Use this feature when you want to constrain against multiple values in a single column. For example, if you want to view the events related to two IP addresses, click **Edit Search**, then modify the appropriate IP address field on the Search page to include both addresses, and then click **Search**.

The search criteria you enter on the search page are listed as the constraints at the top of the page, with the resulting events constrained accordingly. On the Firepower Management Center, the current constraints are also applied when navigating to other workflows, unless they are compound constraints.

When searching, you must pay careful attention to whether your search constraints apply to the table you are searching. For example, client data is not available in connection summaries. If you search for connection events based on the detected client in the connection and then view the results in a connection summary event view, the Firepower Management Center displays connection data as if you had not constrained it at all. Invalid constraints are labeled as not applicable (N/A) and are marked with a strikethrough.

Constraining Events

Procedure

Step 1 Access a workflow by choosing the appropriate menu path and option as described in Workflow Selection, on page 2152.

Step 2 In any workflow, you have the following options:

- To constrain the view to events that match a single value, click the desired value within a row on the page.
- To constrain the view to events that match multiple values, check the check boxes for events with those values, and click **View**.

  Note A compound constraint is added if the row contains multiple non-count values.

- To remove a constraint, click the Search Constraints **Expand Arrow** ( ) and click the name of the constraint in the expanded Search Constraints list.
- To edit constraints using the Search page, click **Edit Search.**
- To save constraints as a saved search, click **Save Search** and give the query a name.

  Note You cannot save queries containing compound constraints.

- To use the same constraints with another event view, click **Jump to** and choose the event view.

  Note You do not retain compound constraints when you switch to another workflow.

- To toggle the display of constraints click the Search Constraints **Expand Arrow** ( ) or the Search Constraints **Collapse Arrow** ( ). This is useful when the list of constraints is large and takes up most of the screen.

Compound Event View Constraints

Compound constraints are based on all non-count values for a specific event. When you select a row with multiple non-count values, you set a compound constraint that retrieves only events matching all the non-count values in that row on that page. For example, if you select a row that has a source IP address of `10.10.31.17` and a destination IP address of `10.10.31.15` and a row that has a source IP address of `172.10.10.17` and a destination IP address of `172.10.10.15`, you retrieve all of the following:

- Events that have a source IP address of 10.10.31.17 AND a destination IP address of 10.10.31.15

  OR

- Events that have a source IP address of 172.10.31.17 AND a destination IP address of 172.10.31.15

When you combine compound constraints with simple constraints, the simple constraints are distributed across each set of compound constraints. If, for example, you added a simple constraint for a protocol value of `tcp` to the compound constraints listed above, you retrieve all of the following:

- Events that have a source IP address of 10.10.31.17 AND a destination IP address of 10.10.31.15 AND a protocol of tcp

  OR

- Events that have a source IP address of 172.10.31.17 AND a destination IP address of 172.10.31.15 AND a protocol of tcp

You cannot perform a search or save a search on a compound constraint. You also cannot retain compound constraints when you use the event view links or click **(switch workflow)** to switch to another workflow. If you bookmark an event view with compound constraints applied, the constraints are not saved with the bookmark.

Using Compound Event View Constraints

Procedure

Step 1 Access a workflow by choosing the appropriate menu path and option as described in Workflow Selection, on page 2152.

Step 2 To manage compound constraints, you have the following options:

- To create a compound constraint, choose one or more rows with multiple non-count values and click **View**.
- To clear compound constraints, click the Search Constraints **Expand Arrow** (▸) and click **Compound Constraints**.

Inter-Workflow Navigation

You can navigate to other workflows using the links in the **Jump to...** drop-down list on a workflow page. Select the drop-down list to view and select additional workflows.

When you select a new workflow, properties shared by the rows you select and the constraints you set are used in the new workflow, if they are applicable. If configured constraints or event properties do not map to fields in the new workflow, they are dropped. In addition, compound constraints are not retained when you switch from one workflow to another. In addition, constraints from the captured files workflow only transfer to file and malware event workflows.



Note When you view event counts over a time range, the total number of events may not reflect the number of events for which more detailed data is available. This occurs because the system sometimes prunes older event details to manage disk space usage. To minimize the occurrence of event detail pruning, you can fine-tune event logging to log only those events most important to your deployment.

Note that unless you have either paused the time window or have configured a static time window, the time window changes when you change workflows.

This feature enhances your ability to investigate suspicious activity. For example, if you are viewing connection data and notice that an internal host is transmitting an abnormally large amount of data to an external site, you can select the responder IP address and the port as constraints and then jump to the **Applications** workflow. The applications workflow will use the responder IP address and port as IP Address and Port constraints and display additional information about the application, such as what kind of application it is. You can also click **Hosts** at the top of the page to view the host profile for the remote host.

After finding more information about the application, you can select **Correlation Events** to return to the connection data workflow, remove the Responder IP from the constraints, add the Initiator IP to constraints, and select **Application Details** to see what client the user on the initiating host used when transferring data to the remote host. Note that the Port constraint is not transferred to the Application Details page. While keeping the local host as a constraint, you can also use other navigation buttons to find additional information:

- To discover if any policies have been violated by the local host, keep the IP address as a constraint and select **Correlation Events** from the **Jump to** drop-down list.
- To find out if an intrusion rule triggered against the host, indicating a compromise, select **Intrusion Events** from the **Jump to** drop-down list.
- To view the host profile for the local host and determine if the host is susceptible to any vulnerabilities that may have been exploited, select **Hosts** from the **Jump to** drop-down list.

Bookmarks

Create a bookmark if you want to return quickly to a specific location and time in an event analysis. Bookmarks retain information about:

- the workflow you are using
- the part of the workflow you are viewing
- the page number within the workflow
- any search constraints
- any disabled columns
- the time range you are using

The bookmarks you create are available to all user accounts with bookmark access. This means that if you uncover a set of events that require more in-depth analysis, you can easily create a bookmark and turn over the investigation to another user with the appropriate privileges.



Note If the events that appear in a bookmark are deleted (either directly by a user or by automatic database cleanup), the bookmark no longer displays the original set of events.

Creating Bookmarks

In a multidomain deployment, you can only view bookmarks created in the current domain.

Procedure

Step 1 During an event analysis, with the events of interest displayed, click **Bookmark This Page**.

Step 2 In the **Bookmark Name** field, enter a name.

Step 3 Click **Save Bookmark**.

Viewing Bookmarks

In a multidomain deployment, you can only view bookmarks created in the current domain.

Procedure

From any event view, you have two options:

- Hover your pointer over **View Bookmarks**, and click on the desired bookmark in the drop-down menu.
- Click on click **View Bookmarks** and on the View Bookmarks page, click on the desired bookmark name or **View** () next to it.

Note If the events that originally appeared in a bookmark are deleted (either directly by a user or by automatic database cleanup), the bookmark no longer displays the original set of events.



CHAPTER 108

Searching for Events

The following topics describe how to search for events within a workflow:

- Event Searches, on page 2177
- Query Overrides Via the Shell, on page 2185

Event Searches

The Firepower System generates information that is stored as events in database tables. Events contain multiple fields that describe the activity that caused the appliance to generate the event. You can create and save searches customized for your environment for any of the different event types and save them to reuse later.

When you save a search you give it a name and specify whether the search will be available to you alone or to all users of the appliance. If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search. If you previously saved a search, you can load it, make any necessary modifications, and then start the search. Custom analysis dashboard widgets, report templates, and custom roles can also use saved searches. If you have saved searches, you can delete them from the Search page.

For some event types, the Firepower System provides predefined searches that serve as examples and can provide quick access to important information about your network. You can modify fields within the predefined searches for your network environment, then save the searches to reuse later.

The search criteria you can use depends on the type of search, but the mechanics are the same. Searches return only records that match the search criteria specified for all fields.

Note Searching a custom table requires a slightly different procedure.

Related Topics

Searching Custom Tables, on page 2202

Search Constraints

Each database table has its own search page where you can enter search constraint values to apply to fields defined for the table. Depending on the type of field, special syntax may be used to specify criteria such as wildcard characters or a range of numeric values.

Search results appear on workflow pages displaying each table field in columnar layout. Some database tables can additionally be searched using fields that are not displayed as columns on workflow pages. To determine whether such a constraint applies to your search results when viewing the results on a workflow page, click

Expand Arrow ( ) to view the active search constraints.

General Search Constraints

When searching for events, observe the following general guidelines:

- All fields accept negation (`!`).
- All fields accept comma-separated lists of search values. Records that contain any of the listed values in the specified field match that search criteria.
- All fields accept comma-separated lists enclosed in quotation marks as search values.
  - For fields that may contain only a single value, records with the specified field containing the exact string specified within the quotation marks match the search criteria. For instance, a search for `A, B, "C, D, E"` will match records where the specified field contains `"A"` or `"B"` or `"C, D, E"`. This permits matching on fields that include the comma in possible values.
  - For fields that may contain multiple values at the same time, records with the specified fields containing all of the values in the quote-enclosed comma-separated list match that search criteria.
  - For fields that may contain multiple values at the same time, search criteria may include single values as well as quote-enclosed comma-separated lists. For instance, a search for `A, B, "C, D, E"` on a field that may contain one of more of these letters matches records where the specified field contains `A` or `B`, or all of `C`, `D`, and `E`.
- Specify `n/a` in any field to identify events where information is not available for that field; use `!n/a` to identify the events where that field is populated.
- You can precede many numeric fields with greater than (>), greater than or equal to (>=), less than (**<**), less than or equal to (<=), equal to (=), or not equal to (<>) operators.

Tip When searching a field with long complicated values (such as SHA-256 hash values), you can copy the search criteria value from source material and paste it into the appropriate field on the search page.

Wildcards and Symbols in Searches

Many text fields on search pages allow you to use an asterisk (*) to match characters in a string. For example, specifying `net*` matches `network`, `netware`, `netscape`, and so on.

Note that in text fields that allow a wildcard, you **must** use the wildcard if you want to match a partial string. For example, if you are searching the audit log for all audit records that involve page views (that is, the message is Page View), searching for `Page` returns no results. Instead, specify `Page*`.

In some fields you can search for all or part of the field contents without using asterisks. In these cases, you must use quotation marks around a search string to make exact matches--otherwise, the system performs a partial match. For example, if you were to search such a field for the string `Scan Completed with Detection` without using quotation marks, the system would return records where the field contains the following strings as well as those where the field exactly matches the search string:

```
Scan Completed, No Detections
Scan completed With Detections
```

If you want to search for non-alphanumeric characters (including the asterisk character), enclose the search string in quotation marks. For example, to search for the string:

```
Find an asterisk (*)
```

enter:

```
"Find an asterisk (*)"
```

Objects and Application Filters in Searches

The Firepower System allows you to create named objects, object groups, and application filters that can be used as part of your network configuration. You can use these objects, groups, and filters as search criteria when performing or saving searches.

When you perform a search, objects, object groups, and application filters appear in the format, `${object_name}`. For example, a network object with the object name `ten_ten_network` appears as `${ten_ten_network}` in a search.

You can click **Object** (+) that appears next to a search field where you can use an object as a search criterion.

Related Topics

The Object Manager, on page 396

Time Constraints in Searches

The formats accepted by search criteria fields that take a time value are shown in the following table.

Table 281: Time Specification in Search Fields

Time Formats	Example
`today [at HH:MMam|pm]`	`today` `today at 12:45pm`
`YYYY-DDMM- HH:MM:SS`	`2006-03-22 14:22:59`

You can precede a time value with one of the following operators:

Table 282: Time Specification Operators

Operator	Example	Explanation
`<`	`< 2006-03-22 14:22:59`	Returns events with a timestamp before 2:23 PM, March 22, 2006.
`>`	`> today at 2:45pm`	Returns events with a timestamp later than today at 2:45 PM.

IP Addresses in Searches

When specifying IP addresses in searches, you can enter an individual IP address, a comma-separated list of addresses, an address block, or a range of IP addresses separated with a hyphen (-). You can also use negation.

For searches that support IPv6 (such as intrusion event, connection data, and correlation event searches) you can enter IPv4 and IPv6 addresses and CIDR/prefix length address blocks in any combination. When you search for hosts by IP address, the results include all hosts for which at least one IP address matches your search conditions, that is, a search for an IPv6 address may return hosts whose primary address in IPv4.

When you use CIDR or prefix length notation to specify a block of IP addresses, the Firepower System uses **only** the portion of the network IP address specified by the mask or prefix length. For example, if you type `10.1.2.3/8`, the Firepower System uses `10.0.0.0/8`.

Because IP addresses can be represented by network objects, you can also click the add network **Object** (+) that appears next to an IP address search field to use a network object as an IP address search criterion.

Table 283: Acceptable IP Address Syntax

To specify...	Type...	For example...
a single IP address	the IP address.	`192.168.1.1` `2001:db8::abcd`
multiple IP addresses using a list	a comma-separated list of IP addresses. Do **not** add a space before or after the commas.	`192.168.1.1,192.168.1.2` `2001:db8::b3ff,2001:db8::0202`
a range of IP addresses that can be specified with a CIDR block or prefix length	the IP address block in IPv4 CIDR or IPv6 prefix length notation.	`192.168.1.0/24` This specifies any IP in the 192.168.1.0 network with a subnet mask of 255.255.255.0, that is, 192.168.1.0 through 192.168.1.255.
a range of IP addresses that cannot be specified with a CIDR block or prefix	the IP address range using a hyphen. Do **not** add a space before or after the hyphen.	`192.168.1.1-192.168.1.5` `2001:db8::0202-2001:db8::8329`

negation of any of the other ways to specify IP addresses or ranges of IP addresses	an exclamation point in front of the IP address, block, or range.	`192.168.0.0/32,!192.168.1.10`
		`!2001:db8::/32`
		`!192.168.1.10,!2001:db8::/32`

hosts that are blocked or monitored (but would have been blocked) See Host Profile Icons, on page 2156.	In connection and Security Intelligence events, in Initiator IP and Responder IP fields: • blacklist • monitor	--

Related Topics

Firepower System IP Address Conventions, on page 15

Managed Devices in Searches

If you group devices—whether just on the FMC, or as actual high availability or scalability configurations—searching for the name for the group correctly returns results for all devices in the group.

If the system finds a match for a group, device high-availability pair, or stack, it replaces the group, device high-availability pair, or stack name with the appropriate member device names for the purpose of performing the search. When you save a search that uses a device group, device high-availability pair, or stack in the device field the system saves the name specified in the device field and performs the device name replacement again each time the search is executed.

Ports in Searches

The Firepower System accepts specific syntax for port numbers in searches. You can enter:

- a single port number
- a comma-separated list of port numbers
- two port numbers separated by a dash to represent a range of port numbers
- a port number followed by a protocol abbreviation, separated by a forward slash (only when searching for intrusion events)
- a port number or range of port numbers preceded by an exclamation mark to indicate a negation of the specified ports

Note Do **not** use spaces when specifying port numbers or ranges.

Table 284: Port Syntax Examples

Example	Description
`21`	Returns all events on port 21, including TCP and UDP events.
`!23`	Returns all events except those on port 23.
`25/tcp`	Returns all TCP-related intrusion events on port 25.
`21/tcp,25/tcp`	Returns all TCP-related intrusion events on ports 21 and 25.
`21-25`	Returns all events on ports 21 through 25.

Event Fields in Searches

When searching for events, you can use the following fields as search criteria:

- Audit Log Workflow Fields, on page 2418
- Application Data Fields, on page 2382
- Application Detail Data Fields, on page 2384
- Captured File Fields, on page 2314

- White List Event Fields, on page 2408
- Connection and Security Intelligence Event Fields, on page 2225
- Correlation Event Fields, on page 2404
- Discovery Event Fields, on page 2364
- The Health Events Table, on page 298
- Host Attribute Data Fields, on page 2372
- Host Data Fields, on page 2366
- File and Malware Event Fields, on page 2300
- Intrusion Event Fields, on page 2254
- Fields in an Intrusion Rule Update Log, on page 150
- Remediation Status Table Fields, on page 2412
- Nmap Scan Results Fields, on page 1871
- Server Data Fields, on page 2379
- Third-Party Vulnerability Data Fields, on page 2390
- User-Related Fields, on page 2392
- Vulnerability Data Fields, on page 2386
- White List Violation Fields, on page 2410

Performing a Search

You must have Admin or Security Analyst privileges to perform a search.

Procedure

Step 1 Select **Analysis** > **Search**.

Tip You may also click **Search** from any page on a workflow.

Step 2 From the table drop-down list, select the type of event or data to search.

Step 3 Enter your search criteria in the appropriate fields. See the following sections for detailed information on the search criteria you can use:

- Search Constraints, on page 2177
- Audit Log Workflow Fields, on page 2418
- Application Data Fields, on page 2382
- Application Detail Data Fields, on page 2384
- Captured File Fields, on page 2314

- White List Event Fields, on page 2408
- Connection and Security Intelligence Event Fields, on page 2225
- Correlation Event Fields, on page 2404
- Discovery Event Fields, on page 2364
- The Health Events Table, on page 298
- Host Attribute Data Fields, on page 2372
- Host Data Fields, on page 2366
- File and Malware Event Fields, on page 2300
- Intrusion Event Fields, on page 2254
- Fields in an Intrusion Rule Update Log, on page 150
- Remediation Status Table Fields, on page 2412
- Nmap Scan Results Fields, on page 1871
- Server Data Fields, on page 2379
- Third-Party Vulnerability Data Fields, on page 2390
- User Data Fields
- User Activity Data Fields
- Vulnerability Data Fields, on page 2386
- White List Violation Fields, on page 2410

Step 4 If you want to use the search again in the future, save the search as described in Saving a Search, on page 2183.

Step 5 Click **Search** to start the search. Your search results appear in the default workflow for the table you are searching, constrained by time (if applicable).

What to do next

- To analyze the search results using workflows, see Using Workflows, on page 2150.

Related Topics

Configuring Event View Settings, on page 33

Saving a Search

You must have Admin or Security Analyst privileges to save a search.

In a multidomain deployment, the system displays saved searches created in the current domain, which you can edit. It also displays searches saved in ancestor domains, which you cannot edit. To view and edit searches created in a lower domain, switch to that domain.

Before you begin

- Establish search criteria as described in Performing a Search, on page 2182, or load a saved search as described in Loading a Saved Search, on page 2184.

Procedure

Step 1 From the Search page, if you want to save the search as private so only you can access it, check the **Private** checkbox.

Tip If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search.

Step 2 You have two options:

- If you want to save a new version of a loaded search, click **Save As New**.
- If you want to save a new search, or overwrite a custom search using the same name, click **Save**. If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Loading a Saved Search

You must have Admin or Security Analyst privileges to load a saved search.

In a multidomain deployment, the system displays saved searches created in the current domain, which you can edit. It also displays searches saved in ancestor domains, which you cannot edit. To view and edit searches created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Analysis** > **Search**.

Tip You may also click **Search** from any page on a workflow.

Step 2 From the table drop-down list, choose the type of event or data to search.

Step 3 Choose the search you want to load from the **Custom Searches** list or the **Predefined Searches** list.

Step 4 If you want to use different search criteria, change the search constraints.

Step 5 If you want to use a changed search again in the future, save the search as described in Saving a Search, on page 2183.

Step 6 Click **Search**.

Query Overrides Via the Shell

System administrators can use a Linux shell-based query management tool to locate and stop long-running queries.

The query management tool allows you to locate queries running longer than a specified number of minutes and stop those queries. The tool logs an event to the audit log and to syslog when you stop a query.

Note that the `admin` internal user can access the FMC CLI. If you use an external authentication object which grants CLI access, users matching the shell access filter can also log into the CLI.



Note Leaving the search page in the web interface does not stop a query. Queries that take a long time to return results impact overall system performance while the query is running.

Shell-Based Query Management Syntax

Use the following syntax to manage long-running queries:

```
query_manager [-v] [-l [minutes]] [-k query_id [...]] [--kill-all minutes]
```

Table 285: `query_manager` Options

Option	Description
`-h, --help`	Prints a brief help message.
`-l, --list [minutes]`	Lists all queries taking longer than passed-in minutes. By default it will show all queries taking longer than 1 minute.
`-k, --kill query_id [...]`	Kills the query with the passed-in id. The option can take multiple ids.
`--kill-all minutes`	Kills all queries taking longer than passed-in minutes.
`-v, --verbose`	Verbose output including full SQL queries.



Caution For system security reasons, Cisco strongly recommends that you not establish additional Linux shell users on any appliance.

Stopping Long-Running Queries

You must be the **admin** user or externally authenticated user with shell access

Procedure

Step 1 Connect to the Firepower Management Center via `ssh`.

Step 2 Run `query_manager` under `sudo` using the syntax described in Shell-Based Query Management Syntax, on page 2185.



CHAPTER 109

Custom Workflows

The following topics describe how to use custom workflows:

- Introduction to Custom Workflows, on page 2187
- Saved Custom Workflows, on page 2187
- Custom Workflow Creation, on page 2188
- Custom Workflow Use and Management, on page 2191

Introduction to Custom Workflows

If the predefined and Cisco-provided custom workflows do not meet your needs, you can create and manage custom workflows.

Custom workflows are workflows that you create to meet the unique needs of your organization. When you create a custom workflow, you choose the kind of event (or database table) on which the workflow is based. On the Firepower Management Center, you can base a custom workflow on a custom table. You can also choose the pages a custom workflow contains; custom workflows can contain drill-down, table view, and host or packet view pages.

If your event evaluation process changes, you can edit custom workflows to meet your new needs. Note that you cannot edit any of the predefined workflows.

Tip You can set a custom workflow as the default workflow for any event type.

Saved Custom Workflows

In addition to predefined workflows, which cannot be modified, the Firepower Management Center includes several saved custom workflows. Each of these workflows is based on a custom table and can be modified.

In a multidomain deployment, these saved workflows belong to the Global domain and cannot be modified in lower domains.

Table 286: Saved Custom Workflows

Workflow Name	Description
Events by Impact, Priority, and Host Criticality	You can use this workflow to quickly pick out and focus in on hosts that are important to your network, currently vulnerable, and possibly currently under attack. This workflow is based on the Intrusion Events with Destination Criticality custom table.
Events by Priority and Classification	This workflow lists events and their type in order of event priority, along with a count showing how many times each event has occurred. This workflow is based on the Intrusion Events custom table.
Events with Destination, Impact, and Host Criticality	You can use this workflow to find the most recent attacks on hosts that are important to your network and currently vulnerable. This workflow is based on the Intrusion Events with Destination Criticality custom table.
Hosts with Servers Default Workflow	You can use this workflow to quickly view the basic information in the Hosts with Servers custom table. This workflow is based on the Hosts with Servers custom table.
Intrusion Events with Destination Criticality Default Workflow	You can use this workflow to quickly view the basic information in the Intrusion Events with Destination Criticality custom table. This workflow is based on the Intrusion Events with Destination Criticality custom table.
Intrusion Events with Source Criticality Default Workflow	You can use this workflow to quickly view the basic information in the Intrusion Events with Source Criticality custom table. This workflow is based on the Intrusion Events with Source Criticality custom table.
Server and Host Details	You can use this workflow to determine what servers are most frequently used on your network and which hosts are running those servers. This workflow is based on the Hosts with Servers custom table.

Custom Workflow Creation

If the predefined and Cisco-provided custom workflows do not meet your needs, you can create custom workflows.



Tip Instead of creating a new custom workflow, you can export a custom workflow from another appliance and then import it onto your appliance. You can then edit the imported workflow to suit your needs.

When you create a custom workflow, you:

- Select a table to be the source of the workflow
- Provide a workflow name
- Add drill-down pages and table view pages to the workflow

For each drill-down page in the workflow, you can:

- Provide a name that appears at the top of the page in the web interface
- Include up to five columns per page
- Specify a default sort order, ascending or descending

You can add table view pages in any position in the sequence of workflow pages. They do not have any editable properties, such as a page name, sort order, or user-definable column positions.

Note You must add at least one drill-down page or a table view of events to a custom workflow.

Note If you selected **Vulnerabilities** as the table type, then add **IP Address** as a table column, the IP Address column does not appear when you are viewing vulnerabilities using your custom workflow, unless you use the search feature to constrain the workflow to view a specific IP address or block of addresses.

The final page of a custom workflow depends on the table on which you base the workflow, as described in the following table. These final pages are added by default when you create the workflow.

Table 287: Custom Workflow Final Pages

Event/Asset Type	Final Page
Discovery events	Hosts
Vulnerabilities	Vulnerability detail
Third-party vulnerabilities	Hosts
Users	Users
Indications of compromise	Hosts or users
Intrusion events	Packets

The system does not add a final page to custom workflows based on other kinds of events (for example, audit log or malware events).

Custom workflows based on connection data are like other custom workflows, except you can include drill-down pages containing connection summary data, and connection data graph pages as well as drill-down pages containing data for individual connections and table view pages.

Creating Custom Workflows Based on Non-Connection Data

You must have Admin or Security Analyst privileges to create a custom workflow based on non-connetion data.

Procedure

Step 1 Choose **Analysis** > **Custom** > **Custom Workflows**.

Step 2 Click **Create Custom Workflow**.

Step 3 Enter a name for the workflow in the **Name** field.

Step 4 Optionally, enter a **Description**.

Step 5 Choose the table you want to include from the **Table** drop-down list.

Step 6 If you want to add one or more drill-down pages to the workflow, click **Add Page**.

Step 7 Enter a name for the page in the **Page Name** field.

Step 8 Under Column 1, choose a sort priority and a table column. This column will appear in the leftmost column of the page.

Example:

For example, to create a page showing the destination ports that are targeted, and to sort the page by count, choose **2** from the **Sort Priority** drop-down list and **Destination Port/ICMP Code** from the **Field** drop-down list.

Step 9 Continue choosing fields to include and setting their sort priority until you have specified all the fields you want to appear on the page.

Step 10 If you want to add a table view page to the workflow, click **Add Table View**.

Step 11 Click **Save**.

Creating Custom Connection Data Workflows

Custom workflows based on connection data are like other custom workflows, except you can include connection data graph pages as well as drill-down pages and table view pages. You can include as many of each type of page in the workflow as you want, in any order. Each connection data graph page contains a single graph, which can be a line graph, bar graph, or pie chart. On line and bar graphs, you may include more than one dataset.

You must have Admin privileges to create a custom workflow based on connection data.

Procedure

Step 1 Choose **Analysis** > **Custom** > **Custom Workflows**.

Step 2 Click **Create Custom Workflow**.

Step 3 Enter a name for the workflow in the **Name** field.

Step 4 Optionally, enter a **Description**.

Step 5 From the **Table** drop-down list, choose **Connection Events**.

Step 6 If you want to add one or more drill-down pages to the workflow, you have two options:

- Click **Add Page** to add a drill-down page that contains data on individual connections,
- Click **Add Summary Page** to add a drill-down page that contains connection summary data.

Step 7 Enter a name for the page in the **Page Name** field.

Step 8 Under **Column 1**, choose a sort priority and a table column. This column will appear in the leftmost column of the page.

Step 9 Continue choosing fields to include and setting their sort priority until you have specified all the fields you want to appear on the page.

Example:

For example, to create a page showing the amount of traffic transmitted over your monitored network and to sort the page by the responders that transmitted the most traffic, choose **1** from the **Sort Priority** drop-down list and **Responder Bytes** from the **Field** drop-down list.

Step 10 If you want to add one or more graph pages to the workflow, click **Add Graph**.

Step 11 Enter a name for the page in the **Graph Name** field.

Step 12 Choose the type of graph you want to include on the page:

- line graph (**Line chart** ( ))
- bar graph(**Bar chart** ( ))
- pie chart (**Pie chart** ( ))

Step 13 Specify what kind of data you want to graph by choosing the x- and y-axes of the graph.

On a pie chart, the x-axis represents the independent variable and the y-axis represents the dependent variable.

Step 14 Choose the datasets you want to include on the graph.

Note that pie charts can include only one data set.

Step 15 If you want to add a table view of connection data, click **Add Table View**.

Table views are not configurable.

Step 16 Click **Save**.

Custom Workflow Use and Management

The method you use to view a workflow depends on whether the workflow is based on one of the predefined event tables or on a custom table.

If your custom workflow is based on a predefined event table, access it in the same way that you would access a workflow that ships with the appliance. For example, to access a custom workflow based on the Hosts table, choose **Analysis** > **Hosts** > **Hosts**. If, on the other hand, your custom workflow is based on a custom table, you must access it from the Custom Tables page.

If your event evaluation process changes, you can edit custom workflows to meet your new needs. Note that you cannot edit any of the predefined workflows.

Tip You can set a custom workflow as the default workflow for any event type.

Viewing Custom Workflows Based on Predefined Tables

You must have Admin, Maintenance, or Security Analyst privileges to view a custom workflow.

Procedure

Step 1 Choose the appropriate menu path and option for the table on which you based your custom workflow, as described in the Workflow Selection, on page 2152.

Step 2 To use a different workflow, including a custom workflow, click **(switch workflow)** next to the current workflow title.

Step 3 If no events appear and the workflow can be constrained by time, you may need to adjust the time range; see Event Time Constraints, on page 2165.

Viewing Custom Workflows Based on Custom Tables

You must have Admin or Security Analyst privileges to view a custom workflow that is based on custom tables.

In a multidomain deployment, the system displays custom workflows created in the current domain, which you can edit. It also displays custom workflows created in ancestor domains, which you cannot edit. To view and edit custom workflows in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Analysis** > **Custom** > **Custom Tables**.

Step 2 Click **View** () next to the custom table you want to view, or click the name of the custom table.

Step 3 To use a different workflow, including a custom workflow, click **(switch workflow)** beside the current workflow title.

Step 4 If no events appear and the workflow can be constrained by time, you may need to adjust the time range; see Event Time Constraints, on page 2165.

Editing Custom Workflows

You must have Admin or Security Analyst privileges to edit a custom workflow.

In a multidomain deployment, the system displays custom workflows created in the current domain, which you can edit. It also displays custom workflows created in ancestor domains, which you cannot edit. To view and edit custom workflows in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Analysis** > **Custom** > **Custom Workflows**.

Step 2 Click **Edit** ( ) next to the name of the workflow that you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Make any changes that you want to the workflow.

Step 4 Click **Save**.



CHAPTER 110

Custom Tables

The following topics describe how to use custom tables:

- Introduction to Custom Tables, on page 2195
- Predefined Custom Tables, on page 2195
- User-Defined Custom Tables, on page 2200
- Searching Custom Tables, on page 2202

Introduction to Custom Tables

As the Firepower System collects information about your network, the Firepower Management Center stores it in a series of database tables. When you use a workflow to view the resulting information, the Firepower Management Center pulls the data from one of these tables. For example, the columns on each page of the Network Applications by Count workflow are taken from the fields in the Applications table.

If you determine that your analysis of the activity on your network would be enhanced by combining fields from different tables, you can create a custom table. For example, you could combine the host criticality information from the predefined Host Attributes table with the fields from the predefined Connection Data table and then examine connection data in a new context.

Note that you can create custom workflows for either predefined or custom tables.

Predefined Custom Tables

Custom tables contain fields from two or more predefined tables. The Firepower System is delivered with a number of system-defined custom tables, but you can create additional custom tables that contain only information that matches your specific needs.

For example, the Firepower System is delivered with system-defined custom tables that correlate intrusion event data with host data, so you can search for events that impact critical systems and view the results of that search in one workflow.

In a multidomain deployment, the predefined custom tables belong to the Global domain and cannot be modified in lower domains.

The following table describes the custom tables provided with the system.

Table 288: System-Defined Custom Tables

Table	Description
Hosts with Servers	Includes fields from the Hosts and Servers tables, providing you with information about the detected applications running on your network, as well as basic operating system information about the hosts running those applications.
Intrusion Events with Destination Criticality	Includes fields from the Intrusion Events table and the Hosts table, providing you with information on the intrusion events, as well as the host criticality of the destination host involved in each intrusion event. You can use this table to search for intrusion events involving destination hosts with high host criticality.
Intrusion Events with Source Criticality	Includes fields from the Intrusion Events table and the Hosts table, providing you with information on the intrusion events and the host criticality of the source host involved in each intrusion event. You can use this table to search for intrusion events involving source hosts with high host criticality.

Possible Table Combinations

When you create a custom table, you can combine fields from predefined tables that have related data. The following table lists the predefined tables you can combine to create a new custom table. Keep in mind that you can create a custom table that combines fields from more than two predefined custom tables.

Table 289: Custom Table Combinations

You can combine fields from...	With fields from...
Applications	• Correlation Events • Intrusion Events • Connection Summary Data • Host Attributes • Application Details • Discovery Events • Connection Events • Hosts • Servers • White List Events

You can combine fields from...	With fields from...
Correlation Events	• Applications • Host Attributes • Hosts
Intrusion Events	• Applications • Host Attributes • Hosts • Servers
Connection Summary Data	• Applications • Host Attributes • Hosts • Servers
Host Indications of Compromise	• Applications • Application Details • Captured Files • Connection Events • Connection Summary Data • Correlation Events • Discovery Events • Host Attributes • Hosts • Intrusion Events • Security Intelligence Events • Servers • White List Events

You can combine fields from...	With fields from...
Host Attributes	• Applications • Correlation Events • Intrusion Events • Connection Summary Data • Application Details • Discovery Events • Connection Events • Hosts • Servers • White List Events
Application Details	• Applications • Host Attributes • Hosts
Discovery Events	• Applications • Host Attributes • Hosts
Connection Events	• Applications • Host Attributes • Hosts • Servers
Security Intelligence Events	• Applications • Host Attributes • Hosts • Servers

You can combine fields from...	With fields from...
Hosts	• Applications • Correlation Events • Intrusion Events • Connection Summary Data • Host Attributes • Application Details • Discovery Events • Connection Events • Servers • White List Events
Servers	• Applications • Intrusion Events • Connection Summary Data • Host Attributes • Connection Events • Hosts
White List Events	• Applications • Host Attributes • Hosts

Sometimes a field in one table maps to more than one field in another table. For example, the predefined **Intrusion Events with Destination Criticality** custom table combines fields from the Intrusion Events table and the Hosts table. Each event in the Intrusion Events table has two IP addresses associated with it—a source IP address and a destination IP address. However, the "events" in the Hosts table each represent a single host IP address (hosts may have multiple IP addresses). Therefore, when you create a custom table based on the Intrusion Events table and the Hosts table, you must choose whether the data you display from the Hosts table applies to the host source IP address or the host destination IP address in the Intrusion Events table.

When you create a new custom table, a default workflow that displays all the columns in the table is automatically created. Also, just as with predefined tables, you can search custom tables for data that you want to use in your network analysis. You can also generate reports based on custom tables, as you can with predefined tables.

User-Defined Custom Tables

Tip Instead of creating a new custom table, you can export a custom table from another Firepower Management Center, then import it onto your Firepower Management Center.

To create a custom table, decide which predefined tables delivered with the Firepower System contain the fields you want to include in your custom table. You can then choose which fields you want to include and, if necessary, configure field mappings for any common fields.

Tip Data involving the Hosts table allows you to view data associated with all IP addresses from one host, rather than one specific IP address.

For example, consider a custom table that combines fields from the Correlation Events table and the Hosts table. You can use this custom table to get detailed information about the hosts involved in violations of any of your correlation policies. Note that you must decide whether to display data from the Hosts table that matches the source IP address or the destination IP address in the Correlation Events table.

If you view the table view of events for this custom table, it displays correlation events, one per row. You can configure the custom table to include the following information:

- the date and time the event was generated
- the name of the correlation policy that was violated
- the name of the rule that triggered the violation
- the IP address associated with the source, or initiating, host involved in the correlation event
- the source host's NetBIOS name
- the operating system and version the source host is running
- the source host criticality

Tip You could create a similar custom table that displays the same information for destination, or responding, hosts.

Creating a Custom Table

Procedure

Step 1 Choose **Analysis** > **Custom** > **Custom Tables**.

Step 2 Click **Create Custom Table**.

Step 3 In the **Name** field, enter a name for the custom table.

Example:

For example, you might enter `Correlation Events with Host Information (Src IP).`

Step 4 From the **Tables** drop-down list, choose **Correlation Events**.

Step 5 Under **Fields**, choose**Time** and click **Add** to add the date and time when a correlation event was generated.

Step 6 Repeat step 5 to add the **Policy** and **Rule** fields.

Tip You can use Ctrl or Shift while clicking to choose multiple fields. You can also click and drag to choose multiple adjacent values. However, if you want to specify the order the fields appear in the table view of events associated with the table, add the fields one at a time.

Step 7 From the **Tables** drop-down list, choose **Hosts**.

Step 8 Add the **IP Address**, **NetBIOS Name**, **OS Name**, **OS Version**, and **Host Criticality** fields to the custom table.

Step 9 Under **Common Fields**, next to **Correlation Events**, choose **Source IP**.

Your custom table is configured to display the host information you chose in step 8 for the source, or initiating, hosts involved in correlation events.

Tip You can create a custom table that displays detailed host information for the destination, or responding, hosts involved in a correlation event by following this procedure but choosing **Destination IP** instead of **Source IP**.

Step 10 Click **Save**.

Modifying a Custom Table

In a multidomain deployment, the system displays custom tables created in the current domain, which you can edit. It also displays custom tables created in ancestor domains, which you cannot edit. To view and edit custom tables in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Analysis** > **Custom** > **Custom Tables**.

Step 2 Click **Edit** ( ) next to the table you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Optionally, remove fields from the table by clicking **Delete** ( ) next to the fields you want to remove.

Note If you delete fields currently in use in reports, the system will prompt you to confirm that you want to remove the sections using those fields from those reports.

Step 4 Make other changes as needed.

Step 5 Click **Save**.

Deleting a Custom Table

In a multidomain deployment, the system displays custom tables created in the current domain, which you can delete. It also displays custom tables created in ancestor domains, which you cannot delete. To delete custom tables in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Analysis** > **Custom** > **Custom Tables**.

Step 2 Click **Delete** ( ) next to the custom table you want to delete.

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Viewing a Workflow Based on a Custom Table

When you create a custom table, the system automatically creates a default workflow for it. The first page of this workflow displays a table view of events. If you include intrusion events in your custom table, the second page of the workflow is the packet view. Otherwise, the second page of the workflow is a hosts page. You can also create your own custom workflows based on your custom table.

Tip If you create a custom workflow based on a custom table, you can specify it as the default workflow for that table.

You can use the same techniques to view events in your custom table that you use for event views based on predefined tables.

In a multidomain deployment, the system displays custom tables created in the current domain, which you can edit. It also displays custom tables created in ancestor domains, which you cannot edit. To view and edit custom tables in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Analysis** > **Custom** > **Custom Tables**.

Step 2 Click **View** ( ) next to the custom table related to the workflow you want to see.

Searching Custom Tables

In a multidomain deployment, the system displays custom tables created in the current domain, which you can edit. It also displays custom tables created in ancestor domains, which you cannot edit. To view and edit custom tables in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Analysis** > **Custom** > **Custom Tables**.

Step 2 Click **View** ( ) next to the custom table you want to search.

Tip To use a different workflow, including a custom workflow, click **(switch workflow)** next the workflow title.

Step 3 Click **Search**.

Tip To search the database for a different kind of event or data, choose it from the table drop-down list.

Step 4 Enter your search criteria in the appropriate fields.

If you enter criteria for multiple fields, the search returns only the records that match search criteria specified for all fields.

Tip Click **Object** ( ) next to a search field to use an object as a search criterion.

Step 5 Optionally, if you plan to save the search, you can check the **Private** check box to save the search as private so only you can access it. Otherwise, leave the check box clear to save the search for all users.

Tip If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search.

Step 6 Optionally, you can save the search to be used again in the future. You have the following options:

- Click **Save** to save the search criteria. The search is visible only to your account if you checked the **Private** check box.
- Click **Save As New** to save a new search or assign a name to a search you created by altering a previously-saved search. The search is saved and visible only to your account if you checked the **Private** check box.

Step 7 Click **Search** to start the search.

Your search results appear in the default workflow for the custom table, constrained by the current time range (if applicable).



PART XXV

Events and Assets

- Connection Logging, on page 2207
- Connection and Security Intelligence Events, on page 2223
- Working with Intrusion Events, on page 2251
- File/Malware Events and Network File Trajectory, on page 2295
- Using Host Profiles, on page 2325
- Working with Discovery Events, on page 2351
- Correlation and Compliance Events, on page 2403
- Auditing the System, on page 2415



CHAPTER 111

Connection Logging

The following topics describe how to configure the Firepower System to log connections made by hosts on your monitored network:

- About Connection Logging, on page 2207
- Limitations of Connection Logging, on page 2215
- Best Practices for Connection Logging, on page 2216
- Requirements and Prerequisites for Connection Logging, on page 2218
- Configure Connection Logging, on page 2218

About Connection Logging

The system can generate logs of the connections its managed devices detect. These logs are called *connection events*. Settings in rules and policies give you granular control over which connections you log, when you log them, and where you store the data. Special connection events, called *Security Intelligence events*, represent connections that were blocked by the reputation-based Security Intelligence feature.

Connection events contain data about the detected sessions. The information available for any individual connection event depends on several factors, but in general includes:

- Basic connection properties: timestamp, source and destination IP address, ingress and egress zones, the device that handled the connection, and so on
- Additional connection properties discovered or inferred by the system: applications, requested URLs, or users associated with the connection, and so on
- Metadata about why the connection was logged: which configuration handled the traffic, whether the connection was allowed or blocked, details about encrypted and decrypted connections, and so on

Log connections according to the security and compliance needs of your organization. When setting up connection logging, keep in mind that the system can log a connection for multiple reasons, and that disabling logging in one place does not mean that matching connections will not be logged.

The information in a connection event depends on several factors, including traffic characteristics, the configuration that ultimately handled the connection, and so on.

Note You can supplement the connection logs gathered by your managed devices with connection data generated from exported NetFlow records. This is especially useful if you have NetFlow-enabled routers or other devices deployed on networks that your Firepower System managed devices cannot monitor.

Related Topics

Netflow Data in the Firepower System, on page 1823

Connections That Are Always Logged

Unless you disable connection event storage, the system automatically saves the following end-of-connection events to the Firepower Management Center database, regardless of any other logging configurations.

Connections Associated with Intrusions

The system automatically logs connections associated with intrusion events, unless the connection is handled by the access control policy's default action.

When an intrusion policy associated with the access control default action generates an intrusion event, the system does *not* automatically log the end of the associated connection. Instead, you must explicitly enable default action connection logging. This is useful for intrusion prevention-only deployments where you do not want to log any connection data.

However, if you enable beginning-of-connection logging for the default action, the system *does* log the end of the connection when an associated intrusion policy triggers, in addition to logging the beginning of the connection.

Connections Associated with File and Malware Events

The system automatically logs connections associated with file and malware events.

Note File events generated by inspecting NetBIOS-SSN (SMB) traffic do not immediately generate connection events because the client and server establish a persistent connection. The system generates connection events after the client or server ends the session.

Connections Associated with Intelligent Application Bypass

The system automatically logs bypassed and would-have-bypassed connections associated with IAB.

Monitored Connections

The system always logs the ends of connections for monitored traffic, even if the traffic matches no other rules and you do not enable default action logging. For more information, see Logging for Monitored Connections, on page 2210.

Other Connections You Can Log

So that you log only critical connections, enable connection logging on a per-rule basis. If you enable connection logging for a rule, the system logs all connections handled by that rule.

You can also log connections handled by policy default actions. Depending on the rule or default action (and for access control, a rule's inspection configuration), your logging options differ.

Prefilter Policy: Rules and Default Action

You can log connections (including entire plaintext, passthrough tunnels) that you fastpath or block with a prefilter policy.

Prefiltering uses outer-header criteria to handle traffic. For tunnels that you log, the resulting connection events contain information from the outer, encapsulation headers.

For traffic subject to further analysis, logging in the prefilter policy is disabled, although matching connections may still be logged by other configurations. The system performs all further analysis using inner headers, that is, the system independently handles and logs each connection within an allowed tunnel.

SSL Policy: Rules and Default Action

You can log connections that match an SSL rule or SSL policy default action.

For blocked connections, the system immediately ends the session and generates an event. For monitored connections and connections that you pass to access control rules, the system generates an event when the session ends.

Access Control Policy: Security Intelligence Decisions

You can log a connection whenever it is blocked by the reputation-based Security Intelligence feature.

Optionally, and recommended in passive deployments, you can use a monitor-only setting for Security Intelligence filtering. This allows the system to further analyze connections that would have been blocked by Security Intelligence, but still log the match. Security Intelligence monitoring also allows you to create traffic profiles using Security Intelligence information.

When the system logs a connection event as the result of Security Intelligence filtering, it also logs a matching Security Intelligence event, which is a special kind of connection event that you can view and analyze separately, and that is also stored and pruned separately. So that you can identify the matching IP address in the connection, host icons beside blocked and monitored IP addresses look slightly different in the tables on the pages under the **Analysis** > **Connections** menus.

Access Control Policy: Rules and Default Action

You can log connections that match an access control rule or access control policy default action.

Related Topics

How Rules and Policy Actions Affect Logging, on page 2209

How Rules and Policy Actions Affect Logging

Connection events contain metadata about why the connection was logged, including which configurations handled the traffic. Where you can configure connection logging, rule actions, and policy default actions

determine not only how the system inspects and handles matching traffic, but also when and how you can log details about matching traffic.

Related Topics

Tunnel and Prefilter Rule Components, on page 1281
TLS/SSL Rule Actions, on page 1359
Access Control Rule Actions, on page 1223
Connection and Security Intelligence Event Fields, on page 2225

Logging for Fastpathed Connections

You can log fastpathed connections and non-encrypted tunnels, which includes traffic matching the following rules and actions in the prefilter policy:

- Tunnel rules—**Fastpath** action (logs the outer session)
- Prefilter rules—**Fastpath** action

Fastpathed traffic bypasses the rest of access control and QoS, so connection events for fastpathed connections contain limited information. You cannot log connections fastpathed with 8000 Series fastpath rules.

Logging for Monitored Connections

The system always logs the ends of connections for traffic matching the following configurations, even if the traffic matches no other rules and you do not enable default action logging:

- Security Intelligence—Block lists set to monitor (also generates a Security Intelligence event)
- SSL rules—**Monitor** action
- Access control rules—**Monitor** action

The system does not generate a separate event each time a single connection matches a Monitor rule. Because a single connection can match multiple Monitor rules, each connection event can include and display information on the first eight Monitor access control rules that the connection matches, as well as the first matching SSL Monitor rule.

Similarly, if you send connection events to an external syslog or SNMP trap server, the system does not send a separate alert each time a single connection matches a Monitor rule. Rather, the alert that the system sends at the end of the connection contains information on the Monitor rules the connection matched.

Logging for Trusted Connections

You can log the beginnings and ends of trusted connections, which includes traffic matching the following rules and actions:

- Access control rules—**Trust** action
- Access control default action—**Trust All Traffic**



Note Although you *can* log trusted connections, we recommend you do not do so because trusted connections are not subject to deep inspection or discovery, so connection events for trusted connections contain limited information.

The system logs TCP connections handled by a Trust access control rule differently depending on the device that detected the connection:

- For 7000 and 8000 Series devices, TCP connections detected by a Trust rule on the first packet generate different events depending on the presence of a preceding enabled Monitor rule. If the Monitor rule is active, the system evaluates the packet and generates both a beginning and end-of-connection event. If no Monitor rule is active, the system generates only an end-of-connection event.
- For all other models, TCP connections detected by a Trust rule on the first packet generate only an end-of-connection event. The system generates the event one hour after the final session packet.

Logging for Blocked Connections

You can log blocked connections, which includes traffic matching the following rules and actions:

- Tunnel rules—**Block**
- Prefilter rules—**Block**
- Prefilter default action—**Block all tunnel traffic**
- Security Intelligence—Block lists not set to Monitor (also generates a Security Intelligence event)
- SSL rules—**Block** and **Block with reset**
- SSL default action—**Block** and **Block with reset**
- Access control rules—**Block**, **Block with reset**, and **Interactive Block**
- Access control default action—**Block All Traffic**

Only devices deployed inline (that is, using routed, switched, or transparent interfaces, or inline interface pairs) can block traffic. Because blocked connections are not actually blocked in passive deployments, the system may report multiple beginning-of-connection events for each blocked connection.



Caution Logging blocked TCP connections during a Denial of Service (DoS) attack can affect system performance and overwhelm the database with multiple similar events. Before you enable logging for an Block rule, consider whether the rule monitors traffic on an Internet-facing interface or other interface vulnerable to DoS attack.

Beginning vs End-of-Connection Logging for Blocked Connections

When you log a blocked connection, how the system logs it depends on why the connection was blocked; this is important to keep in mind when configuring correlation rules based on connection logs:

- For SSL rules and SSL policy default actions that block encrypted traffic, the system logs **end**-of-connection events. This is because the system cannot determine if a connection is encrypted using the first packet in the session.
- For other blocking actions, the system logs **beginning**-of-connection events. Matching traffic is denied without further inspection.

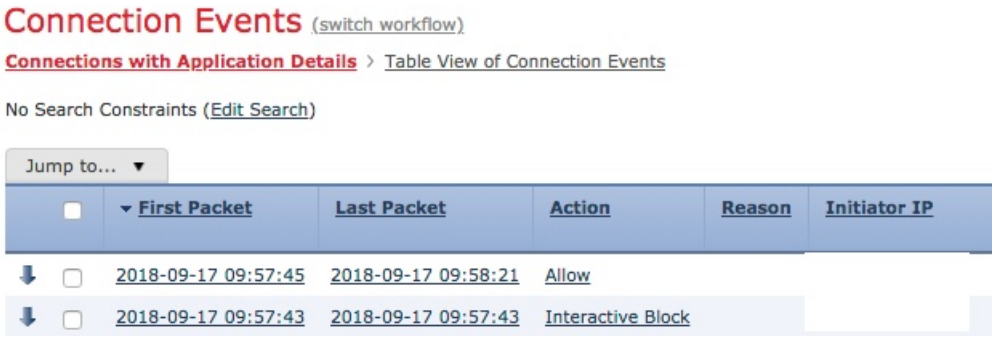
Logging Bypassed Interactive Blocks

Interactive blocking access control rules, which cause the system to display a warning page when a user browses to a prohibited website, allow you to configure end-of-connection logging. This is because if the user clicks through the warning page, the connection is considered a new, allowed connection which the system can monitor and log.

Therefore, for packets that match an Interactive Block or Interactive Block with Reset rule, the system can generate the following connection events:

- A beginning-of-connection event when a user's request is initially blocked and the warning page is displayed; this event has an associated action of `Interactive Block` or `Interactive Block with Reset`
- Multiple beginning- or end-of-connection events if the user clicks through the warning page and loads the originally requested page; these events have an associated action of `Allow` and a reason of `User Bypass`

  The following figure shows an example of an interactive block followed by allow.



Logging for Allowed Connections

You can log allowed connections, which includes traffic matching the following rules and actions:

- SSL rules—**Decrypt** action
- SSL rules—**Do not decrypt** action
- SSL default action—**Do not decrypt**
- Access control rules—**Allow** action
- Access control default action—**Network Discovery Only** and any intrusion prevention option

Enabling logging for these configurations ensures the connection is logged, while also permitting (or specifying) the next phase of inspection and traffic handling. SSL logging is always end-of-connection; access control configurations also allow beginning-of-connection logging.

Although the **Analyze** action in tunnel and prefilter rules also allows connections to continue with access control, logging is disabled for rules with this action. Matching connections may still be logged by other configurations. Allowed tunnels might have their encapsulated sessions evaluated and logged individually.

When you allow traffic with an access control rule or default action, you can use an associated intrusion policy to further inspect traffic and block intrusions. For access control rules, you can also use a file policy to detect and block prohibited files, incuding malware. Unless you disable connection event storage, the system

automatically logs most allowed connections associated with intrusion, file, and malware events. For detailed information, see Connections That Are Always Logged, on page 2208.

Connections with encrypted payloads are not subject to deep inspection, so connection events for encrypted connections contain limited information.

File and Malware Event Logging for Allowed Connections

When a file policy detects or blocks a file, it logs one of the following events to the Firepower Management Center database:

- *File events*, which represent detected or blocked files, including malware files
- *Malware events*, which represent detected or blocked malware files only
- *Retrospective malware events*, which are generated when the malware disposition for a previously detected file changes

You can disable this logging on a per-access-control-rule basis. You can also disable file and malware event storage entirely.

Note We recommend you leave file and malware event logging enabled.

Beginning vs End-of-Connection Logging

You can log a connection at its beginning or its end, with the following exceptions for blocked traffic:

- Blocked traffic—Because blocked traffic is immediately denied without further inspection, usually you can log only beginning-of-connection events for blocked traffic. There is no unique end of connection to log.
- Blocked encrypted traffic—When you enable connection logging in an SSL policy, the system logs end-of-connection rather than beginning-of-connection events. This is because the system cannot determine if a connection is encrypted using the first packet in the session, and thus cannot immediately block encrypted sessions.

To optimize performance, log either the beginning or the end of any connection, but not both. Monitoring a connection for any reason forces end-of-connection logging. For a single non-blocked connection, the end-of-connection event contains all of the information in the beginning-of-connection event, as well as information gathered over the duration of the session.

The following table details the differences between beginning and end-of-connection events, including the advantages to logging each.

Table 290: Comparing Beginning and End-of-Connection Events

	Beginning-of-Connection Events	End-of-Connection Events
Can be generated...	When the system detects the beginning of a connection (or, after the first few packets if event generation depends on application or URL identification).	When the system: • Detects the close of a connection. • Does not detect the end of a connection aft of time. • Can no longer track the session due to mer constraints.
Can be logged for...	All connections except those blocked by the SSL policy.	Most connections
Contain...	Only information that can be determined in the first packet (or the first few packets, if event generation depends on application or URL identification).	All information in the beginning-of-connection information determined by examining traffic ov duration of the session; for example, the total am transmitted or the timestamp of the last packet i connection. **Note** The connection event does not cou amount of data transmitted after th defense returns a snort verdict for t connection or if you fastpath the co
Are useful...	If you want to log: • Blocked connections. • Only the beginning of a connection because the end-of-connection information does not matter to you.	If you want to: • Log encrypted connections handled by an • Perform any kind of detailed analysis on, c correlation rules using, information collect duration of the session. • View connection summaries (aggregated c data) in custom workflows, view connectio graphical format, or create and use traffic p

Firepower Management Center vs External Logging

If you store connection and Security Intelligence event logs on the Firepower Management Center, you can use the Firepower System's reporting, analysis, and data correlation features. For example:

- Dashboards and the Context Explorer provide you with graphical, at-a-glance views of the connections logged by the system.
- Event views (most of the options available under the Analysis menu) present detailed information on the connections logged by the system, which you can display in a graphical or tabular format or summarize in a report.
- Traffic profiling uses connection data to create a profile of your normal network traffic that you can then use as a baseline against which to detect and track anomalous behavior.

- Correlation policies allow you to generate events and trigger responses (such as alerts or external remediations) to specific types of connections or traffic profile changes.

The number of events the Firepower Management Center can store depends on its model.



Note To use these features, you **must** log connections (and in most cases, the end of those connections rather than the beginning). This is why the system automatically logs critical connections—those associated with logged intrusions, prohibited files, and malware.

You can also log events to an external syslog or SNMP trap server using the following:

- For external logging on any device:

  A connection you configure called an *alert response*.

- For external logging on FTD devices:

  See About Configuring Syslog, on page 962 and Configure SNMP Traps, on page 960.

Related Topics

Firepower Management Center Alert Responses, on page 2075

Limitations of Connection Logging

You cannot log:

- Connections fastpathed with 8000 Series fastpath rules
- The outer session of a plaintext, passthrough tunnel whose encapsulated connections are inspected by access control
- TCP connections if the three-way handshake is not completed.

  These connections are not logged as doing so would provide an opportunity for a denial-of-service attack against your Firepower deployment.

  However, you can use the following workaround to monitor or debug failed connections:

  - Use the use **show asp drops** command in the command-line interface.
  - Use the packet capture feature to get more details about these connections. See Packet Capture Overview, on page 327 its and subtopics.

If a connection event does not contain the information you think it should, see Requirements for Populating Connection Event Fields, on page 2240 and Information Available in Connection Event Fields, on page 2242.

When Events Appear in the Event Viewer

The following points are applicable to all types of events:

- If you are looking at a page under the Analysis menu, you must refresh the page to display new events.

- Events generally are available for viewing within a few seconds of the time the traffic was detected. However, there can be an arbitrary delay under situations such as: Exceptionally heavy traffic conditions; the FMC is managing many devices on a low-bandwidth network; or during operations such as event backup which pause event processing.
- All the connection events logged as per the defined rules appear in the event viewer. An option to filter the events is not available for the unified logging of the connection events.

Best Practices for Connection Logging

Use the following best practices to ensure that you log *only* the connections you want to log.

So that you log only critical connections, enable connection logging on a per-access-control-rule basis.

Connections that are always logged

The system automatically logs the following:

- Some connections associated with detected files, malware, intrusions, and Intelligent Application Bypass (IAB).

  For more information, see Connections That Are Always Logged, on page 2208.

- Monitored connections.

  For more information, see Logging for Monitored Connections, on page 2210.

Connections to never log

Do not enable logging for the following:

- Access control rules with a Trust action.

  Trusted connections are not subject to deep inspection or discovery, so connection events for trusted connections contain limited information.

- Do not enable logging for Block rules in passive deployments. To log connections that the system *would have blocked* if your devices were deployed inline, use a Monitor rule instead of a Block rule.

  Only devices deployed inline (that is, using routed, switched, or transparent interfaces, or inline interface pairs) can block traffic. Because blocked connections are not actually blocked in passive deployments, the system may report multiple beginning-of-connection events for each blocked connection.

- Traffic you're not interested in. Examples follow:
  - Specific allowed traffic, such as DNS requests to a trusted DNS host.
  - Infrastructure traffic that is not related to your service offering.

  (As previously mentioned, you can still monitor this traffic for threats.)

As discussed in Connections That Are Always Logged, on page 2208, even if you disable logging for the preceding, intrusion events, malware, and IAB are still logged.

Avoid logging what's being logged elsewhere

If another device or service is logging connection data for a network segment, disable logging for that segment's data in the Firepower Management Center. Examples follow:

- If a router logs connection events on the same network segment as the Firepower Management Center, avoid logging the same connections on the Firepower Management Center unless you need those connection events for something else, such as correlation policies or traffic profiles.

  For more information about correlation policies, see Introduction to Correlation Policies and Rules, on page 1985. For more information about traffic profiles, see Introduction to Traffic Profiles, on page 2021.

- If you use Secure Network Analytics to leverage NetFlow records reported from switches and routers to identify potential behavioral anomalies and suspicious traffic patterns, you can disable connection logging for rules monitoring those segments and instead rely on Secure Network Analytics for behavioral analytics for those parts of your network.

  For more information, consult the Stealthwatch documentation.

Log either the beginning or end of the connection (not both)

If you have a choice between beginning and end-of-connection logging, enable end-of-connection logging. This is because end-of-connection logs information from beginning-of-connection events, as well as information gathered over the duration of the session.

Log the beginning of connections *only* if you want to log blocked connections, or if end-of-connection information does not matter to you.

For more information, see Beginning vs End-of-Connection Logging, on page 2213.

Logging for blocked traffic

Because blocked traffic is immediately denied without further inspection, usually you can log only beginning-of-connection events.

For more information, see Logging for Blocked Connections, on page 2211.

Log events to an external location

If your company's security policy permits it, you can save disk space on your Firepower Management Center by streaming logs to an external source using any of the following:

- eStreamer, which enables you to stream logs from a Firepower Management Center or 7000 or 8000 Series device to a custom-developed client application. For more information, see the *Firepower eStreamer Integration Guide*.

- Syslog or SNMP trap, which are referred to as *alert responses*. For more information, see Firepower Management Center Alert Responses, on page 2075.

Specify the maximum number of event records

Consider the minimum and maximum number of records that can be stored in the database. For example, a virtual Firepower Management Center by default stores 10 million events but the maximum number of events is 50 million. Go to **System** > **Configuration** > **Database** to adjust the size to meet your needs.

For a list of all Firepower Management Center models and their event database sizes, see Database Event Limits, on page 879.

Control what is displayed in connection events

To specify the number of rows displayed in connection events, click your username in the upper right of the Firepower Management Center and click **User Preferences** > **Event View Settings**. The maximum you can set is 1000 events per page.

Set up connection event reports

To make sure you do not miss connection events, you can set up automated reports in .csv format and optionally schedule them to occur at a regular interval. For more information, see the following:

- Use the report designer (**Analysis** > **Connection** > **Events** > **Report Designer**): About Designing Reports, on page 2051.
- Schedule tasks (**System** > **Tools** > **Scheduling**): About Task Scheduling, on page 179.

Requirements and Prerequisites for Connection Logging

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Configure Connection Logging

The following sections describe how to set up connection logging to match various rules and conditions.

Logging Connections with Tunnel and Prefilter Rules

The prefilter policy applies to Firepower Threat Defense devices only.

Before you begin

- Set the rule action to **Block** or **Fastpath**. Logging is disabled for the **Analyze** action, which allows connections to continue with access control, where other configurations determine their handling and logging.

Procedure

Step 1 In the prefilter policy editor, click **Edit** ( ) next to the rule where you want to configure logging.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 2 Click **Logging**.

Step 3 Specify whether you want to **Log at Beginning of Connection** or **Log at End of Connection**.

To optimize performance, log either the beginning or the end of any connection, but not both. Because blocked traffic is immediately denied without further inspection, you cannot log end-of-connection events for Block rules.

Step 4 Specify where to send connection events:

Send events to the event viewer if you want to perform Firepower Management Center-based analysis on these connection events.

Step 5 Click **Save** to save the rule.

Step 6 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Logging Decryptable Connections with SSL Rules

SSL rules do not apply to NGIPSv devices.

Procedure

Step 1 In the SSL policy editor, click **Edit** ( ) next to the rule where you want to configure logging.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 2 Click **Logging**.

Step 3 Check **Log at End of Connection**.

For monitored traffic, end-of-connection logging is required.

Step 4 Specify where to send connection events.

Send events to the event viewer if you want to perform Firepower Management Center-based analysis on these connection events. For monitored traffic, this is required.

Step 5 Click **Save** to save the rule.

Step 6 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Logging Connections with Security Intelligence

The Security Intelligence policy requires the Threat Smart License or Protection Classic License.

Procedure

Step 1 In the access control policy editor, click **Security Intelligence**.

Step 2 Click **Logging** ( ) to enable Security Intelligenge logging using the following critera:

- By IP address—Click logging next to **Networks**.
- By URL—Click logging next to **URLs**.
- By Domain Name—Click logging next to the **DNS Policy** drop-down list.

If the controls are dimmed, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Check the **Log Connections** check box.

Step 4 Specify where to send connection and Security Intelligence events.

Send events to the event viewer if you want to perform Firepower Management Center-based analysis, or if you set a Block list to monitor-only.

Step 5 Click **OK** to set logging options.

Step 6 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Logging Connections with Access Control Rules

Depending on your choices for the rule action and deep inspection options, your logging options differ; see How Rules and Policy Actions Affect Logging, on page 2209.

Procedure

Step 1 In the access control policy editor, click **Edit** ( ) next to the rule where you want to configure logging.

If **View** ( ) appears instead, the configuration is inherited from an ancestor policy, belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 2 Click the **Logging** tab.

Step 3 Specify whether you want to **Log at Beginning of Connection** or **Log at End of Connection**.

To optimize performance, log either the beginning or the end of any connection, but not both.

Step 4 (Optional) Check the **Log Files** check box to log file and malware events associated with the connection.

Cisco recommends you leave this option enabled.

Step 5 Specify where to send the connection events:

- **Event Viewer**: Send connection events to Firepower Management Center web interface if you want to perform Firepower Management Center-based analysis on these connection events, or if the rule action is **Monitor**.
- **Syslog Server**: Send connection events to the syslog server configured in the Logging tab in Access Control Policy, unless overridden.
- **SNMP Trap**: Connection events are sent to the selected SNMP trap.

Step 6 Click **Save** to save the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Logging Connections with a Policy Default Action

A policy's default action determines how the system handles traffic that matches none of the rules in the policy (except Monitor rules in access control and SSL policies, which match and log—but do not handle or inspect—traffic).

Logging settings for the SSL policy default action also govern how the system logs undecryptable sessions.

Before you begin

- For prefilter default action logging, set the default action to **Block all tunnel traffic**. Logging is disabled for the **Allow all tunnel traffic** action, which allows connections to continue with access control, where other configurations determine their handling and logging.

Procedure

Step 1 In the policy editor, click **Logging** ( ) next to the **Default Action** drop-down list.

Step 2 Specify when you want to log matching connections:

- Log at Beginning of Connection—Not supported for SSL default actions.

- Log at End of Connection—Not supported if you choose the access control **Block All Traffic** default action or the prefilter **Block all tunnel traffic** default action.

To optimize performance, log either the beginning or the end of any connection, but not both.

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration. In an access control policy, the configuration may also be inherited from an ancestor policy.

Step 3 Specify where to send connection events.

Send events to the event viewer if you want to perform Firepower Management Center-based analysis on these connection events.

Step 4 Click **OK**.

Step 5 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.

Limiting Logging of Long URLs

End-of-connection events for HTTP traffic record the URL requested by monitored hosts. Disabling or limiting the number of stored URL characters may improve system performance. Disabling URL logging (storing zero characters) does not affect URL filtering. The system filters traffic based on requested URLs even though the system does not record them.

Procedure

Step 1 In the access control policy editor, click **Advanced**, then click **Edit** () next to **General Settings**.

If **View** () appears instead, the configuration is inherited from an ancestor policy, belongs to an ancestor domain, or you do not have permission to modify the configuration. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 Enter the **Maximum URL characters to store in connection events**.

Step 3 Click **OK**.

Step 4 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 346.



CHAPTER 112

Connection and Security Intelligence Events

The following topics describe how to use connection and security events tables.

- About Connection Events, on page 2223
- Connection and Security Intelligence Event Fields, on page 2225
- Using Connection and Security Intelligence Event Tables, on page 2245
- Viewing the Connection Summary Page, on page 2249

About Connection Events

The system can generate logs of the connections its managed devices detect. These logs are called *connection events*. Connection events include *Security Intelligence events* (connections blocked by the reputation-based Security Intelligence feature.)

Connection events generally include transactions detected by:

- Access Control policies
- SSL policies
- Prefilter policies (captured by prefilter or tunnel rules)
- DNS Block lists
- URL Block lists
- Network (IP address) Block lists

Settings in rules and policies give you granular control over which connections you log, when you log them, and where you store the data.

For detailed information, see Connection Logging, on page 2207.

Related Topics

About Security Intelligence, on page 1249

Connection vs. Security Intelligence Events

A *Security Intelligence event* is a connection event that is generated whenever a session is blocked or monitored by the reputation-based Security Intelligence feature.

However, for every Security Intelligence event, there is an identical connection event. You can view and analyze Security Intelligence events independently. The system also stores and prunes Security Intelligence events separately.

Note that the system enforces Security Intelligence before more resource-intensive evaluations. When a connection is blocked by Security Intelligence, the resulting event does not contain the information that the system would have gathered from subsequent evaluation, for example, user identity.



Note In this guide, information about connection events also pertains to Security Intelligence events, unless otherwise noted.

NetFlow Connections

To supplement the connection data gathered by your managed devices, you can use records broadcast by NetFlow exporters to generate connection events. This is especially useful if the NetFlow exporters are monitoring different networks than those monitored by your managed devices.

The system logs NetFlow records as unidirectional end-of-connection events in the Firepower Management Center database. The available information for these connections differs somewhat from connections detected by your access control policy; see Differences between NetFlow and Managed Device Data, on page 1824.

Related Topics

Netflow Data in the Firepower System, on page 1823

Connection Summaries (Aggregated Data for Graphs)

The Firepower System aggregates connection data collected over five-minute intervals into connection summaries, which the system uses to generate connection graphs and traffic profiles. Optionally, you can create custom workflows based on connection summary data, which you use in the same way as you use workflows based on individual connection events.

Note that there are no connection summaries specifically for Security Intelligence events, although corresponding end-of-connection events can be aggregated into connection summary data.

To be aggregated, multiple connections must:

- represent the end of connections
- have the same source and destination IP addresses, and use the same port on the responder (destination) host
- use the same protocol (TCP or UDP)
- use the same application protocol
- either be detected by the same Firepower System managed device or by the same NetFlow exporter

Each connection summary includes total traffic statistics, as well as the number of connections in the summary. Because NetFlow exporters generate unidirectional connections, a summary's connection count is incremented by two for every connection based on NetFlow data.

Note that connection summaries do not contain all of the information associated with the summaries' aggregated connections. For example, because client information is not used to aggregate connections into connection summaries, summaries do not contain client information.

Long-Running Connections

If a monitored session spans two or more five-minute intervals over which connection data is aggregated, the connection is considered a *long-running connection*. When calculating the number of connections in a connection summary, the system increments the count only for the five-minute interval in which a long-running connection was initiated.

Also, when calculating the number of packets and bytes transmitted by the initiator and responder in a long-running connection, the system does not report the number of packets and bytes that were actually transmitted during each five-minute interval. Instead, the system assumes a constant rate of transmission and calculates estimated figures based on the total number of packets and bytes transmitted, the length of the connection, and what portion of the connection occurred during each five-minute interval.

Combined Connection Summaries from External Responders

To reduce the space required to store connection data and speed up the rendering of connection graphs, the system combines connection summaries when:

- one of the hosts involved in the connection is not on your monitored network
- other than the IP address of the external host, the connections in the summaries meet the summary aggregation criteria

When viewing connection summaries in the Analysis > Connections submenu pages, and when working with connection graphs, the system displays `external` instead of an IP address for the non-monitored hosts.

As a consequence of this aggregation, if you attempt to drill down to the table view of connection data (that is, access data on individual connections) from a connection summary or graph that involves an external responder, the table view contains no information.

Connection and Security Intelligence Event Fields



Note You cannot use the connection/Security Intelligence events Search page to search for events associated with a connection.

Access Control Policy (Syslog: ACPolicy)

The access control policy that monitored the connection.

Access Control Rule (Syslog: AccessControlRuleName)

The access control rule or default action that handled the connection, as well as up to eight Monitor rules matched by that connection.

If the connection matched one Monitor rule, the Firepower Management Center displays the name of the rule that handled the connection, followed by the Monitor rule name. If the connection matched more than one Monitor rule, the number of matching Monitor rules is displayed, for example, `Default Action + 2 Monitor Rules`.

To display a pop-up window with a list of the first eight Monitor rules matched by the connection, click ***N* Monitor Rules**.

Action (Syslog: AccessControlRuleAction)

The action associated with the configuration that logged the connection.

For Security Intelligence-monitored connections, the action is that of the first non-Monitor access control rule triggered by the connection, or the default action. Similarly, because traffic matching a Monitor rule is always handled by a subsequent rule or by the default action, the action associated with a connection logged due to a Monitor rule is never Monitor. However, you can still trigger correlation policy violations on connections that match Monitor rules.

Action	Description
Allow	Connections either allowed by access control explicitly, or allowed because a user bypassed an interactive block.
Block, Block with reset	Blocked connections, including: • tunnels and other connections blocked by the prefilter policy • connections blocked by Security Intelligence • encrypted connections blocked by an SSL policy • connections where an exploit was blocked by an intrusion policy • connections where a file (including malware) was blocked by a file policy For connections where the system blocks an intrusion or file, system displays `Block`, even though you use access control Allow rules to invoke deep inspection.
Fastpath	Non-encrypted tunnels and other connections fastpathed by the prefilter policy.
Interactive Block, Interactive Block with reset	Connections logged when the system initially blocks a user's HTTP request using an Interactive Block rule. If the user clicks through the warning page that the system displays, additional connections logged for the session have an action of Allow.
Trust	Connections trusted by access control. The system logs trusted TCP connections differently depending on the device model.
Default Action	Connections handled by the access control policy's default action.
(Blank/empty)	The connection closed before enough packets had passed to match a rule. This can happen only if a facility other than access control, such as intrusion prevention, causes the connection to be logged.

Application Protocol (Syslog: ApplicationProtocol)

In the Firepower Management Center web interface, this value constrains summaries and graphs.

The application protocol, which represents communications between hosts, detected in the connection.

Application Protocol Category and Tag

Criteria that characterize the application to help you understand the application's function.

Application Risk

The risk associated with the application traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of application detected in the connection has an associated risk; this field displays the highest of those.

Business Relevance

The business relevance associated with the application traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of application detected in the connection has an associated business relevance; this field displays the lowest (least relevant) of those.

Client and Client Version (Syslog: Client, ClientVersion)

The client application and version of that client detected in the connection.

If the system cannot identify the specific client used in the connection, the field displays the word "client" appended to the application protocol name to provide a generic name, for example, FTP client.

Client Category and Tag

Criteria that characterize the application to help you understand the application's function.

Connections

The number of connections in a connection summary. For long-running connections, that is, connections that span multiple connection summary intervals, only the first connection summary interval is incremented. To view meaningful results for searches using the **Connections** criterion, use a custom workflow that has a connection summary page.

Count

The number of connections that match the information that appears in each row. Note that the **Count** field appears only after you apply a constraint that creates two or more identical rows. If you create a custom workflow and do not add the **Count** column to a drill-down page, each connection is listed individually and packets and bytes are not summed.

Destination Port/ICMP Code (Syslog: Separate fields - DstPort, ICMPCode)

In the Firepower Management Center web interface, these values constrain summaries and graphs.

The port or ICMP code used by the session responder.

Detection Type

This field shows the source of detection of a client.

Device

In the Firepower Management Center web interface, this value constrains summaries and graphs.

The managed device that detected the connection or, for connections generated from NetFlow data, the managed device that processed the data.

DNS Query (Syslog: DNSQuery)

The DNS query submitted in a connection to the name server to look up a domain name.

DNS Record Type (Syslog: DNSRecordType)

The type of the DNS resource record used to resolve a DNS query submitted in a connection.

DNS Response (Syslog: DNSResponseType)

The DNS response returned in a connection to the name server when queried.

DNS Sinkhole Name (Syslog: DNS_Sinkhole)

The name of the sinkhole server where the system redirected a connection.

DNS TTL (Syslog: DNS_TTL)

The number of seconds a DNS server caches the DNS resource record.

Domain

The domain of the managed device that detected the connection or, for connections generated from NetFlow data, the domain of the managed device that processed the data. This field is only present if you have ever configured the Firepower Management Center for multitenancy.

Endpoint Location

The IP address of the network device that used ISE to authenticate the user, as identified by ISE.

Endpoint Profile (Syslog: Endpoint Profile)

The user's endpoint device type, as identified by ISE.

Files (Syslog: FileCount)

The number of files (including malware files) detected or blocked in a connection associated with one or more file events.

In the Firepower Management Center web interface, the **View Files icon** links to a list of files. The number on the icon indicates the number of files (including malware files) detected or blocked in that connection.

First Packet or Last Packet (Syslog: See the ConnectionDuration field)

The date and time the first or last packet of the session was seen.

HTTP Referrer (Syslog: HTTPReferer)

The HTTP referrer, which represents the referrer of a requested URL for HTTP traffic detected in the connection (such as a website that provided a link to, or imported a link from, another URL).

HTTP Response Code (Syslog: HTTPResponse)

The HTTP status code sent in response to a client's HTTP request over a connection. It indicates the reason behind successful and failed HTTP request.

For more details about HTTP response codes, see RFC 2616 (HTTP), Section 10.

Ingress/Egress Interface (Syslog: IngressInterface, EgressInterface)

The ingress or egress interface associated with the connection. If your deployment includes an asymmetric routing configuration, the ingress and egress interface may not belong to the same inline pair.

Ingress/Egress Security Zone (Syslog: IngressZone, EgressZone)

The ingress or egress security zone associated with the connection.

For rezoned encapsulated connections, the ingress field displays the tunnel zone you assigned, instead of the original ingress security zone. The egress field is blank.

Initiator/Responder Bytes (Syslog: InitiatorBytes, ResponderBytes)

The total number of bytes transmitted by the session initiator or received by the session responder.

Initiator/Responder Continent

When a routable IP is detected, the continent associated with the IP address for the session initiator or responder.

Initiator/Responder Country

When a routable IP is detected, the country associated with the IP address of the session initiator or responder. The system displays an icon of the country's flag, and the country's ISO 3166-1 alpha-3 country code. Hover your pointer over the flag icon to view the country's full name.

Initiator/Responder IP (Syslog: SrcIP, DstIP)

In the Firepower Management Center web interface, these values constrain summaries and graphs.

The IP address (and host name, if DNS resolution is enabled) of the session initiator or responder.

See also A Note About Initiator/Responder, Source/Destination, and Sender/Receiver Fields, on page 2238.

In the Firepower Management Center web interface, the host icon identifies the IP address that caused the connection to be blocked.

For plaintext, passthrough tunnels either blocked or fastpathed by the prefilter policy, initiator and responder IP addresses represent the tunnel endpoints—the routed interfaces of the network devices on either side of the tunnel.

Initiator/Responder Packets (Syslog: InitiatorPackets, ResponderPackets)

The total number of packets transmitted by the session initiator or received by the session responder.

Initiator User (Syslog: User)

In the Firepower Management Center web interface, this value constrains summaries and graphs.

The user logged into the session initiator. If this field is populated with **No Authentication**, the user traffic:

- matched an access control policy without an associated identity policy
- did not match any rules in the identity policy

See also A Note About Initiator/Responder, Source/Destination, and Sender/Receiver Fields, on page 2238.

Intrusion Events (Syslog: IPSCount)

The number of intrusion events, if any, associated with the connection.

In the Firepower Management Center web interface, the **View Intrusion Events icon** links to a list of events.

IOC

Whether the event triggered an indication of compromise (IOC) against a host involved in the connection.

NetBIOS Domain (Syslog: NetBIOSDomain)

The NetBIOS domain used in the session.

NetFlow SNMP Input/Output

For connections generated from NetFlow data, the interface index for the interface where connection traffic entered or exited the NetFlow exporter.

NetFlow Source/Destination Autonomous System

For connections generated from NetFlow data, the border gateway protocol autonomous system number for the source or destination of traffic in the connection.

NetFlow Source/Destination Prefix

For connections generated from NetFlow data, the source or destination IP address ANDed with the source or destination prefix mask.

NetFlow Source/Destination TOS

For connections generated from NetFlow data, the setting for the type-of-service (TOS) byte when connection traffic entered or exited the NetFlow exporter.

Network Analysis Policy (Syslog: NAPPolicy)

The network analysis policy (NAP), if any, associated with the generation of the event.

Original Client Country

The country where the original client IP address belongs. To obtain this value, the system extracts the original client IP address from an X-Forwarded-For (XFF), True-Client-IP, or custom-defined HTTP header, then maps it to the country using the geolocation database (GeoDB). To populate this field, you must enable an access control rule that handles proxied traffic based on its original client.

Original Client IP (Syslog: originalClientSrcIP)

The original client IP address from an X-Forwarded-For (XFF), True-Client-IP, or custom-defined HTTP header. To populate this field, you must enable an access control rule that handles proxied traffic based on its original client.

Prefilter Policy (Syslog: Prefilter Policy)

The prefilter policy that handled the connection.

Protocol (Syslog: Protocol)

In the Firepower Management Center web interface:

- This value constrains summaries and graphs.
- This field is available only as a search field.

The transport protocol used in the connection. To search for a specific protocol, use the name or number protocol as listed in http://www.iana.org/assignments/protocol-numbers.

QoS-Applied Interface

For rate-limited connections, the name of the interface where you applied rate limiting.

QoS-Dropped Initiator/Responder Bytes

The number of bytes dropped from the session initiator or session responder due to rate limiting.

QoS-Dropped Initiator/Responder Packets

The number of packets dropped from the session initiator or session responder due to rate limiting.

QoS Policy

The QoS policy that rate limited the connection.

QoS Rule

The QoS rule that rate limited the connection.

Reason (Syslog: AccessControlRuleReason)

The reason or reasons the connection was logged, in many situations. For a full list, see Connection Event Reasons, on page 2239.

Connections with a Reason of IP Block, DNS Block, and URL Block have a threshold of 15 seconds per unique initiator-responder pair. After the system blocks one of those connections, it does not generate connection events for additional blocked connections between those two hosts for the next 15 seconds, regardless of port or protocol.

Referenced Host (Syslog: ReferencedHost)

If the protocol in the connection is HTTP or HTTPS, this field displays the host name that the respective protocol was using.

Security Context (Syslog: Context)

For connections handled by ASA FirePOWER in multiple context mode, the metadata identifying the virtual firewall group through which the traffic passed.

Security Group Tag (Syslog: Security Group)

The Security Group Tag (SGT) attribute of the packet involved in the connection. The SGT specifies the privileges of a traffic source within a trusted network. Security Group Access (a feature of both Cisco TrustSec and Cisco ISE) applies the attribute as packets enter the network.

Security Intelligence Category (Syslog: URLSICategory, DNSSICategory)

The name of the object that represents or contains the IP address that caused the connection to be blocked. The Security Intelligence category can be the name of a network object or group, a Block list, a custom Security Intelligence list or feed, or one of the categories in the Intelligence Feed.

In the Firepower Management Center web interface, DNS, Network (IP address), and URL Security Intelligence connection events are combined into a single category field. In syslog messages, those events are specific by type.

For more information about the categories in the Intelligence Feed, see Security Intelligence Categories, on page 1255.

Source Device

In the Firepower Management Center web interface, this value constrains summaries and graphs.

The IP address of the NetFlow exporter that broadcast the data used to generate for the connection. If the connection was detected by a managed device, this field displays `Firepower`.

Source Port/ICMP Type (Syslog: SrcPort, ICMPType)

In the Firepower Management Center web interface, these values constrain summaries and graphs.

The port or ICMP type used by the session initiator.

SSL Actual Action (Syslog: SSLActualAction)

In the Firepower Management Center web interface, this field is a search field only.

The system displays field values in the **SSL Status** field on search workflow pages.

The action the system applied to encrypted traffic in the SSL policy.

Action	Description
Block/Block with reset	Represents blocked encrypted connections.
Decrypt (Resign)	Represents an outgoing connection decrypted using a re-signed server certificate.
Decrypt (Replace Key)	Represents an outgoing connection decrypted using a self-signed server certificate with a substituted public key.
Decrypt (Known Key)	Represents an incoming connection decrypted using a known private key.
Default Action	Indicates the connection was handled by the default action.
Do not Decrypt	Represents a connection the system did not decrypt.

SSL Certificate Information (Syslog: SSLCertificate)

In the Firepower Management Center web interface, this field is a search field only.

The information stored on the public key certificate used to encrypt traffic, including:

- Subject/Issuer Common Name
- Subject/Issuer Organization
- Subject/Issuer Organization Unit
- Not Valid Before/After
- Serial Number
- Certificate Fingerprint
- Public Key Fingerprint

SSL Certificate Status (Syslog: SSLServerCertStatus)

This applies only if you configured a Certificate Status SSL rule condition. If encrypted traffic matches an SSL rule, this field displays one or more of the following server certificate status values:

- Self Signed
- Valid
- Invalid Signature
- Invalid Issuer
- Expired
- Unknown

- Not Valid Yet
- Revoked

If undecryptable traffic matches an SSL rule, this field displays `Not Checked`.

SSL Cipher Suite (Syslog: SSSLCipherSuite)

A macro value representing a cipher suite used to encrypt the connection. See www.iana.org/assignments/tls-parameters/tls-parameters.xhtml for cipher suite value designations.

SSL Encryption applied to the connection

This field is available only as a search field in the Firepower Management Center web interface.

Enter **`yes`** or **`no`** in the **SSL** search field to view TLS/SSL-encrypted or non-encrypted connections.

SSL Expected Action (Syslog: SSLExpectedAction)

In the Firepower Management Center web interface, this field is a search field only.

The action the system expected to apply to encrypted traffic, given the SSL rules in effect.

Enter any of the values listed for **SSL Actual Action**.

SSL Failure Reason (Syslog: SSLFlowStatus)

The reason the system failed to decrypt encrypted traffic:

- Unknown
- No Match
- Success
- Uncached Session
- Unknown Cipher Suite
- Unsupported Cipher Suite
- Unsupported SSL Version
- SSL Compression Used
- Session Undecryptable in Passive Mode
- Handshake Error
- Decryption Error
- Pending Server Name Category Lookup
- Pending Common Name Category Lookup
- Internal Error
- Network Parameters Unavailable
- Invalid Server Certificate Handle
- Server Certificate Fingerprint Unavailable
- Cannot Cache Subject DN

- Cannot Cache Issuer DN
- Unknown SSL Version
- External Certificate List Unavailable
- External Certificate Fingerprint Unavailable
- Internal Certificate List Invalid
- Internal Certificate List Unavailable
- Internal Certificate Unavailable
- Internal Certificate Fingerprint Unavailable
- Server Certificate Validation Unavailable
- Server Certificate Validation Failure
- Invalid Action

Field values are displayed in the **SSL Status** field on the search workflow pages.

SSL Flow Error

The error name and hexadecimal code if an error occurred during the TLS/SSL session; `Success` if no error occurred.

SSL Flow Flags

The first ten debugging level flags for an encrypted connection. On a workflow page, to view all flags, click the ellipsis (**...**).

SSL Flow Messages

The keywords below indicate encrypted traffic is associated with the specified message type exchanged between client and server during the TLS/SSL handshake. See http://tools.ietf.org/html/rfc5246 for more information.

- HELLO_REQUEST
- CLIENT_ALERT
- SERVER_ALERT
- CLIENT_HELLO
- SERVER_HELLO
- SERVER_CERTIFICATE
- SERVER_KEY_EXCHANGE
- CERTIFICATE_REQUEST
- SERVER_HELLO_DONE
- CLIENT_CERTIFICATE
- CLIENT_KEY_EXCHANGE
- CERTIFICATE_VERIFY

- CLIENT_CHANGE_CIPHER_SPEC
- CLIENT_FINISHED
- SERVER_CHANGE_CIPHER_SPEC
- SERVER_FINISHED
- NEW_SESSION_TICKET
- HANDSHAKE_OTHER
- APP_DATA_FROM_CLIENT
- APP_DATA_FROM_SERVER

SSL Policy (Syslog: SSLPolicy)

The SSL policy that handled the connection.

SSL Rule (Syslog: SSLRuleName)

The SSL rule or default action that handled the connection, as well as the first Monitor rule matched by that connection. If the connection matched a Monitor rule, the field displays the name of the rule that handled the connection, followed by the Monitor rule name.

SSLServerName (Syslog Only)

This field exists ONLY as a syslog field; it does not exist in the Firepower Management Center web interface.

Hostname of the server with which the client established an encrypted connection.

SSL Session ID (Syslog: SSLSessionID)

The hexadecimal Session ID negotiated between the client and server during the TLS/SSL handshake.

SSL Status

The action associated with the **SSL Actual Action** (SSL rule, default action, or undecryptable traffic action) that logged the encrypted connection. The **Lock icon** links to SSL certificate details. If the certificate is unavailable (for example, for connections blocked due to TLS/SSL handshake error), the lock icon is dimmed.

If the system fails to decrypt an encrypted connection, it displays the **SSL Actual Action** (undecryptable traffic action) taken, as well as the **SSL Failure Reason**. For example, if the system detects traffic encrypted with an unknown cipher suite and allows it without further inspection, this field displays `Do Not Decrypt (Unknown Cipher Suite)`.

When searching this field, enter one or more of the **SSL Actual Action** and **SSL Failure Reason** values to view encrypted traffic the system handled or failed to decrypt.

SSL Subject/Issuer Country

This field is available only in the Firepower Management Center web interface, and only as a search field.

A two-character ISO 3166-1 alpha-2 country code for the subject or issuer country associated with the encryption certificate.

SSL Ticket ID (Syslog: SSLTicketID)

A hexadecimal hash value of the session ticket information sent during the TLS/SSL handshake.

SSLURLCategory (Syslog Only)

URL categories for the URL visited in the encrypted connection.

This field exists ONLY as a syslog field; in the Firepower Management Center web interface, values in this field are included in the URL Category column.

See also **URL**.

SSL Version (Syslog: SSLVersion)

The TLS/SSL protocol version used to encrypt the connection:

- Unknown
- SSLv2.0
- SSLv3.0
- TLSv1.0
- TLSv1.1
- TLSv1.2

TCP Flags (Syslog: TCPFlags)

For connections generated from NetFlow data, the TCP flags detected in the connection.

When searching this field, enter a list of comma-separated TCP flags to view all connections that have *at least* one of those flags.

Time

The ending time of the five-minute interval that the system used to aggregate connections in a connection summary. This field is not searchable.

TLS Fingerprint Process Name

Process or client in the TLS client hello packet that was analyzed by the encrypted visibility engine.

TLS Fingerprint Process Confidence Score

The confidence value in the range 0-100% that the encrypted visibility engine has detected the right process. For example, if the process name is Firefox and if the confidence score is 80%, it means that the engine is 80% confident that the process it has detected is Firefox.

TLS Fingerprint Malware Confidence

The probability level that the process detected by the encrypted visibility engine contains malware. This field indicates the bands (Very High, High, Medium, Low, or Very Low) based on the value in the malware confidence score.

TLS Fingerprint Malware Confidence Score

The confidence value in the range 0-100% that the process detected by the encrypted visibility engine contains malware. If the malware confidence score is very high, say 90%, then the TLS Fingerprint Process Name field will display "Malware."

Total Packets

This field is available only as a search field.

The total number of packets transmitted in the connection.

Traffic (KB)

This field is available only as a search field.

The total amount of data transmitted in the connection, in kilobytes.

Tunnel/Prefilter Rule (Syslog: Tunnel or Prefilter Rule)

The tunnel rule, prefilter rule, or prefilter policy default action that handled the connection.

URL, URL Category, and URL Reputation (Syslog: URL, URLCategory and SSLURLCategory, URLReputation)

The URL requested by the monitored host during the session and its associated category and reputation, if available.

For the connection event to display URL category and reputation, you must include the applicable URL rules in an access control policy and configure the rule with URL category and URL reputation under the URL's Tab.

URL category and reputation do not appear in an event if the connection is processed before it matches a URL rule.

If the system identifies or blocks a TLS/SSL application, the requested URL is in encrypted traffic, so the system identifies the traffic based on an SSL certificate. For TLS/SSL applications, therefore, this field indicates the common name contained in the certificate.

See also **SSLURLCategory**, above.

User Agent (Syslog: UserAgent)

The user-agent string application information extracted from HTTP traffic detected in the connection.

VLAN ID

The innermost VLAN ID associated with the packet that triggered the connection.

Web Application (Syslog: WebApplication)

The web application, which represents the content or requested URL for HTTP traffic detected in the connection.

If the web application does not match the URL for the event, the traffic is probably referred traffic, such as advertisement traffic. If the system detects referred traffic, it stores the referring application (if available) and lists that application as the web application.

If the system cannot identify the specific web application in HTTP traffic, this field displays `Web Browsing`.

Web Application Category and Tag

Criteria that characterize the application to help you understand the application's function.

About Connection and Security Intelligence Event Fields

In the Firepower Management Center web interface, you can view and search connection and security intelligence events using tabular and graphical workflows under the **Analysis > Connections** submenus.

Note For each Security Intelligence event, there is an identical, separately stored connection event. All Security Intelligence events have a populated **Security Intelligence Category** field.

The information available for any individual event can vary depending on how, why, and when the system logged the connection.

Search Constraints

Fields marked with an asterisk (*****) on search pages constrain connection graphs and connection summaries. Because connection graphs are based on connection summaries, the same criteria that constrain connection summaries also constrain connection graphs. If you search connection summaries using invalid search constraints and view your results using a connection summary page in a custom workflow, the invalid constraints are labeled as not applicable (N/A) and are marked with a strikethrough.

Syslog Fields

Most fields appear both in the Firepower Management Center web interface and as syslog messages. Fields without a listed syslog equivalent are not available in syslog messages. A few fields are syslog-only, as noted, and few others are separate fields in syslog messages but are consolidated fields in the web interface or vice-versa.

A Note About Initiator/Responder, Source/Destination, and Sender/Receiver Fields

Table 291: Comparison of Terms

Fields	Event Type	Description
Initiator/Responder	Connection	Initiator/responder of the connection. The initiator of a connection is not necessarily the same as the source of an intrusion or the sender of a malware file.
Source/Destination	Intrusion	Source/destination of the attack. The source of an intrusion event can be the initiator or the responder of the connection.
Sender/Receiver (Sending..., Receiving...)	File, Malware	Sender/receiver of a file or malware. The sender of a file is not necessarily the initiator of the connection, as a file may be uploaded or downloaded.

Connection Event Reasons

The Reason field in a connection event displays the reason or reasons the connection was logged, in the following situations:

Reason	Description
Content Restriction	The system modified the packet to enforce content restrictions related to either the Safe Search or YouTube EDU feature.
DNS Block	The system denied the connection without inspection, based on the domain name and Security Intelligence data. A reason of DNS Block is paired with an action of Block, Domain not found, or Sinkhole, depending on the DNS rule action.
DNS Monitor	The system would have denied the connection based on the domain name and Security Intelligence data, but you configured the system to monitor, rather than deny, the connection.
File Block	The connection contained a file or malware file that the system prevented from being transmitted. A reason of File Block is always paired with an action of Block.
File Custom Detection	The connection contained a file on the custom detection list that the system prevented from being transmitted.
File Monitor	The system detected a particular type of file in the connection.
File Resume Allow	File transmission was originally blocked by a Block Files or Block Malware file rule. After a new access control policy allowing the file was deployed, the HTTP session automatically resumed. This reason only appears in inline deployments.
File Resume Block	File transmission was originally allowed by a Detect Files or Malware Cloud Lookup file rule. After a new access control policy blocking the file was deployed, the HTTP session automatically stopped. This reason only appears in inline deployments.
Intelligent App Bypass	The Intelligent Application Bypass (IAB) mode: • If the action is Trust, IAB was in bypass mode. Matching traffic passed without further inspection. • If the action is Allow, IAB was in test mode. Matching traffic was available for further inspection.
Intrusion Block	The system blocked or would have blocked an exploit (intrusion policy violation) detected in the connection. A reason of Intrusion Block is paired with an action of Block for blocked exploits and Allow for would-have-blocked exploits.
Intrusion Monitor	The system detected, but did not block, an exploit detected in the connection. This occurs when the state of the triggered intrusion rule is set to Generate Events.
IP Block	The system denied the connection without inspection, based on the IP address and Security Intelligence data. A reason of IP Block is always paired with an action of Block.

Reason	Description
IP Monitor	The system would have denied the connection based on the IP address and Security Intelligence data, but you configured the system to monitor, rather than deny, the connection.
SSL Block	The system blocked an encrypted connection based on the TLS/SSL inspection configuration. A reason of SSL Block is always paired with an action of Block.
URL Block	The system denied the connection without inspection, based on the URL and Security Intelligence data. A reason of URL Block is always paired with an action of Block.
URL Monitor	The system would have denied the connection based on the URL and Security Intelligence data, but you configured the system to monitor, rather than deny, the connection.
User Bypass	The system initially blocked a user's HTTP request, but the user clicked through a warning page to view the site. A reason of User Bypass is always paired with an action of Allow.

Requirements for Populating Connection Event Fields

The information available for a connection event, Security Intelligence event, or connection summary depends on several factors.

Appliance Model and License

Many features require that you enable specific licensed capabilities on target devices, and many features are only available on some models.

For example, NGIPSv devices do not support TLS/SSL inspection. They cannot inspect encrypted traffic; logged connection events do not contain information about encrypted connections.

Traffic Characteristics

The system only reports information present (and detectable) in network traffic. For example, there could be no user associated with an initiator host, or no referenced host detected in a connection where the protocol is not DNS, HTTP, or HTTPS.

Origin/Detection Method: Traffic-Based Detection vs NetFlow

With the exception of NetFlow-only fields, the information available in NetFlow records is more limited than the information generated by traffic-based detection; see Differences between NetFlow and Managed Device Data, on page 1824.

Evaluation Stage

Each type of traffic inspection and control occurs where it makes the most sense for maximum flexibility and performance.

For example, the system enforces Security Intelligence before more resource-intensive evaluations. When a connection is blocked by Security Intelligence, the resulting event does not contain the information that the system would have gathered from subsequent evaluation, for example, user identity.

Logging Method: Beginning or End of Connection

When the system detects a connection, whether you can log it at its beginning or its end (or both) depends on how you configure the system to detect and handle it.

Beginning-of-connection events do not have information that must be determined by examining traffic over the duration of the session (for example, the total amount of data transmitted or the timestamp of the last packet in the connection). Beginning-of-connection events are also not guaranteed to have information about application or URL traffic in the session, and do not contain any details about the session's encryption. Beginning-of-connection logging is usually the only option for blocked connections.

Connection Event Type: Individual vs Summary

Connection summaries do not contain all of the information associated with their aggregated connections. For example, because client information is not used to aggregate connections into connection summaries, summaries do not contain client information.

Keep in mind that connection graphs are based on connection summary data, which use only end-of-connection logs. If your system is configured to log only beginning-of-connection data, connection graphs and connection summary event views contain no data.

Other Configurations

Other configurations that affect connection logging include, but are not limited to:

- ISE-related fields are populated only if you configure ISE, in connections associated with users who authenticate via an Active Directory domain controller. Connection events do not contain ISE data for users who authenticate via LDAP, RADIUS, or RSA domain controllers.
- The Security Group Tag field is populated only if you configure ISE as an identity source or add custom SGT rule conditions.
- Prefilter-related fields (including tunnel zone information in security zone fields) are populated only in connections handled by a prefilter policy.
- TLS/SSL-related fields are populated only in encrypted connections handled by an SSL policy. You can view the values of the fields using a Do Not Decrypt rule action if you do not need to decrypt the traffic.
- File information fields are populated only in connections logged by access control rules associated with file policies.
- Intrusion information fields are populated only in connections logged by access control rules either associated with intrusion policies or using the default action.
- QoS-related fields are populated only in connections subject to rate limiting.
- The Reason field is populated only in specific situations, such as when a user bypasses an Interactive Block configuration.
- The Domain field is only present if you have ever configured the Firepower Management Center for multitenancy.
- An advanced setting in the access control policy controls the number of characters the system stores in the connection log for each URL requested by monitored hosts in HTTP sessions. If you use this setting to disable URL logging, the system does not display individual URLs in the connection log, although you can still view category and reputation data, if it exists.

- For the connection event to display URL category and reputation, you must include the applicable URL rules in an access control policy and configure the rule with URL category and URL reputation under the URL's Tab. URL category and reputation do not appear in an event if the connection is processed before it matches a URL rule.

Related Topics

Differences between NetFlow and Managed Device Data, on page 1824

Information Available in Connection Event Fields

The table in this topic indicates when the system can populate connection and Security Intelligence fields. The columns in the table represent the following event types:

- Origin: Direct—Events that represent connections detected and handled by a managed device.
- Origin: NetFlow—Events that represent connections exported by a NetFlow exporter.
- Logging: Start—Events that represent connections logged at their beginning.
- Logging: End—Events that represent connections logged at their end.

A "yes" in the table does not mean that the system must populate a connection event field, rather, that it can. The system only reports information present (and detectable) in network traffic. For example, TLS/SSL-related fields are populated only for records of encrypted connections handled by an SSL policy.

Connection Event Field	Origin: Direct	Origin: NetFlow	Logging: Start	Logging: End
Access Control Policy	yes	no	yes	yes
Access Control Rule	yes	no	yes	yes
Action	yes	no	yes	yes
Application Protocol	yes	yes	if available	yes
Application Protocol Category & Tag	yes	no	if available	yes
Application Risk	yes	no	if available	yes
Business Relevance	yes	no	if available	yes
Client	yes	no	if available	yes
Client Category & Tag	yes	no	if available	yes
Client Version	yes	no	if available	yes
Connections	yes	yes	no	yes
Count	yes	yes	yes	yes
Destination Port/ICMP Type	yes	yes	yes	yes
Device	yes	yes	yes	yes

Connection Event Field	Origin: Direct	Origin: NetFlow	Logging: Start	Logging: End
Domain	yes	yes	yes	yes
DNS Query	yes	no	yes	yes
DNS Record Type	yes	no	yes	yes
DNS Response	yes	no	yes	yes
DNS Sinkhole Name	yes	no	yes	yes
DNS TTL	yes	no	yes	yes
Egress Interface	yes	no	yes	yes
Egress Security Zone	yes	no	yes	yes
Endpoint Location	yes	no	yes	yes
Endpoint Profile	yes	no	yes	yes
Files	yes	no	no	yes
First Packet	yes	yes	yes	yes
HTTP Referrer	yes	no	no	yes
HTTP Response Code	yes	no	yes	yes
Ingress Interface	yes	no	yes	yes
Ingress Security Zone	yes	no	yes	yes
Initiator Bytes	yes	yes	not useful	yes
Initiator Country	yes	no	yes	yes
Initiator IP	yes	yes	yes	yes
Initiator Packets	yes	yes	not useful	yes
Initiator User	yes	yes	yes	yes
Intrusion Events	yes	no	no	yes
Intrusion Policy	yes	no	yes	yes
IOC (Indication of Compromise)	yes	no	yes	yes
Last Packet	yes	yes	no	yes
NetBIOS Domain	yes	no	yes	yes
NetFlow Source/Destination Autonomous System	no	yes	no	yes

Connection Event Field	Origin: Direct	Origin: NetFlow	Logging: Start	Logging: End
NetFlow Source/Destination Prefix	no	yes	no	yes
NetFlow Source/Destination TOS	no	yes	no	yes
NetFlow SNMP Input/Output	no	yes	no	yes
Network Analysis Policy	yes	no	yes	yes
Original Client Country	yes	no	yes	yes
Original Client IP	yes	no	yes	yes
Prefilter Policy	yes	no	yes	yes
QoS-Applied Interface	yes	no	no	yes
QoS-Dropped Initiator Bytes	yes	no	no	yes
QoS-Dropped Initiator Packets	yes	no	no	yes
QoS-Dropped Responder Bytes	yes	no	no	yes
QoS-Dropped Responder Packets	yes	no	no	yes
QoS Policy	yes	no	no	yes
QoS Rule	yes	no	no	yes
Reason	yes	no	yes	yes
Referenced Host	yes	no	no	yes
Responder Bytes	yes	yes	not useful	yes
Responder Country	yes	no	yes	yes
Responder IP	yes	yes	yes	yes
Responder Packets	yes	yes	not useful	yes
Security Context (ASA only)	yes	no	yes	yes
Security Group Tag (SGT)	yes	no	yes	yes
Security Intelligence Category	yes	no	yes	yes
Source Device	yes	yes	yes	yes
Source Port/ICMP Type	yes	yes	yes	yes
SSL Certificate Status	yes	no	no	yes
SSL Cipher Suite	yes	no	no	yes
SSL Flow Error	yes	no	no	yes

Connection Event Field	Origin: Direct	Origin: NetFlow	Logging: Start	Logging: End
SSL Flow Flags	yes	no	no	yes
SSL Flow Messages	yes	no	no	yes
SSL Policy	yes	no	no	yes
SSL Rule	yes	no	no	yes
SSL Session ID	yes	no	no	yes
SSL Status	yes	no	no	yes
SSL Version	yes	no	no	yes
TCP Flags	no	yes	no	yes
Time	yes	yes	no	yes
Tunnel/Prefilter Rule	yes	no	yes	yes
URL	yes	no	if available	yes
URL Category	yes	no	if available	yes
URL Reputation	yes	no	if available	yes
User Agent	yes	no	no	yes
VLAN ID	yes	no	yes	yes
Web Application	yes	no	if available	yes
Web Application Category & Tag	yes	no	if available	yes

Using Connection and Security Intelligence Event Tables

You can use the Firepower Management Center to view a table of connection or Security Intelligence events. Then, you can manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access connection graphs differs depending on the workflow you use. You can use a predefined workflow, which terminates in a table view of events. You can also create a custom workflow that displays only the information that matches your specific needs.

When you are using a connection or Security Intelligence workflow table, you can perform many common actions.

Note that when you constrain connection events on a drill-down page, the packets and bytes from identical events are summed. However, if you are using a custom workflow and did not add a **Count** column to a drill-down page, the events are listed individually and packets and bytes are not summed.

Note that **Connection Events** table view displays **1 of Many** instead of how many pages of events are available if your system generates more than 25 connection events.

Before you begin

You must be an Admin or Security Analyst user to perform this task.

Procedure

Step 1 Choose either of the following:

- **Analysis** > **Connections** > **Events** (for connection events)
- **Analysis** > **Connections** > **Security Intelligence Events**

Note If a connection graph appears instead of a table, click **(switch workflow)** by the workflow title, and choose the predefined **Connection Events** workflow, or a custom workflow. Note that all predefined connection event workflows—including connection graphs—terminate in a table view of connections.

Step 2 You have the following choices:

- Time Range — To adjust the time range, which is useful if no events appear, see Changing the Time Window, on page 2168.
- Field Names — To learn more about the contents of the columns in the table, see Connection and Security Intelligence Event Fields, on page 2225.

  Tip In the table view of events, several fields are hidden by default, including the Category and Tag fields for each type of application, NetFlow-related fields, TLS/SSL-related fields, and others. To show a hidden field in an event view, expand the search constraints, then click the field name under **Disabled Columns**.

- Host Profile — To view the host profile for an IP address, click **Host Profile** or, for hosts with active indications of compromise (IOC) tags, **Compromised Host** that appears next to the IP address.
- User Profile — To view user identity information, click the user icon that appears next to the **User Identity**, or for users associated with IOCs,**Red User**.
- Files and Malware —To view the files, including malware, detected or blocked in a connection, click **View Files** and proceed as described in Viewing Files and Malware Detected in a Connection, on page 2247.
- Intrusion Events — To view the intrusion events associated with a connection, as well as their priority and impact, click **Intrusion Events** in the **Intrusion Events** column and proceed as described in Viewing Intrusion Events Associated with a Connection, on page 2248.

  Tip To quickly view intrusion, file, or malware events associated with one or more connections, check the connections using the check boxes in the table, then choose the appropriate option from the **Jump to** drop-down list. Note that because they are blocked before access control rule evaluation, there can be no files or intrusions associated with connections blocked by Security Intelligence. You can only see this information for a Security Intelligence event if you configured Security Intelligence to monitor, rather than block, connections.

- Certificate — To view details about an available certificate used to encrypt a connection, click **Enabled Lock** in the **SSL Status** column.
- Constrain — To constrain the columns that appear, click **Close** ( ) in the column heading that you want to hide. In the pop-up window that appears, click **Apply**.

  Tip To hide or show other columns, check or clear the appropriate check boxes before you click **Apply**. To add a disabled column back to the view, expand the search constraints, then click the column name under Disabled Columns.

- Delete Events — To delete some or all items in the current constrained view, check the check boxes next to items you want to delete and click **Delete** or click **Delete All**.
- Drill Down — See Using Drill-Down Pages, on page 2158.

  Tip To drill down using one of several Monitor rules that matched a logged connection, click an *N* **Monitor Rules** value. In the pop-up window that appears, click the Monitor rule you want to use to constrain connection events.

- Navigate This Page — See Workflow Page Traversal Tools, on page 2155.
- Navigate Between Pages — To navigate between pages in the current workflow, keeping the current constraints, click the appropriate page link at the top left of the workflow page.
- Navigate Between Event Views — To navigate to other event views to view associated events, click **Jump to** and choose the event view from the drop-down list.
- Sort — To sort data in a workflow, click the column title. Click the column title again to reverse the sort order.

Related Topics

Overview: Workflows, on page 2141
Configuring Event View Settings, on page 33

Viewing Files and Malware Detected in a Connection

If you associate a file policy with one or more access control rules, the system can detect files (including malware) in matching traffic. Use the Analysis > Connections menu options to see the file events, if any, associated with the connections logged by those rules. Instead of a list of files, the Firepower Management Center displays view files ( ) in the **Files** column. The number on the view files indicates the number of files (including malware files) detected or blocked in that connection.

Not all file and malware events are associated with connections. Specifically:

- Malware events detected by AMP for Endpoints ("endpoint-based malware events") are not associated with connections. Those events are imported from your AMP for Endpoints deployment.
- Many IMAP-capable email clients use a single IMAP session, which ends only when the user exits the application. Although long-running connections are logged by the system, files downloaded in the session are not associated with the connection until the session ends.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

You must be an Admin or Security Analyst user to perform this task.

Procedure

Step 1 Go to **Analysis > Connections** and choose the relevant option.

Step 2 While using a connection event table, click **View Files**.
A pop-up window appears with a list of the files detected in the connection as well as their types, and if applicable, their malware dispositions.

Step 3 You have the following choices:

- View — To view a table view of file events, click a **File's View**.
- View — To view details in a table view of malware events, click a **Malware File's View**.
- Track — To track the file's transmission through your network, click a **File's Trajectory**.
- View — To view details on all of the connection's detected file or malware events detected by AMP for Networks ("network-based malware events"), click **View File Events** or **View Malware Events**.

Viewing Intrusion Events Associated with a Connection

If you associate an intrusion policy with an access control rule or default action, the system can detect exploits in matching traffic. Use the Analysis > Connections menu options to see the intrusion events, if any, associated with logged connections, as well as their priority and impact.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

You must be an Admin or Security Analyst user to perform this task.

Procedure

Step 1 Go to **Analysis > Connections** and choose the relevant option.

Step 2 While using a connection event table, click **Intrusion Events** in the **Intrusion Events** column.

Step 3 In the pop-up window that appears, you have the following options:

- Click a **Listed Event's View** to view details in the packet view.
- Click **View Intrusion Events** to view details on all of the connection's associated intrusion events.

Encrypted Connection Certificate Details

You can use options under the Analysis > Connections menu to display the public key certificate (if available) used to encrypt a connection handled by the system. The certificate contains the following information.

Table 292: Encrypted Connection Certificate Details

Attribute	Description
Subject/Issuer Common Name	The host and domain name of the certificate subject or certificate issuer.
Subject/Issuer Organization	The organization of the certificate subject or certificate issuer.
Subject/Issuer Organization Unit	The organizational unit of the certificate subject or certificate issuer.
Not Valid Before/After	The dates when the certificate is valid.
Serial Number	The serial number assigned by the issuing CA.
Certificate Fingerprint	The SHA hash value used to authenticate the certificate.
Public Key Fingerprint	The SHA hash value used to authenticate the public key contained within the certificate.

Viewing the Connection Summary Page

The Connection Summary page is visible only to users who have custom roles that are restricted by searches on connection events and who have been granted explicit menu-based access to the Connection Summary page. This page provides graphs of the activity on your monitored network organized by different criteria. For example, the Connections over Time graph displays the total number of connections on your monitored network over the interval that you choose.

You can perform almost all the same actions on connection summary graphs that you can perform on connection graphs. However, because the graphs on the Connection Summary page are based on aggregated data, you cannot examine the individual connection events on which the graphs are based. In other words, you cannot drill down to a connection data table view from a connection summary graph.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Overview** > **Summary** > **Connection Summary**.

Step 2 From the **Select Device** list, choose the device whose summary you want to view, or choose **All** to view a summary of all devices.

Step 3 To manipulate and analyze the connection graphs, proceed as described in Using Connection Event Graphs, on page 2160.

Tip To detach a connection graph so you can perform further analysis without affecting the default time range, click **View**.

Related Topics

User Role Escalation, on page 59



CHAPTER 113

Working with Intrusion Events

The following topics describe how to work with intrusion events.

- About Intrusion Events, on page 2251
- Tools for Reviewing and Evaluating Intrusion Events, on page 2251
- License Requirements for Intrusion Events, on page 2252
- Requirements and Prerequisites for Intrusion Events, on page 2252
- Viewing Intrusion Events, on page 2252
- Intrusion Event Workflow Pages, on page 2268
- The Intrusion Events Clipboard, on page 2286
- Viewing Intrusion Event Statistics, on page 2288
- Viewing Intrusion Event Performance Graphs, on page 2290
- Viewing Intrusion Event Graphs, on page 2294

About Intrusion Events

The Firepower System can help you monitor your network for traffic that could affect the availability, integrity, and confidentiality of a host and its data. By placing managed devices on key network segments, you can examine the packets that traverse your network for malicious activity. The system has several mechanisms it uses to look for the broad range of exploits that attackers have developed.

When the system identifies a possible intrusion, it generates an *intrusion event* (sometimes called by a legacy term, "IPS event"), which is a record of the date, time, the type of exploit, and contextual information about the source of the attack and its target. For packet-based events, a copy of the packet or packets that triggered the event is also recorded. Managed devices transmit their events to the Firepower Management Center where you can view the aggregated data and gain a greater understanding of the attacks against your network assets.

You can also deploy a managed device as an inline, switched, or routed intrusion system, which allows you to configure the device to drop or replace packets that you know to be harmful.

Tools for Reviewing and Evaluating Intrusion Events

You can use the following tools to review intrusion events and evaluate whether they are important in the context of your network environment and your security policies.

- An event summary page that gives you an overview of the current activity on your managed devices

- Text-based and graphical reports that you can generate for any time period you choose; you can also design your own reports and configure them to run at scheduled intervals
- An incident-handling tool that you can use to gather event data related to an attack; you can also add notes to help you track your investigation and response
- Automated alerting that you can configure for SNMP, email, and syslog
- Automated correlation policies that you can use to respond to and remediate specific intrusion events
- Predefined and custom workflows that you can use to drill down through the data to identify the events that you want to investigate further
- External tools for managing and analyzing data. You can send data to those tools using syslog or eStreamer.

To search for a particular message string and retrieve documentation for the rule that generated an event, see https://www.snort.org/rule_docs/.

License Requirements for Intrusion Events

FTD License

Threat

Classic License

Protection

Requirements and Prerequisites for Intrusion Events

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Intrusion Admin

Viewing Intrusion Events

You view an intrusion event to determine whether there is a threat to your network security.

The initial intrusion events view differs depending on the workflow you use to access the page. You can use one of the predefined workflows, which includes one or more drill-down pages, a table view of intrusion events, and a terminating packet view, or you can create your own workflow. You can also view workflows based on custom tables, which may include intrusion events.

An event view may be slow to display if it contains a large number of IP addresses and you have enabled the **Resolve IP Addresses** event view setting.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Intrusions** > **Events**.

Step 2 You have the following choices:

- Adjust time range — Adjust the time range for the event view as described in Changing the Time Window, on page 2168.
- Change workflows — If you are using a custom workflow that does not include the table view of intrusion events, choose any of the system-provided workflows by clicking **(switch workflow)** next to the workflow title.
- Constrain — To narrow your view to the intrusion events that are important to your analysis, see Using Intrusion Event Workflows, on page 2269.
- Delete event — To delete an event from the database, click **Delete** to delete the event whose packet you are viewing or click **Delete All** to delete all the events whose packets you previously selected.
- Mark reviewed — To mark intrusion events reviewed, see Marking Intrusion Events Reviewed, on page 2265.
- View connection data —To view connection data associated with intrusion events, see Viewing Connection Data Associated with Intrusion Events, on page 2264.
- View contents — To view the contents of the columns in the table as described in Intrusion Event Fields, on page 2254.

Related Topics

Using the Intrusion Event Packet View, on page 2272

About Intrusion Event Fields

When the system identifies a possible intrusion, it generates an *intrusion event*, which is a record of the date, time, the type of exploit, and contextual information about the source of the attack and its target. For packet-based events, a copy of the packet or packets that triggered the event is also recorded.

You can view intrusion event data in the Firepower Management Center web interface at **Analysis > Intrusions > Events** or emit data from certain fields as syslog messages for consumption by an external tool. Syslog fields are indicated in the list below; fields without a listed syslog equivalent are not available in syslog messages.

When searching intrusion events, keep in mind that the information available for any individual event can vary depending on how, why, and when system logged the event. For example, only intrusion events triggered on decrypted traffic contain TLS/SSL information.

Note In the Firepower Management Center web interface, some fields in the table view of intrusion events are disabled by default. To enable a field for the duration of your session, expand the search constraints, then click the column name under **Disabled Columns**.

Intrusion Event Fields

Access Control Policy (Syslog: ACPolicy)

The access control policy associated with the intrusion policy where the intrusion, preprocessor, or decoder rule that generated the event is enabled.

Access Control Rule

The access control rule that invoked the intrusion policy that generated the event. `Default Action` indicates that the intrusion policy where the rule is enabled is not associated with a specific access control rule but, instead, is configured as the default action of the access control policy.

This field is empty if there is:

- No associated rule/default action: Intrusion inspection was associated with neither an access control rule nor the default action, for example, if the packet was examined by the intrusion policy specified to handle packets that must pass before the system can determine which rule to apply. (This policy is specified in the Advanced tab of the access control policy.)
- No associated connection event: The connection event logged for the session has been purged from the database, for example, if connection events have higher turnover than intrusion events.

Application Protocol (Syslog: ApplicationProtocol)

The application protocol, if available, which represents communications between hosts detected in the traffic that triggered the intrusion event.

Application Protocol Category and Tag

Criteria that characterize the application to help you understand the application's function.

Application Risk

The risk associated with detected applications in the traffic that triggered the intrusion event: Very High, High, Medium, Low, and Very Low. Each type of application detected in a connection has an associated risk; this field displays the highest risk of those.

Business Relevance

The business relevance associated with detected applications in the traffic that triggered the intrusion event: Very High, High, Medium, Low, and Very Low. Each type of application detected in a connection has an associated business relevance; this field displays the lowest (least relevant) of those.

Classification (Syslog: Classification)

The classification where the rule that generated the event belongs.

See a list of possible classification values in Intrusion Event Details, on page 1551.

When searching this field, enter the classification number, or all or part of the classification name or description for the rule that generated the events you want to view. You can also enter a comma-separated list of numbers, names, or descriptions. Finally, if you add a custom classification, you can also search using all or part of its name or description.

Client (Syslog: Client)

The client application, if available, which represents software running on the monitored host detected in the traffic that triggered the intrusion event.

Client Category and Tag

Criteria that characterize the application to help you understand the application's function.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows. This field is not searchable.

Destination Continent

The continent of the receiving host involved in the intrusion event.

Destination Country

The country of the receiving host involved in the intrusion event.

Destination IP (Syslog: DstIP)

The IP address used by the receiving host involved in the intrusion event.

See also A Note About Initiator/Responder, Source/Destination, and Sender/Receiver Fields, on page 2238.

Destination Port / ICMP Code (Syslog: DstPort, ICMPCode)

The port number for the host receiving the traffic. For ICMP traffic, where there is no port number, this field displays the ICMP code.

Destination User

The username associated with the Responder IP of the connection event. This host may or may not be the host receiving the exploit. This value is typically known only for users on your network.

.

See also A Note About Initiator/Responder, Source/Destination, and Sender/Receiver Fields, on page 2238.

Device

The managed device where the access control policy was deployed.

Note that the primary and secondary devices in a stacked configuration report intrusion events as if they were separate devices.

Domain

The domain of the device that detected the intrusion. This field is only present if you have ever configured the Firepower Management Center for multitenancy.

Egress Interface (Syslog: EgressInterface)

The egress interface of the packet that triggered the event. This interface column is not populated for a passive interface.

Egress Security Zone (Syslog: EgressZone)

The egress security zone of the packet that triggered the event. This security zone field is not populated in a passive deployment.

Email Attachments

The MIME attachment file name that was extracted from the MIME Content-Disposition header. To display attachment file names, you must enable the SMTP preprocessor **Log MIME Attachment Names** option. Multiple attachment file names are supported.

Email Headers

This field is a search field only.

The data that was extracted from the email header.

To associate email headers with intrusion events for SMTP traffic, you must enable the SMTP preprocessor **Log Headers** option.

Email Recipient

The address of the email recipient that was extracted from the SMTP RCPT TO command. To display a value for this field, you must enable the SMTP preprocessor **Log To Addresses** option. Multiple recipient addresses are supported.

Email Sender

The address of the email sender that was extracted from the SMTP MAIL FROM command. To display a value for this field, you must enable the SMTP preprocessor **Log From Address** option. Multiple sender addresses are supported.

Generator

The component that generated the event.

See also information about the following intrusion event fields: GID, Message, and Snort ID.

GID (Syslog Only)

Generator ID; the ID of the component that generated the event.

See also information about the following intrusion event fields: Generator, Message, and Snort ID.

HTTP Hostname

The host name, if present, that was extracted from the HTTP request Host header. Note that request packets do not always include the host name.

To associate host names with intrusion events for HTTP client traffic, you must enable the HTTP Inspect preprocessor **Log Hostname** option.

In table views, this column displays the first fifty characters of the extracted host name. You can hover your pointer over the displayed portion of an abbreviated host name to display the complete name, up to 256 bytes. You can also display the complete host name, up to 256 bytes, in the packet view.

HTTP Response Code (Syslog: HTTPResponse)

The HTTP status code sent in response to a client's HTTP request over the connection that triggered the event. It indicates the reason behind sucessful and failed HTTP request.

For more details about HTTP response codes, see RFC 2616, Section 10.

HTTP URI

The raw URI, if present, associated with the HTTP request packet that triggered the intrusion event. Note that request packets do not always include a URI.

To associate URIs with intrusion events for HTTP traffic, you must enable the HTTP Inspect preprocessor **Log URI** option.

To see the associated HTTP URI in intrusion events triggered by HTTP responses, you should configure HTTP server ports in the **Perform Stream Reassembly on Both Ports** option; note, however, that this increases resource demands for traffic reassembly.

This column displays the first fifty characters of the extracted URI. You can hover your pointer over the displayed portion of an abbreviated URI to display the complete URI, up to 2048 bytes. You can also display the complete URI, up to 2048 bytes, in the packet view.

Impact

The impact level in this field indicates the correlation between intrusion data, network discovery data, and vulnerability information.

When searching this field, do not specify impact icon colors or partial strings. For example, do not use **`blue`**, **`level 1`**, or **`0`**. Valid case-insensitive values are:

- Impact 0, Impact Level 0
- Impact 1, Impact Level 1
- Impact 2, Impact Level 2
- Impact 3, Impact Level 3
- Impact 4, Impact Level 4

Because no operating system information is available for hosts added to the network map from NetFlow data, the system cannot assign Vulnerable (impact level 1: red) impact levels for intrusion events involving those hosts. In such cases, use the host input feature to manually set the operating system identity for the hosts.

Ingress Interface (Syslog: IngressInterface)

The ingress interface of the packet that triggered the event. Only this interface column is populated for a passive interface.

Ingress Security Zone (Syslog: IngressZone)

The ingress security zone or tunnel zone of the packet that triggered the event. Only this security zone field is populated in a passive deployment.

Inline Result

In workflow and table views, this field displays one of the following:

Table 293: Inline Result Field Contents in Workflow and Table Views

This Icon	Indicates
A black down arrow	The system dropped the packet that triggered the rule.
A gray down arrow	IPS would have dropped the packet if you enabled the **Drop when Inline** intrusion policy option (in an inline deployment), or if a Drop and Generate rule generated the event while the system was pruning.
No icon (blank)	The triggered rule was not set to Drop and Generate Events

In a passive deployment, the system does not drop packets, including when an inline interface is in tap mode, regardless of the rule state or the inline drop behavior of the intrusion policy.

When searching this field, enter either of the following:

- **`dropped`** to specify whether the packet is dropped in an inline deployment.
- **`would have dropped`** to specify whether the packet would have dropped if the intrusion policy had been set to drop packets in an inline deployment.

Intrusion Policy

The intrusion policy where the intrusion, preprocessor, or decoder rule that generated the event was enabled. You can choose an intrusion policy as the default action for an access control policy, or you can associate an intrusion policy with an access control rule.

IOC (Syslog: NumIOC)

Whether the traffic that triggered the intrusion event also triggered an indication of compromise (IOC) for a host involved in the connection.

When searching this field, specify **`triggered`** or **`n/a`**.

Message (Syslog: Message)

The explanatory text for the event. For rule-based intrusion events, the event message is pulled from the rule. For decoder- and preprocessor-based events, the event message is hard coded.

The Generator and Snort IDs (GID and SID) and the SID version (Revision) are appended in parentheses to the end of each message in the format of numbers separated by colons (GID:SID:version). For example `(1:36330:2)`.

MPLS Label

The Multiprotocol Label Switching label associated with the packet that triggered the intrusion event.

Network Analysis Policy (Syslog: NAPPolicy)

The network analysis policy, if any, associated with the generation of the event.

This field displays the first fifty characters of the extracted URI. You can hover your pointer over the displayed portion of an abbreviated URI to display the complete URI, up to 2048 bytes. You can also display the complete URI, up to 2048 bytes, in the packet view.

Original Client IP

The original client IP address that was extracted from an X-Forwarded-For (XFF), True-Client-IP, or custom-defined HTTP header.

To display a value for this field, you must enable the HTTP preprocessor **Extract Original Client IP Address** option in the network analysis policy. Optionally, in the same area of the network analysis policy, you can also specify up to six custom client IP headers, as well as set the priority order in which the system selects the value for the Original Client IP event field.

Priority (Syslog: Priority)

The event priority as determined by the Cisco Talos Intelligence Group (Talos). The priority corresponds to either the value of the `priority` keyword or the value for the `classtype` keyword. For other intrusion events, the priority is determined by the decoder or preprocessor. Valid values are high, medium, and low.

Protocol (Syslog: Protocol)

In the Firepower Management Center web interface, this field is a search field only.

The name or number of the transport protocol used in the connection as listed in http://www.iana.org/assignments/protocol-numbers. This is the protocol associated with the source and destination port/ICMP column.

Reviewed By

The name of the user who reviewed the event. When searching this field, you can enter **`unreviewed`** to search for events that have not been reviewed.

Revision (Syslog Only)

The version of the signature that was used to generate the event.

See also information about the following intrusion event fields: Generator, GID, Message, SID, and Snort ID.

Security Context (Syslog: Context)

The metadata identifying the virtual firewall group through which the traffic passed. The system only populates this field for ASA FirePOWER in multiple context mode.

SID (Syslog Only)

The signature ID (also known as the Snort ID) of the rule that generated the event.

See also information about the following intrusion event fields: Generator, GID, Message, Revision, and Snort ID.

Snort ID

This field is a search field only.

(For the syslog field, see SID.)

When performing your search: Specify the Snort ID (SID) of the rule that generated the event or, optionally, specify the combination Generator ID (GID) and SID of the rule, where the GID and SID are separated with a colon (:) in the format GID:SID. You can specify any of the values in the following table:

Table 294: Snort ID Search Values

Value	Example
a single SID	10000
a SID range	10000-11000
greater than a SID	>10000
greater than or equal to a SID	>=10000
less than a SID	<10000
less than or equal to a SID	<=10000
a comma-separated list of SIDs	10000,11000,12000
a single GID:SID combination	1:10000
a comma-separated list of GID:SID combinations	1:10000,1:11000,1:12000
a comma-separated list of SIDs and GID:SID combinations	10000,1:11000,12000

The SID of the events you are viewing is listed in the Message column. For more information, see the description in this section for the Message field.

Source Continent

The continent of the sending host involved in the intrusion event.

Source Country

The country of the sending host involved in the intrusion event.

Source IP (Syslog: SrcIP)

The IP address used by the sending host involved in the intrusion event.

See also A Note About Initiator/Responder, Source/Destination, and Sender/Receiver Fields, on page 2238.

Source Port / ICMP Type (Syslog: SrcPort, ICMPType)

The port number on the sending host. For ICMP traffic, where there is no port number, this field displays the ICMP type.

Source User (Syslog: User)

The username associated with the IP address of the host that initiated the connection, which may or may not be the source host of the exploit. This user value is typically known only for users on your network.

SSL Actual Action (Syslog: SSLActualAction)

In the Firepower Management Center web interface, this field is a search field only.

The action the system applied to encrypted traffic:

Block/Block with reset

Represents blocked encrypted connections.

Decrypt (Resign)

Represents an outgoing connection decrypted using a re-signed server certificate.

Decrypt (Replace Key)

Represents an outgoing connection decrypted using a self-signed server certificate with a substituted public key.

Decrypt (Known Key)

Represents an incoming connection decrypted using a known private key.

Default Action

Indicates the connection was handled by the default action.

Do not Decrypt

Represents a connection the system did not decrypt.

Field values are displayed in the **SSL Status** field on the search workflow pages.

SSL Certificate Information

This field is a search field only.

The information stored on the public key certificate used to encrypt traffic, including:

- Subject/Issuer Common Name
- Subject/Issuer Organization
- Subject/Issuer Organization Unit
- Not Valid Before/After
- Serial Number
- Certificate Fingerprint
- Public Key Fingerprint

SSL Failure Reason

This field is a search field only.

The reason the system failed to decrypt encrypted traffic:

- Unknown
- No Match
- Success
- Uncached Session
- Unknown Cipher Suite
- Unsupported Cipher Suite
- Unsupported SSL Version
- SSL Compression Used
- Session Undecryptable in Passive Mode
- Handshake Error
- Decryption Error
- Pending Server Name Category Lookup
- Pending Common Name Category Lookup
- Internal Error
- Network Parameters Unavailable
- Invalid Server Certificate Handle
- Server Certificate Fingerprint Unavailable
- Cannot Cache Subject DN
- Cannot Cache Issuer DN
- Unknown SSL Version
- External Certificate List Unavailable
- External Certificate Fingerprint Unavailable
- Internal Certificate List Invalid
- Internal Certificate List Unavailable
- Internal Certificate Unavailable
- Internal Certificate Fingerprint Unavailable
- Server Certificate Validation Unavailable
- Server Certificate Validation Failure
- Invalid Action

Field values are displayed in the **SSL Status** field on the search workflow pages.

SSL Status

The action associated with the **SSL Actual Action** (SSL rule, default action, or undecryptable traffic action) that logged the encrypted connection.

If the system fails to decrypt an encrypted connection, it displays the **SSL Actual Action** (undecryptable traffic action) taken, as well as the **SSL Failure Reason**. For example, if the system detects traffic encrypted with an unknown cipher suite and allows it without further inspection, this field displays `Do Not Decrypt (Unknown Cipher Suite)`.

Click the **Lock icon** to view certificate details.

When searching this field, enter one or more of the **SSL Actual Action** and **SSL Failure Reason** values to view encrypted traffic the system handled or failed to decrypt.

SSL Subject/Issuer Country

This field is a search field only.

A two-character ISO 3166-1 alpha-2 country code for the subject or issuer country associated with the encryption certificate.

Time

The date and time of the event. This field is not searchable.

VLAN ID

The innermost VLAN ID associated with the packet that triggered the intrusion event.

Web Application (Syslog: WebApplication)

The web application, which represents the content or requested URL for HTTP traffic detected in the traffic that triggered the intrusion event.

If the system detects an application protocol of HTTP but cannot detect a specific web application, the system supplies a generic web browsing designation instead.

Web Application Category and Tag

Criteria that characterize the application to help you understand the application's function.

Related Topics

Event Searches, on page 2177

Intrusion Event Impact Levels

To help you evaluate the impact an event has on your network, the Firepower Management Center displays an impact level in the table view of intrusion events. For each event, the system adds an impact level icon whose color indicates the correlation between intrusion data, network discovery data, and vulnerability information.

Note Because no operating system information is available for hosts added to the network map from NetFlow data, the system cannot assign Vulnerable (impact level 1: red) impact levels for intrusion events involving those hosts. In such cases, use the host input feature to manually set the operating system identity for the hosts.

The following table describes the possible values for the impact levels.

Table 295: Impact Levels

Impact Level	Vulnerability	Color	Description
Unknown (0)	Unknown	gray	Neither the source nor the destination host is on a network that is monitored by network discovery.
Vulnerable (1)	Vulnerable	red	Either: • the source or the destination host is in the network map, and a vulnerability is mapped to the host • the source or destination host is potentially compromised by a virus, trojan, or other piece of malicious software
Potentially Vulnerable (2)	Potentially Vulnerable	orange	Either the source or the destination host is in the network map and one of the following is true: • for port-oriented traffic, the port is running a server application protocol • for non-port-oriented traffic, the host uses the protocol
Currently Not Vulnerable (3)	Currently Not Vulnerable	yellow	Either the source or the destination host is in the network map and one of the following is true: • for port-oriented traffic (for example, TCP or UDP), the port is not open • for non-port-oriented traffic (for example, ICMP), the host does not use the protocol
Unknown Target (4)	Unknown Target	blue	Either the source or destination host is on a monitored network, but there is no entry for the host in the network map.

Viewing Connection Data Associated with Intrusion Events

The system can log the connections where intrusion events are detected. Although this logging is automatic for intrusion policies associated with access control rules, you must manually enable connection logging to see associated connection data for the default action.

Viewing associated data is most useful when navigating between table views of events.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Intrusions** > **Events**.

Step 2 Choose the intrusion events using the check boxes in the table, then choose **Connections** from the **Jump to** drop-down list.

Tip You can view the intrusion events associated with particular connections in a similar way. For more information, see Inter-Workflow Navigation, on page 2173.

Related Topics

Logging for Allowed Connections, on page 2212
Using Intrusion Event Workflows, on page 2269
Using Connection and Security Intelligence Event Tables, on page 2245

Marking Intrusion Events Reviewed

If you are confident that an intrusion event is not malicious, you can mark the event reviewed.

If you have examined an intrusion event and are confident that the event does not represent a threat to your network security (for example, because you know that none of the hosts on your network are vulnerable to the detected exploit), you can mark the event reviewed. Reviewed events are stored in the event database and are included in the event summary statistics, but no longer appear in the default intrusion event pages. Your name appears as the reviewer.

In a multidomain deployment, if you mark an event reviewed, the system marks it reviewed in all domains that can view that event.

If you perform a backup and then delete reviewed intrusion events, restoring your backup restores the deleted intrusion events but does not restore their reviewed status. You view those restored intrusion events under **Intrusion Events**, not under **Reviewed Events**.

Procedure

On a page that displays intrusion events, you have two options:

- To mark one or more intrusion events from the list of events, check the check boxes next to the events and click **Review**.
- To mark all intrusion events from the list of events, click **Review All**.

Related Topics

Using Intrusion Event Workflows, on page 2269

Viewing Previously Reviewed Intrusion Events

In a multidomain deployment, if you mark an event reviewed, the system marks it reviewed in all domains that can view that event.

Procedure

Step 1 Choose **Analysis** > **Intrusions** > **Reviewed Events**.

Step 2 You have the following choices:

- Adjust the time range as described in Changing the Time Window, on page 2168.
- If you are using a custom workflow that does not include the table view of intrusion events, choose any of the system-provided workflows by clicking **(switch workflow)** next to the workflow title.
- To learn more about the events that appear, see Intrusion Event Fields, on page 2254.

Related Topics

Using Intrusion Event Workflows, on page 2269

Marking Reviewed Intrusion Events Unreviewed

You can return a reviewed event to the default intrusion events view by marking the event unreviewed.

In a multidomain deployment, if you mark an event reviewed, the system marks it reviewed in all domains that can view that event.

Procedure

On a page that displays reviewed events, you have two choices:

- To remove individual intrusion events from the list of reviewed events, check the check boxes next to specific events and click **Unreview**.
- To remove all intrusion events from the list of reviewed events, click **Unreview All**.

Preprocessor Events

Preprocessors provide two functions: performing the specified action on the packet (for example, decoding and normalizing HTTP traffic) and reporting the execution of specified preprocessor options by generating an event whenever a packet triggers that preprocessor option and the associated preprocessor rule is enabled. For example, you can enable the `Double Encoding` HTTP Inspect option and the associated preprocessor rule with the HTTP Inspect Generator (GID) 119 and the Snort ID (SID) 2 to generate an event when the preprocessor encounters IIS double-encoded traffic.

Generating events to report the execution of preprocessors helps you detect anomalous protocol exploits. For example, attackers can craft overlapping IP fragments to cause a DoS attack on a host. The IP defragmentation preprocessor can detect this type of attack and generate an intrusion event for it.

Preprocessor events differ from rule events in that the packet display does not include a detailed rule description for the event. Instead, the packet display shows the event message, the GID, SID, the packet header data, and the packet payload. This allows you to analyze the packet's header information, determine if its header options are being used and if they can exploit your system, and inspect the packet payload. After the preprocessors analyze each packet, the rules engine executes appropriate rules against it (if the preprocessor was able to defragment it and establish it as part of a valid session) to further analyze potential content-level threats and report on them.

Preprocessor Generator IDs

Each preprocessor has its own Generator ID number, or GID, that indicates which preprocessor was triggered by the packet. Some of the preprocessors also have related SIDs, which are ID numbers that classify potential attacks. This helps you analyze events more effectively by categorizing the type of event much the way a rule's Snort ID (SID) can offer context for packets triggering rules. You can list preprocessor rules by preprocessor in the Preprocessors filter group on the intrusion policy Rules page; you can also list preprocessor rules in the preprocessor and packet decoder sub-groupings in the Category filter group.



Note Events generated by standard text rules have a generator ID of 1. For shared object rules, the events have a generator ID of 3 . For both, the event's SID indicates which specific rule triggered.

The following table describes the types of events that generate each GID.

Table 296: Generator IDs

ID	Component	Description
1	Standard Text Rule	The event was generated when the packet triggered a standard text rule.
2	Tagged Packets	The event was generated by the Tag generator, which generates packets from a tagged session. This occurs when the `tag` rule option is used.
3	Shared Object Rule	The event was generated when the packet triggered a shared object rule.
102	HTTP Decoder	The decoder engine decoded HTTP data within the packet.
105	Back Orifice Detector	The Back Orifice Detector identified a Back Orifice attack associated with the packet.
106	RPC Decoder	The RPC decoder decoded the packet.
116	Packet Decoder	The event was generated by the packet decoder.
119, 120	HTTP Inspect Preprocessor	The event was generated by the HTTP Inspect preprocessor. GID 120 rules relate to server-specific HTTP traffic.
122	Portscan Detector	The event was generated by the portscan flow detector.
123	IP Defragmentor	The event was generated when a fragmented IP datagram could not be properly reassembled.
124	SMTP Decoder	The event was generated when the SMTP preprocessor detected an exploit against an SMTP verb.

ID	Component	Description
125	FTP Decoder	The event was generated when the FTP/Telnet decoder detected an exploit within FTP traffic.
126	Telnet Decoder	The event was generated when the FTP/Telnet decoder detected an exploit within telnet traffic.
128	SSH Preprocessor	The event was generated when the SSH preprocessor detected an exploit within SSH traffic.
129	Stream Preprocessor	The event was generated during stream preprocessing by the stream preprocessor.
131	DNS Preprocessor	The event was generated by the DNS preprocessor.
133	DCE/RPC Preprocessor	The event was generated by the DCE/RPC preprocessor.
134	Rule Latency Packet Latency	The event was generated when rule latency suspended (134:1) or re-enabled (134:2) a group of intrusion rules, or when the system stopped inspecting a packet because the packet latency threshold was exceeded (134:3).
135	Rate-Based Attack Detector	The event was generated when a rate-based attack detector identified excessive connections to hosts on the network.
137	SSL Preprocessor	The event was generated by the TLS/SSL preprocessor.
138, 139	Sensitive Data Preprocessor	The event was generated by the sensitive data preprocessor.
140	SIP Preprocessor	The event was generated by the SIP preprocessor.
141	IMAP Preprocessor	The event was generated by the IMAP preprocessor.
142	POP Preprocessor	The event was generated by the POP preprocessor.
143	GTP Preprocessor	The event was generated by the GTP preprocessor.
144	Modbus Preprocessor	The event was generated by the Modbus SCADA preprocessor.
145	DNP3 Preprocessor	The event was generated by the DNP3 SCADA preprocessor.
1000 - 2000	Standard Text Rule	The event was generated when the packet triggered a standard text rule.

Intrusion Event Workflow Pages

The preprocessor, decoder, and intrusion rules that are enabled in the current intrusion policy generate intrusion events whenever the traffic that you monitor violates the policy.

The Firepower System provides a set of predefined workflows, populated with event data, that you can use to view and analyze intrusion events. Each of these workflows steps you through a series of pages to help you pinpoint the intrusion events that you want to evaluate.

The predefined intrusion event workflows contain three different types of pages, or event views:

- one or more drill-down pages
- the table view of intrusion events
- a packet view

Drill-down pages generally include two or more columns in a table (and, for some drill-down views, more than one table) that allow you to view one specific type of information.

When you "drill down" to find more information for one or more destination ports, you automatically select those events and the next page in the workflow appears. In this way, drill-down tables help you reduce the number of events you are analyzing at one time.

The initial *table view* of intrusion events lists each intrusion event in its own row. The columns in the table list information such as the time, the source IP address and port, the destination IP address and port, the event priority, the event message, and more.

When you select events on a table view, instead of selecting events and displaying the next page in the workflow, you add to what are called *constraints*. Constraints are limits that you impose on the types of events that you want to analyze.

For example, if you click **Close** ( ) in any column and clear **Time** from the drop-down list, you can remove Time as one of the columns. To narrow the list of events in your analysis, you can click the link for a value in one of the rows in the table view. For example, to limit your analysis to the events generated from one of the source IP addresses (presumably, a potential attacker), click the IP address in the **Source IP Address** column.

If you select one or more rows in a table view and then click **View**, the packet view appears. A *packet view* provides information about the packet that triggered the rule or the preprocessor that generated the event. Each section of the packet view contains information about a specific layer in the packet. You can expand collapsed sections to see more information.

Note Because each portscan event is triggered by multiple packets, portscan events use a special version of the packet view.

If the predefined workflows do not meet your specific needs, you can create custom workflows that display only the information you are interested in. Custom intrusion event workflows can include drill-down pages, a table view of events, or both; the system automatically includes a packet view as the last page. You can easily switch between the predefined workflows and your own custom workflows depending on how you want to investigate events.

Using Intrusion Event Workflows

The drill-down views and table view of events share some common features that you can use to narrow a list of events and then concentrate your analysis on a group of related events.

To avoid displaying the same intrusion events on different workflow pages, the time range pauses when you click a link at the bottom of the page to display another page of events, and resumes when you click to take any other action on the subsequent page.



Tip At any point in the process, you can save the constraints as a set of search criteria. For example, if you find that over the course of a few days your network is being probed by an attacker from a single IP address, you can save your constraints during your investigation and then use them again later. You cannot, however, save compound constraints as a set of search criteria.

Procedure

Step 1 Access an intrusion event workflow using **Analysis** > **Intrusions** > **Events**.

Step 2 Optionally, constrain the number of intrusion events that appear on the event views as described in Intrusion Event Drill-Down Page Constraints, on page 2271 or Intrusion Event Table View Constraints, on page 2271.

Step 3 You have the following choices:

- To learn more about the columns that appear, see Intrusion Event Fields, on page 2254.
- To view a host's profile, click **Host Profile** that appears next to the host IP address.
- To view geolocation details, click flag that appears in the Source Country or Destination Country columns.
- To modify the time and date range for displayed events, see Changing the Time Window, on page 2168.

  Tip If no intrusion events appear on the event views, adjusting the specified time range might return results. If you specified an older time range, events in that time range might have been deleted. Adjusting the rule thresholding configuration might generate events.

  Note Events generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.

- To sort events on the current workflow page or navigate within the current workflow page, see Using Workflows, on page 2150.
- To navigate between pages in the current workflow, keeping the current constraints, click the appropriate page link at the top left of the workflow page.
- To add events to the clipboard so you can transfer them to an incident at a later time, click **Copy** or **Copy All**.
- To delete events from the event database, check the check boxes next to events you want to delete, then click **Delete**, or click **Delete All**.
- To mark events reviewed to remove them from intrusion event pages, but not the event database, see Marking Intrusion Events Reviewed, on page 2265.
- To download a local copy of the packet (a packet capture file in libpcap format) that triggered each selected event, check the check boxes next to events triggered by the packets you want to download, then click **Download Packets**, or click **Download All Packets**. Captured packets are saved in libpcap format. This format is used by several popular protocol analyzers.
- To navigate to other event views to view associated events, see Inter-Workflow Navigation, on page 2173.
- To temporarily use a different workflow, click **(switch workflow)**.

- To bookmark the current page so that you can quickly return to it, click **Bookmark This Page**.
- To view the Intrusion Events section of the Summary Dashboard, click **Dashboards**.
- To navigate to the bookmark management page, click **View Bookmarks**.
- To generate a report based on the data in the current view, see Creating a Report Template from an Event View, on page 2055.

Related Topics

Event Searches, on page 2177
Bookmarks, on page 2174

Intrusion Event Drill-Down Page Constraints

The following table describes how to use the drill-down pages.

Table 297: Constraining Events on Drill-Down Pages

To...	You can...
drill down to the next workflow page constraining on a specific value	click the value. For example, on the Destination Port workflow, to constrain the events to those with a destination of port 80, click **80/tcp** in the **DST Port/ICMP Code** column. The next page of the workflow, Events, appears and contains only port 80/tcp events.
drill down to the next workflow page constraining on selected events	select the check boxes next to the events you want to view on the next workflow page, then click **View**. For example, on the Destination Port workflow, to constrain the events to those with destination ports 20/tcp and 21/tcp, select the check boxes next to the rows for those ports and click **View**. The next page of the workflow, Events, appears and contains only port 20/tcp and 21/tcp events. Note that if you constrain on multiple rows and the table has more than one column (not including a Count column), you build what is called a compound constraint. Compound constraints ensure that you do not include more events in your constraint than you mean to. For example, if you use the Event and Destination workflow, each row that you select on the first drill-down page creates a compound constraint. If you pick event 1:100 with a destination IP address of 10.10.10.100 and you also pick event 1:200 with a destination IP address of 192.168.10.100, the compound constraint ensures that you do not also select events with 1:100 as the event type and 192.168.10.100 as the destination IP address or events with 1:200 as the event type and 10.10.10.100 as the destination IP address.
drill down to the next workflow page keeping the current constraints	click **View All**.

Intrusion Event Table View Constraints

The following table describes how to use the table view.

Table 298: Constraining Events on the Table View of Events

To...	You can...
constrain the view to events with a single attribute	click the attribute. For example, to constrain the view to events with a destination of port 80, click **80/tcp** in the **DST Port/ICMP Code** column.
remove a column from the table	click **Close** ( ) in the column heading that you want to hide. In the pop-up window that appears, click **Apply**. If you want to hide or show other columns, select or clear the appropriate check boxes before you click **Apply**. To add a disabled column back to the view, click the **expand arrow** to expand the search constraints, then click the column name under **Disabled Columns**.
view the packets associated with one or more events	either: • click the **down arrow** next to the event whose packets you want to view. • select one or more events whose packets you want to view, and, at the bottom of the page, click **View.** • at the bottom of the page, click **View All** to view the packets for all events that match the current constraints.

Using the Intrusion Event Packet View

A packet view provides information about the packet that triggered the rule that generated an intrusion event.

Tip The packet view on a Firepower Management Center does not contain packet information when the **Transfer Packet** option is disabled for the device detecting the event.

The packet view indicates why a specific packet was captured by providing information about the intrusion event that the packet triggered, including the event's time stamp, message, classification, priority, and, if the event was generated by a standard text rule, the rule that generated the event. The packet view also provides general information about the packet, such as its size.

In addition, the packet view has a section that describes each layer in the packet: data link, network, and transport, as well as a section that describes the bytes that comprise the packet. If the system decrypted the packet, you can view the decrypted bytes. You can expand collapsed sections to display detailed information.

Note Because each portscan event is triggered by multiple packets, portscan events use a special version of the packet view.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 On the table view of intrusion events, choose packets to view as described in Intrusion Event Table View Constraints, on page 2271.

Step 2 Optionally, if you chose more than one event, you can page through the packets in the packet view by using the page numbers at the bottom of the page.

Step 3 You also have the following options:

- Adjust — To modify the date and time range in the packet views, see Changing the Time Window, on page 2168.
- Clipboard — To add an event to the clipboard so you can transfer it to the incidents at a later time, click **Copy** to copy the event whose packet you are viewing or click **Copy All** to copy all the events whose packets you previously selected.
- Configure — To configure the intrusion rule that triggered the event, click the arrow next to Actions and continue as described in Configuring Intrusion Rules within the Packet View, on page 2276.
- Delete — To delete an event from the database, click **Delete** to delete the event whose packet you are viewing or click **Delete All** to delete all the events whose packets you previously selected.
- Download — To download a local copy of the packet (a packet capture file in libpcap format) that triggered the event, click **Download Packet** to save a copy of the captured packet for the event you are viewing or click **Download All Packets** to save copies of the captured packets for all the events whose packets you previously selected. The captured packet is saved in libpcap format. This format is used by several popular protocol analyzers.

  Note You cannot download a portscan packet because single portscan events are based on multiple packets; however, the portscan view provides all usable packet information. You must have at least 15% available disk space in order to download.

- Mark reviewed — To mark an event reviewed to remove it from event views, but not the event database, click **Review** to mark the event whose packet you are viewing or click **Review All** to mark all the events whose packets you previously selected. For more information, see Marking Intrusion Events Reviewed, on page 2265.
- View additional information — To expand or collapse a page section, click the arrow next to the section. For details, see Event Information Fields, on page 2273, Frame Information Fields, on page 2279, and Data Link Layer Information Fields, on page 2280.
- View network layer information — See Viewing Network Layer Information, on page 2281.
- View packet byte information — See Viewing Packet Byte Information, on page 2286.
- View transport layer information — See Viewing Transport Layer Information, on page 2283

Related Topics

Portscan Detection, on page 1793
The Intrusion Events Clipboard, on page 2286

Event Information Fields

On the packet view, you can view information about the packet in the Event Information section.

Event

The event message. For rule-based events, this corresponds to the rule message. For other events, this is determined by the decoder or preprocessor.

The ID for the event is appended to the message in the format (GID:SID:Rev). GID is the generator ID of the rules engine, the decoder, or the preprocessor that generated the event. SID is the identifier for the rule, decoder message, or preprocessor message. Rev is the revision number of the rule.

Timestamp

The time that the packet was captured, in UTC time zone.

Classification

The event classification. For rule-based events, this corresponds to the rule classification. For other events, this is determined by the decoder or preprocessor.

Priority

The event priority. For rule-based events, this corresponds to either the value of the `priority` keyword or the value for the `classtype` keyword. For other events, this is determined by the decoder or preprocessor.

Ingress Security Zone

The ingress security zone of the packet that triggered the event. Only this security zone field is populated in a passive deployment.

Egress Security Zone

The egress security zone of the packet that triggered the event. This field is not populated in a passive deployments

Domain

The domain where the managed device belongs. This field is only present if you have ever configured the Firepower Management Center for multitenancy.

Device

The managed device where the access control policy was deployed.

Note that the primary and secondary devices in a stacked configuration report intrusion events as if they were separate devices.

Security Context

The metadata identifying the virtual firewall group through which the traffic passed. Note that the system only populates this field for ASA FirePOWER in multiple context mode.

Ingress Interface

The ingress interface of the packet that triggered the event. Only this interface column is populated for a passive interface.

Egress Interface

For an inline set, the egress interface of the packet that triggered the event.

Source/Destination IP

The host IP address or domain name where the packet that triggered the event (source) originated, or the target (destination) host of the traffic that triggered the event.

Source Port/ICMP Type

Source port of the packet that triggered the event. For ICMP traffic, where there is no port number, the system displays the ICMP type.

Destination Port/ICMP Code

The port number for the host receiving the traffic. For ICMP traffic, where there is no port number, the system displays the ICMP code.

Email Headers

The data that was extracted from the email header. Note that email headers do not appear in the table view of intrusion events, but you can use email header data as a search criterion.

To associate email headers with intrusion events for SMTP traffic, you must enable the SMTP preprocessor **Log Headers** option. For rule-based events, this row appears when email data is extracted.

HTTP Hostname

The host name, if present, extracted from the HTTP request Host header. This row displays the complete host name, up to 256 bytes. You can expand the complete host name if it is longer than a single row.

To display host names, you must enable the HTTP Inspect preprocessor **Log Hostname** option.

Note that HTTP request packets do not always include a host name. For rule-based events, this row appears when the packet contains the HTTP host name or the HTTP URI.

HTTP URI

The raw URI, if present, associated with the HTTP request packet that triggered the intrusion event. This row displays the complete URI, up to 2048 bytes. You can expand the complete URI if it is longer than a single row.

To display the URI, you must enable the HTTP Inspect preprocessor **Log URI** option.

Note that HTTP request packets do not always include a URI. For rule-based events, this row appears when the packet contains the HTTP host name or the HTTP URI.

To see the associated HTTP URI in intrusion events triggered by HTTP responses, you should configure HTTP server ports in the **Perform Stream Reassembly on Both Ports** option; note, however, that this increases resource demands for traffic reassembly.

Intrusion Policy

The intrusion policy, if present, where the intrusion, preprocessor, or decoder rule that generated the intrusion event was enabled. You can choose an intrusion policy as the default action for an access control policy or associate an intrusion policy with an access control rule.

Access Control Policy

The access control policy that includes the intrusion policy where the intrusion, preprocessor, or decoder rule that generated the event is enabled.

Access Control Rule

The access control rule associated with an intrusion rule that generated the event. Default Action indicates that the intrusion policy where the rule is enabled is not associated with an access control rule but, instead, is configured as the default action of the access control policy.

Rule

For standard text rule events, the rule that generated the event.

Note that if the event is based on a shared object rule, a decoder, or a preprocessor, the rule is not available.

Because rule data may contain sensitive information about your network, administrators may toggle users' ability to view rule information in the packet view with the View Local Rules permission in the user role editor.

Actions

For standard text and custom rule events, expand **Actions** to take any of the following actions on the rule that triggered the event:

- edit the rule
- view documentation for the revision of the rule
- add a comment to the rule
- change the state of the rule
- set a threshold for the rule
- suppress the rule

Note that if the event is based on a shared object rule, a decoder, or a preprocessor, the rule is not available.

Configuring Intrusion Rules within the Packet View

Within the packet view of an intrusion event, you can take several actions on the rule that triggered the event. Note that if the event is based on a shared object rule, a decoder, or a preprocessor, the rule is not available.

Procedure

Step 1 Within the packet view of an intrusion event that was generated by an intrusion rule, expand **Actions** in the Event Information section.

Step 2 You have the following choices:

- Comment — For standard text rule events, click **Rule Comment** to add a text comment to the rule that generated the event. This allows you to provide additional context and information about the rule and the exploit or policy violation it identifies. You can also add and view rule comments in the intrusion rules editor.

- Disable — Click **Disable this rule...** to disable the rule.

  If this event is generated by a standard text rule, you can disable the rule, if necessary. You can set the rule in all policies that you can edit locally. Alternately, you can set the rule only in the current policy (that is, the policy that generated the event) if you can edit the current policy locally.

  Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by the system.

  Note You **cannot** disable shared object rules from the packet view, nor can you disable rules in the default policies.

- Drop packets — Click **Set this rule to drop the triggering packet...** to set the rule to drop packets that trigger it.

  If your managed device is deployed inline on your network, you can set the rule that triggered the event to drop packets that trigger the rule in all policies that you can edit locally. Alternately, you can set the rule only in the current policy (that is, the policy that generated the event) if you can edit the current policy locally.

  Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by the system. Note also that this option appears only when **Drop when Inline** is enabled in the current policy.

- Edit — For standard text rule events, click **Edit** to modify the rule that generated the event. If the event is based on a shared object rule, a decoder, or a preprocessor, the rule is not available.

  Note If you edit a system-provided rule (as opposed to a custom standard text rule), you actually create a new local rule. Make sure you set the local rule to generate events and also disable the original rule in the current intrusion policy. Note, however, that you **cannot** enable local rules in the default policies.

- Generate events — Click **Set this rule to generate events...** to set the rule to generate events.

  If this event is generated by a standard text rule, you can set the rule to generate events in all policies that you can edit locally. Alternately, you can set the rule only in the current policy (that is, the policy that generated the event) if you can edit the current policy locally.

  Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by the system.

  Note You **cannot** set shared object rules to generate events from the packet view, nor can you disable rules in the default policies.

- Set suppression options — Expand **Set Suppression Options** and continue as described in Setting Suppression Options within the Packet View, on page 2279.

  You can use this option to suppress the rule that triggered this event in all policies that you can edit locally. Alternately, you can suppress the rule only in the current policy (that is, the policy that generated the event) if you can edit the current policy locally.

  Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by Cisco.

- Set threshold options — Expand **Set Thresholding Options** and continue as described in Setting Threshold Options within the Packet View, on page 2278.

You can use this option to create a threshold for the rule that triggered this even in all policies that you can edit locally. Alternately, you create a threshold only for the current policy (that is, the policy that generated the event) if you can edit the current policy locally.

Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default intrusion policy provided by the system.

- View documentation — Click **View Documentation** to learn more about the rule that generated the event.

Setting Threshold Options within the Packet View

You can control the number of events that are generated per rule over time by setting the threshold options in the packet view of an intrusion event. You can set threshold options in all policies that you can edit locally or, when it can be edited locally, only in the in the current policy (that is, the policy that caused the event to be generated).

Procedure

Step 1 Within the packet view of an intrusion event that was generated by an intrusion rule, expand **Actions** in the Event Information section.

Step 2 Expand **Set Thresholding Options** and choose one of the two possible options:

- **in the current policy**
- **in all locally created policies**

Note The current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by the system.

Step 3 Choose the type of threshold you want to set:

- Click **limit** to limit notification to the specified number of event instances per time period.
- Click **threshold** to provide notification for each specified number of event instances per time period.
- Click **both** to provide notification once per time period after a specified number of event instances.

Step 4 Click the appropriate threshold to indicate whether you want the event instances tracked by **Source** or **Destination** IP address.

Step 5 In the **Count** field, enter the number of event instances you want to use as your threshold.

Step 6 In the **Seconds** field, enter a number between 1 and 86400 that specifies the time period for which event instances are tracked.

Step 7 If you want to override any current thresholds for this rule in existing intrusion policies, check the **Override any existing settings for this rule** check box.

Step 8 Click **Save Thresholding**.

Setting Suppression Options within the Packet View

You can use the suppression options to suppress intrusion events altogether, or based on the source or destination IP address. You can set suppression options in all policies that you can edit locally. Alternately, you can set suppression options only in the current policy (that is, the policy that generated the event) when the current policy can be edited locally.

Procedure

Step 1 Within the packet view of an intrusion event that was generated by an intrusion rule, expand **Actions** in the Event Information section.

Step 2 Expand **Set Suppression Options** and click one of the two possible options:

- **in the current policy**
- **in all locally created policies**

Note The current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by Cisco.

Step 3 Choose one of the following **Track By** options:

- Click **Source** to suppress events generated by packets originating from a specified source IP address.
- Click **Destination** to suppress events generated by packets going to a specified destination IP address.
- Click **Rule** to completely suppress events for the rule that triggered this event.

Step 4 In the **IP address or CIDR block** field, enter the IP address or CIDR block/prefix length you want to specify as the source or destination IP address.

Step 5 Click **Save Suppression**.

Related Topics

Firepower System IP Address Conventions, on page 15

Frame Information Fields

On the packet view, click the arrow next to **Frame** to view information about the captured frame. The packet view may display a single frame or multiple frames. Each frame provides information about an individual network packet. You would see multiple frames, for example, in the case of tagged packets or packets in reassembled TCP streams.

Frame n

The captured frame, where *n* is 1 for single-frame packets and the incremental frame number for multi-frame packets. The number of captured bytes in the frame is appended to the frame number.

Arrival Time

The date and time the frame was captured.

Time delta from previous captured frame

For multi-frame packets, the elapsed time since the previous frame was captured.

Time delta from previous displayed frame

For multi-frame packets, the elapsed time since the previous frame was displayed.

Time since reference or first frame

For multi-frame packets, the elapsed time since the first frame was captured.

Frame Number

The incremental frame number.

Frame Length

The length of the frame in bytes.

Capture Length

The length of the captured frame in bytes.

Frame is marked

Whether the frame is marked (true or false).

Protocols in frame

The protocols included in the frame.

Related Topics

The tag Keyword, on page 1640
TCP Stream Reassembly, on page 1779

Data Link Layer Information Fields

On the packet view, click the arrow next to the data link layer protocol (for example, **Ethernet II**) to view the data link layer information about the packet, which contains the 48-bit media access control (MAC) addresses for the source and destination hosts. It may also display other information about the packet, depending on the hardware protocol.



Note Note that this example discusses Ethernet link layer information; other protocols may also appear.

The packet view reflects the protocol used at the data link layer. The following listing describes the information you might see for an Ethernet II or IEEE 802.3 Ethernet packet in the packet view.

Destination

The MAC address for the destination host.

Note Ethernet can also use multicast and broadcast addresses as the destination address.

Source

The MAC address for the source host.

Type

For Ethernet II packets, the type of packet that is encapsulated in the Ethernet frame; for example, IPv6 or ARP datagrams. Note that this item only appears for Ethernet II packets.

Length

For IEEE 802.3 Ethernet packets, the total length of the packet, in bytes, not including the checksum. Note that this item only appears for IEEE 802.3 Ethernet packets.

Viewing Network Layer Information

Procedure

On the packet view, click the arrow next to the network layer protocol (for example, **Internet Protocol**) to view more detailed information about network layer information related to the packet.

Note Note that this example discusses IP packets; other protocols may also appear.

IPv4 Network Layer Information Fields

The following listing describes protocol-specific information that might appear in an IPv4 packet.

Version

The Internet Protocol version number.

Header Length

The number of bytes in the header, including any IP options. An IP header with no options is 20 bytes long.

Differentiated Services Field

The values for differentiated services that indicate how the sending host supports Explicit Congestion Notification (ECN):

- `0x0` — does not support ECN-Capable Transport (ECT)
- `0x1` and `0x2` — supports ECT
- `0x3` — Congestion Experienced (CE)

Total Length

The length of the IP packet, in bytes, minus the IP header.

Identification

The value that uniquely identifies an IP datagram sent by the source host. This value is used to trace fragments of the same datagram.

Flags

The values that control IP fragmentation, where:

values for the Last Fragment flag indicate whether there are more fragments associated with the datagram:

- `0` — there are no more fragments associated with the datagram
- `1` — there are more fragments associated with the datagram

values for the Don't Fragment flag control whether the datagram can be fragmented:

- `0` — the datagram can be fragmented
- `1` — the datagram must **not** be fragmented

Fragment Offset

The value for the fragment offset from the beginning of the datagram.

Time to Live (ttl)

The remaining number of hops that the datagram can make between routers before the datagram expires.

Protocol

The transport protocol that is encapsulated in the IP datagram; for example, ICMP, IGMP, TCP, or UDP.

Header Checksum

The indicator for whether the IP checksum is valid. If the checksum is invalid, the datagram may have been corrupted during transit or may be being used in an intrusion evasion attempt.

Source/Destination

The IP address or domain name for the source (or destination) host.

Note that to display the domain name, you must enable IP address resolution.

Click the address or domain name to view the context menu, then select **Whois** to do a whois search on the host, **View Host Profile** to view host information, or choose an option to add the address to a global Block list or Do-Not-Block list.

IPv6 Network Layer Information Fields

The following listing describes protocol-specific information that might appear in an IPv6 packet.

Traffic Class

An experimental 8-bit field in the IPv6 header for identifying IPv6 packet classes or priorities similar to the differentiated services functionality provided for IPv4. When unused, this field is set to zero.

Flow Label

A optional 20-bit IPv6 hexadecimal value 1 to FFFFF that identifies a special flow such as non-default quality of service or real-time service. When unused, this field is set to zero.

Payload Length

A 16-bit field identifying the number of octets in the IPv6 payload, which is comprised of all of the packet following the IPv6 header, including any extension headers.

Next Header

An 8-bit field identifying the type of header immediately following the IPv6 header, using the same values as the IPv4 Protocol field.

Hop Limit

An 8-bit decimal integer that each node that forwards the packet decrements by one. The packet is discarded if the decremented value reaches zero.

Source

The 128-bit IPv6 address for the source host.

Destination

The 128-bit IPv6 address for the destination host.

Viewing Transport Layer Information

Procedure

Step 1 On the packet view, click the arrow next to the transport layer protocol (for example, **TCP**, **UDP**, or **ICMP**).

Step 2 Optionally, click **Data** when present to view the first twenty-four bytes of the payload for the protocol immediately above it in the Packet Information section of the packet view.

Step 3 View the contents of the transport layer for TCP, UDP, and ICMP protocols as described in TCP Packet View Fields, on page 2283, UDP Packet View Fields, on page 2285, or ICMP Packet View Fields, on page 2285.

Note Note that these examples discuss TCP, UDP, and ICMP packets; other protocols may also appear.

TCP Packet View Fields

This section describes the protocol-specific information for a TCP packet.

Source port

The number that identifies the originating application protocol.

Destination port

The number that identifies the receiving application protocol.

Sequence number

The value for the first byte in the current TCP segment, keyed to initial sequence number in the TCP stream.

Next sequence number

In a response packet, the sequence number of the next packet to send.

Acknowledgement number

The TCP acknowledgement, which is keyed to the sequence number of the previously accepted data.

Header Length

The number of bytes in the header.

Flags

The six bits that indicate the TCP segment's transmission state:

- `U` — the urgent pointer is valid
- `A` — the acknowledgement number is valid
- `P` — the receiver should push data
- `R` — reset the connection
- `S` — synchronize sequence numbers to start a new connection
- `F` — the sender has finished sending data

Window size

The amount of unacknowledged data, in bytes, that the receiving host will accept.

Checksum

The indicator for whether the TCP checksum is valid. If the checksum is invalid, the datagram may have been corrupted during transit or may be being used in an in evasion attempt.

Urgent Pointer

The position, if present, in the TCP segment where the urgent data ends. Used in conjunction with the `U` flag.

Options

The values, if present, for TCP options.

UDP Packet View Fields

This section describes the protocol-specific information for a UDP packet.

Source port

The number that identifies the originating application protocol.

Destination port

The number that identifies the receiving application protocol.

Length

The combined length of the UDP header and data.

Checksum

The indicator for whether the UDP checksum is valid. If the checksum is invalid, the datagram may have been corrupted during transit.

ICMP Packet View Fields

This section describes the protocol-specific information for an ICMP packet.

Type

The type of ICMP message:

- 0 — echo reply
- 3 — destination unreachable
- 4 — source quench
- 5 — redirect
- 8 — echo request
- 9 — router advertisement
- 10 — router solicitation
- 11 — time exceeded
- 12 — parameter problem
- 13 — timestamp request
- 14 — timestamp reply
- 15 — information request (obsolete)
- 16 — information reply (obsolete)
- 17 — address mask request
- 18 — address mask reply

Code

The accompanying code for the ICMP message type. ICMP message types 3, 5, 11, and 12 have corresponding codes as described in RFC 792.

Checksum

The indicator for whether the ICMP checksum is valid. If the checksum is invalid, the datagram may have been corrupted during transit.

Viewing Packet Byte Information

Procedure

On the packet view, click the arrow next to **Packet Bytes** to view hexadecimal and ASCII versions of the bytes that comprise the packet. If the system decrypted traffic, you can view the decrypted packet bytes.

Internally Sourced Intrusion Events

Intrusion events coming from internal sources indicate a compromised host on your network. If the source IP address is on your network, this is a sign that you shoud investigate this host.

The Intrusion Events Clipboard

The clipboard is a holding area where you can copy intrusion events from any of the intrusion event views.

The contents of the clipboard are sorted by the date and time that the events were generated. After you add intrusion events to the clipboard, you can delete them from the clipboard as well as generate reports on the contents of the clipboard.

You can also add intrusion events from the clipboard to incidents, which are compilations of events that you suspect are involved in a possible violation of your security policies.

Related Topics

Using Intrusion Event Workflows, on page 2269
Using the Intrusion Event Packet View, on page 2272
Creating an Incident, on page 2132

Generating Clipboard Reports

You can generate a report for the events on the clipboard just as you would from any of the event views.

Before you begin

- Add one or more events to the clipboard as described in Using Intrusion Event Workflows, on page 2269 or Using the Intrusion Event Packet View, on page 2272.

Procedure

Step 1 Choose **Analysis** > **Intrusions** > **Clipboard**.

Step 2 You have the following options:

- To include specific events from a page on the clipboard, navigate to that page, check the check box next to the events, and click **Generate Report**.
- To include all the events from the clipboard, click **Generate Report All**.

Step 3 Specify how you want your report to look, then click **Generate**.

Step 4 Choose one or more output formats and, optionally, modify any of the other settings.

Step 5 Click **Generate**, then click **Yes**.

Step 6 You have the following choices:

- Click a report link to display the report in a new window.
- Click **OK** to return to the Report Templates page where you can modify your report design.

Related Topics

Report Templates, on page 2051

Deleting Events from the Clipboard

If you have intrusion events on the clipboard that you do not want to add to an incident, you can delete the events.

Note Deleting an event from the clipboard does **not** delete the event from the event database. However, deleting an event from the event database does delete the event from the clipboard.

Procedure

Step 1 Choose **Analysis** > **Intrusions** > **Clipboard**.

Step 2 You have the following options:

- Delete specific events — To delete specific intrusion events from a page on the clipboard, navigate to the page, check the check box next to the events, and click **Delete**.
- Delete all events — To delete all the intrusion events from the clipboard, click **Delete All**. Note that if you choose the **Confirm 'All' Actions** option in the Event Preferences, you are first prompted to confirm that you want to delete all the events.

Viewing Intrusion Event Statistics

The Intrusion Event Statistics page provides you with a quick summary of the current state of your appliance and any intrusion events generated for your network.

Each of the IP addresses, ports, protocols, event messages, and so on shown on the page is a link. Click any link to view the associated event information. For example, if one of the top 10 destination ports is `80 (http)/tcp`, clicking that link displays the first page in the default intrusion events workflow, and lists the events targeting that port. Note that only the events (and the managed devices that generate events) in the current time range appear. Also, intrusion events that you have marked reviewed continue to appear in the statistics. For example, if the current time range is the past hour but the first event was generated five hours ago, when you click the **First Event** link, the resulting event pages will not show the event until you change the time range.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Overview** > **Summary** > **Intrusion Event Statistics**.

Step 2 From the two selection boxes at the top of the page, choose the zones and devices whose statistics you want to view, or choose **All Security Zones** and **All Devices** to view statistics for all the devices that are collecting intrusion events.

Step 3 Click **Get Statistics**.

Tip To view data from a custom time range, click the link in the upper right page area and follow the directions in Changing the Time Window, on page 2168.

Host Statistics

The Host Statistics section of the Intrusion Event Statistics page provides information about the appliance itself. On the Firepower Management Center, this section also provides information about any managed devices.

This information includes the following:

Time

The current time on the appliance.

Uptime

The number of days, hours, and minutes since the appliance itself was restarted. On the Firepower Management Center, the uptime also shows the last time each managed device was rebooted, the number of users logged in, and the load average.

Disk Usage

The percentage of the disk that is being used.

Memory Usage

The percentage of system memory that is being used.

Load Average

The average number of processes in the CPU queue for the past 1 minute, 5 minutes, and 15 minutes.

Event Overview

The Event Overview section of the Intrusion Event Statistics page provides an overview of the information in the intrusion event database.

These statistics include the following:

Events

The number of events in the intrusion event database.

Events in Time Range

The currently selected time range as well as the number and percentage of events from the database that fall within the time range.

First Event

The event message for the first event in the event database.

Last Event

The event message for the last event in the event database.



Note If you select a managed device while viewing intrusion event data on the Firepower Management Center, the Event Overview section for that device appears instead.

Event Statistics

The Event Statistics section of the Intrusion Event Statistics page provides more specific information about of the information in the intrusion event database.

This information includes details on:

- the top 10 event types
- the top 10 source IP addressees
- the top 10 destination IP addresses
- the top 10 destination ports
- the protocols, ingress and egress security zones, and devices with the greatest number of events

Note In a multidomain deployment, the system builds a separate network map for each leaf domain. As a result, a leaf domain can contain an IP address that is unique within its network, but identical to an IP address in another leaf domain. When you view event statistics in an ancestor domain, the system may display multiple instances of that repeated IP address. At first glance, they might appear to be duplicate entries. However, if you drill down to the host profile information for each IP address, the system shows that they belong to different leaf domains.

Viewing Intrusion Event Performance Graphs

The intrusion event performance page allows you to generate graphs that depict performance statistics for intrusion events over a specific period of time for a Firepower Management Center or a managed device. Graphs can be generated to reflect number of intrusion events per second, number of megabits per second, average number of bytes per packet, the percent of packets uninspected by Snort, and the number of packets blocked as the result of TCP normalization. These graphs can show statistics for the last hour, last day, last week, or last month of operation.

Note New data is accumulated for statistics graphs every five minutes. Therefore, if you reload a graph quickly, the data may not change until the next five-minute increment occurs. Each graph displays *average* values in the intervals shown (day, hour, or five minutes) for the selected time period (last month, week, day, or hour). Decimal values are displayed when the average is less than one.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Overview** > **Summary** > **Intrusion Event Performance**.

Step 2 From the **Select Device** list, choose the devices whose data you want to view.

Step 3 From the **Select Graph(s)** list, choose the type of graph you want to create as described in Intrusion Event Performance Statistics Graph Types, on page 2290.

Step 4 From the **Select Time Range** list, choose the time range you would like to use for the graph.

Step 5 Click **Graph**.

Step 6 To save the graph, right-click it and follow the instructions for your browser to save the image.

Intrusion Event Performance Statistics Graph Types

The following table lists the available graph types. Note that graph types display differently if they are populated with data affected by the network analysis policy **Inline Mode** setting. If **Inline Mode** is disabled, the graph types marked with an asterisk (`*`) in the web interface (a `yes` in the column below) populate with data about the traffic the system would have modified or dropped if **Inline Mode** was enabled..

Table 299: Intrusion Event Performance Graph Types

To generate data for...	You must...	Which represents...	Affected by Inline Mode?
Avg Bytes/Packet	n/a	the average number of bytes included in each packet.	no
ECN Flags Normalized in TCP Traffic/Packet	enable **Explicit Congestion Notification** and select **Packet**	the number of packets for which ECN flags have been cleared on a per-packet basis regardless of negotiation.	yes
ECN Flags Normalized in TCP Traffic/Session	enable **Explicit Congestion Notification** and select **Stream**	the number of times that ECN flags have been cleared on a per-stream basis when ECN use was not negotiated.	yes
Events/Sec	n/a	the number of events per second generated on the device.	no
ICMPv4 Echo Normalizations	enable **Normalize ICMPv4**	the number of ICMPv4 packets for which the 8-bit Code field in Echo (Request) or Echo Reply messages were cleared.	yes
ICMPv6 Echo Normalizations	enable **Normalize ICMPv6**	the number of ICMPv6 packets for which the 8-bit Code field in Echo (Request) or Echo Reply messages was cleared.	yes
IPv4 DF Flag Normalizations	enable **Normalize IPv4** and **Normalize Don't Fragment Bit**	the number of IPv4 packets for which the single-bit Don't Fragment subfield of the IPv4 Flags header field was cleared.	yes
IPv4 Options Normalizations	enable **Normalize IPv4**	the number of IPv4 packets for which the option octet was set to 1 (No Operation).	yes
IPv4 Reserved Flag Normalizations	enable **Normalize IPv4** and **Normalize Reserved Bit**	the number of IPv4 packets for which the single-bit Reserved subfield of the IPv4 Flags header field was cleared.	yes
IPv4 Resize Normalizations	enable **Normalize IPv4**	the number of IPv4 packets with excessive-length payload that have been truncated to the datagram length specified in the IP header.	yes
IPv4 TOS Normalizations	enable **Normalize IPv4** and **Normalize TOS Bit**	the number of IPv4 packets for which the one-byte Differentiated Services (DS) field (formerly known as the Type of Service (TOS) field) was cleared.	yes
IPv4 TTL Normalizations	enable **Normalize IPv4**, **Maximum TTL**, and **Reset TTL**	the number of IPv4 Time to Live normalizations.	yes
IPv6 Options Normalizations	enable **Normalize IPv6**	the number of IPv6 packets for which the Option Type field in the Hop-by-Hop Options or Destination Options extension header was set to `00` (Skip and continue processing).	yes

To generate data for...	You must...	Which represents...	Affected by Inline Mode?
IPv6 TTL Normalizations	enable **Normalize IPv6**, **Minimum TTL**, and **Reset TTL**	the number of IPv6 Hop Limit (TTL) normalizations.	yes
Mbits/Sec	n/a	the number of megabits per second of traffic that passes through the device.	no
Packet Resized to Fit MSS Normalizations	enable **Trim Data to MSS**	the number of packets for which the payload was longer than the TCP Data field, so the payload was trimmed to the Maximum Segment Size.	yes
Packet Resized to Fit TCP Window Normalizations	enable **Trim Data to Window**	the number of packets for which the TCP Data field was trimmed to fit the receiving host's TCP window.	yes
Percent Packets Dropped	n/a	the average percentage of uninspected packets across all selected devices. For example, if you select two devices, then an average of 50% may indicate that one device has a 90% drop rate and the other has a 10% drop rate. It may also indicate that both devices have a drop rate of 50%. The graph only represents the total % drop when you select a single device.	no
RST Packets With Data Stripped Normalizations	enable **Remove Data on RST**	the number of packets for which data was removed from a TCP reset (RST) packet.	yes
SYN Packets With Data Stripped Normalizations	enable **Remove Data on SYN**	the number of packets for which data was removed from SYN packets when the TCP operating system was not Mac OS.	yes
TCP Header Padding Normalizations	enable **Normalize/Clear Option Padding Bytes**	the number of TCP packets in which option padding bytes were set to 0.	yes
TCP No Option Normalizations	enable **Allow These TCP Options** and set to an option other than `any`	the number of packets from which the Time Stamp option was stripped.	yes
TCP NS Flag Normalizations	enable **Explicit Congestion Notification** and select **Packet**	the number of ECN Nonce Sum (NS) option normalizations.	yes
TCP Options Normalizations	enable **Allow These TCP Options** and set to an option other than `any`	the number of options (excluding MSS, Window Scale, Time Stamp, and explicitly allowed options) for which the option field is set to No Operation (TCP Option 1).	yes
TCP Packets Blocked By Normalizations	enable **Normalize TCP Payload** (segment reassembly must fail)	the number of packets dropped because the TCP segments could not be properly reassembled.	yes
TCP Reserved Flags Normalizations	enable **Normalize/Clear Reserved Bits**	the number of TCP packets where the Reserved bits have been cleared.	yes

To generate data for...	You must...	Which represents...	Affected by Inline Mode?
TCP Segment Reassembly Normalizations	enable **Normalize TCP Payload** (segment reassembly must be successful)	the number of packets for which the TCP Data field was normalized to ensure consistency in retransmitted data (any segments that cannot be properly reassembled are dropped).	yes
TCP SYN Option Normalizations	enable **Allow These TCP Options** and set to an option other than `any`	the number of options for which the Maximum Segment Size or Window Scale option was set to No Operation (TCP Option 1) because the SYN control bit was not set.	yes
TCP Timestamp ECR Normalizations	enable **Allow These TCP Options** and set to an option other than `any`	the number of packets for which the Time Stamp Echo Reply (TSecr) option field was cleared because the Acknowledgment (ACK) control bit was not set.	yes
TCP Urgent Pointer Normalizations	enable **Normalize Urgent Pointer**	the number of packets for which the two-byte TCP header Urgent Pointer field was greater than the payload length and was set to the payload length.	yes
Total Blocked Packets	configure **Inline Mode** or **Drop when Inline**	the total number of dropped packets, including rule, decoder, and preprocessor drops.	no
Total Injected Packets	configure **Inline Mode**	the number of packets that were resized before being retransmitted.	no
Total TCP Filtered Packets	configure TCP Stream Preprocessing	the number of packets skipped by the stream because of TCP port filtering.	no
Total UDP Filtered Packets	configure UDP Stream Preprocessing	the number of packets skipped by the stream because of UDP port filtering.	no
Urgent Flag Cleared Normalizations	enable **Clear URG if Urgent Pointer is Not Set**	the number of packets for which the TCP header URG control bit was cleared because the urgent pointer was not set.	yes
Urgent Pointer and Urgent Flag Cleared Normalizations	enable **Clear Urgent Pointer/URG on Empty Payload**	the number of packets for which the TCP header Urgent Pointer field and the URG control bit have been cleared because there was no payload.	yes
Urgent Pointer Cleared Normalizations	enable **Clear Urgent Pointer if URG=0**	the number of packets for which the 16-bit TCP header Urgent Pointer field was cleared because the urgent (URG) control bit was not set.	yes

Related Topics

The Inline Normalization Preprocessor, on page 1762
Preprocessor Traffic Modification in Inline Deployments, on page 1683
Drop Behavior in an Inline Deployment, on page 1486

Viewing Intrusion Event Graphs

The Firepower System provides graphs that show you intrusion event trends over time. You can generate intrusion event graphs over time ranging from the last hour to the last month, for one or all managed devices.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Overview** > **Summary** > **Intrusion Event Graphs**.

Step 2 Under **Select Device**, choose **all** to include all devices, or choose the specific device you want to include in the graph.

Step 3 Under **Select Graph(s)**, choose the type of graph you want to generate:

- Top 10 Destination Ports
- Top 10 Source IP Addresses
- Top 10 Event Messages

Step 4 Under **Select Time Range**, choose the time range for the graph:

- Last Hour
- Last Day
- Last Week
- Last Month

Step 5 Click **Graph**.



CHAPTER 114

File/Malware Events and Network File Trajectory

The following topics provide an overview of file and malware events, local malware analysis, dynamic analysis, captured files, and network file trajectories.

- About File/Malware Events and Network File Trajectory, on page 2295
- File and Malware Events, on page 2296
- View Details About Analyzed Files, on page 2312
- Using Captured File Workflows, on page 2314
- Manually Submit Files for Analysis, on page 2317
- Network File Trajectory, on page 2318

About File/Malware Events and Network File Trajectory

File policies automatically generate file and malware events for matched traffic, and log captured file information. When a file policy generates a file or malware event, or captures a file, the system also automatically logs the end of the associated connection to the Firepower Management Center database. You can analyze this data to address any negative impacts and block future attacks.

Based on the file analysis results, you can review captured files and generated malware and file events using tables on pages available under the Analysis > Files menu. When available, you can examine a file's composition, disposition, threat score, and dynamic analysis summary report for further insight into the malware analysis.

To further target your analysis, you can use a malware file's *network file trajectory* (a map of how the file traversed your network, passing among hosts, as well as various file properties) to track the spread of an individual threat across hosts over time, allowing you to concentrate outbreak control and prevention efforts where most useful.

If you configure local malware analysis or dynamic analysis in a file rule, the system preclassifies files matching the rule and generates a file composition report.

If your organization has deployed *AMP for Endpoints* and integrated that deployment with your Firepower Management Center, you can also import records of scans, malware detections, and quarantines, as well as indications of compromise (IOC) identified by that product. This data is displayed alongside event data gathered by Firepower for a more complete picture of malware on your network.

The Context Explorer, dashboards, and reporting features can also aid a deeper understanding of the files and malware detected, captured, and blocked. You can also use events to trigger correlation policy violations, or alert you via email, SMTP, or syslog.

Note To configure your system to detect malware and generate file and malware events, see File Policies and Malware Protection, on page 1401.

File and Malware Events

The Firepower Management Center can log various types of file and malware events. The information available for any individual event can vary depending on how and why it was generated:

- *File events* represent files, including malware, detected by the Firepower system (AMP for Networks.) File events do not contain AMP for Endpoints-related fields.
- *Malware events* represent malware detected by either AMP for Networks or AMP for Endpoints; malware events can also record data other than threats from your AMP for Endpoints deployment, such as scans and quarantines.
- *Retrospective malware events* represent files detected by AMP for Networks whose dispositions (whether the files are malware) have changed.

Note

- Files identified as malware by AMP for Networks generate both a file event and a malware event. Malware events generated by AMP for Endpoints do not have corresponding file events.
- File events generated by inspecting NetBIOS-ssn (SMB) traffic do not immediately generate connection events because the client and server establish a persistent connection. The system generates connection events after the client or server ends the session.
- The Firepower System supports the display and input of file names that use Unicode (UTF-8) characters. However, Unicode file names appear in PDF reports in transliterated form. Additionally, the SMB protocol replaces unprintable characters in file names with periods.

File and Malware Event Types

File Events

The system logs the file events generated when a managed device detects or blocks a file in network traffic, according to the rules in currently deployed file policies.

When the system generates a file event, the system also logs the end of the associated connection to the Firepower Management Center database, regardless of the logging configuration of the invoking access control rule.

Malware Events

The Firepower system (specifically the AMP for Networks feature) generates malware events when it detects malware in network traffic as part of your overall access control configuration. Malware events contain the disposition of the resulting event and contextual data about how, where, and when the malware was detected.

Table 300: Malware Event Generation Scenarios

When the system detects a file and...	Disposition
successfully queries the AMP cloud (performs a malware cloud lookup) for the file's disposition	Malware, Clean, or Unknown
queries the AMP cloud but cannot establish a connection or the cloud is otherwise unavailable	Unavailable You may see a small percentage of events with this disposition; this is expected behavior.
the threat score associated with a file exceeds the malware threshold threat score defined in the file policy that detected the file, or local malware analysis identifies malware	Malware
it is on the custom detection list (manually marked as malware)	Custom Detection
it is on the on the clean list (manually marked as clean),	Clean

File Disposition and File Action in Malware Events

Each file rule has an associated action that determines how the system handles traffic that matches the conditions of the rule. If you select *Block Malware* or *Malware Cloud Lookup* as file rule action, the system queries the AMP cloud to determine if files traversing your network contain malware, then block files that represent threats. Cloud lookup allows you to obtain and log the file's disposition based on its SHA-256 hash value.

The following table describes the file action that associates with the file disposition returned by the AMP cloud:

Table 301: File Disposition and File Action in Malware Events

File Rule Action Selected	File Disposition	File Action in the Malware Event
• Block Malware • Malware Cloud Lookup	Malware	Block
	• Clean • Unknown • Unavailable • NA	Cloud Lookup **Note** Under the file policy editor Advanced Settings, you can set a threshold threat score for **If AMP Cloud disposition is Unknown, override disposition based upon threat score** option. If you set a threshold threat score, files with an AMP cloud verdict of Unknown are considered malware if their dynamic analysis score is equal to or worse than the threshold.

Retrospective Malware Events

For malware detected in network traffic, dispositions can change. For example, the AMP cloud can determine that a file that was previously thought to be clean is now identified as malware, or the reverse—that a malware-identified file is actually clean. When the disposition changes for a file you queried in the last week, the AMP cloud notifies the system. Then, two things happen:

- The Firepower Management Center generates a new retrospective malware event.

  This new retrospective malware event represents a disposition change for all files detected in the last week that have the same SHA-256 hash value. For that reason, these events contain limited information: the date and time the Firepower Management Center was notified of the disposition change, the new disposition, the SHA-256 hash value of the file, and the threat name. They do not contain IP addresses or other contextual information.

- The Firepower Management Center changes the file disposition for previously detected files with the retrospective event's associated SHA-256 hash value.

  If a file's disposition changes to Malware, the Firepower Management Center logs a new malware event to its database. Except for the new disposition, the information in this new malware event is identical to that in the file event generated when the file was initially detected.

  If a file's disposition changes to Clean, the Firepower Management Center does not delete the malware event. Instead, the event reflects the change in disposition. This means that files with clean dispositions can appear in the malware table, but only if they were originally thought to be malware. Files that were never identified as malware appear only in the files table.

Malware Events Generated by AMP for Endpoints

If your organization uses AMP for Endpoints, individual users install lightweight connectors on *endpoints*: computers and mobile devices. Connectors can inspect files upon upload, download, execution, open, copy, move, and so on. These connectors communicate with the AMP cloud to determine if inspected files contain malware.

When a file is positively identified as malware, the AMP cloud sends the threat identification to the Firepower Management Center. The AMP cloud can also send other kinds of information to the Firepower Management Center, including data on scans, quarantines, blocked executions, and cloud recalls. The Firepower Management Center logs this information as malware events.



Note The IP addresses reported in malware events generated by AMP for Endpoints may not be in your network map—and may not even be in your monitored network at all. Depending on your deployment, level of compliance, and other factors, endpoints in your organization monitored by AMP for Endpoints may not be the same hosts as those monitored by AMP for Networks.

Malware Event Analysis with AMP for Endpoints

If your organization has deployed Cisco AMP for Endpoints:

- You can configure the system to display malware events detected by AMP for Endpoints on Firepower Management Center event pages, alongside events detected by AMP for Networks.
- If you are using the AMP public cloud, you can view file trajectory and other information about a particular SHA in AMP for Endpoints.

To configure the above functionality, see Integrate Firepower and AMP for Endpoints, on page 1437.

Event Data from AMP for Endpoints

If your organization has deployed AMP for Endpoints for malware protection, you can configure the system to let you work in FMC with file and malware data from AMP for Endpoints.

However, you should be aware of the differences between file and malware data from AMP for Endpoints and file and malware data from AMP for Networks (malware protection using the Firepower system.)

Because AMP for Endpoints malware detection is performed at the endpoint at download or execution time, while managed devices detect malware in network traffic, the information in the two types of malware events is different. For example, malware events detected by AMP for Endpoints ("endpoint-based malware") contain information on file path, invoking client application, and so on, while malware detections in network traffic contain port, application protocol, and originating IP address information about the connection used to transmit the file.

As another example, for malware events detected by AMP for Networks ("network-based malware events"), user information represents the user most recently logged into the host where the malware was destined, as determined by network discovery. But AMP for Endpoints-reported users represent the user currently logged into the endpoint where the malware was detected.

Note Depending on your deployment, endpoints monitored by AMP for Endpoints may not be the same hosts as those monitored by AMP for Networks. For this reason, malware events generated by AMP for Endpoints do not add hosts to the network map. However, the system uses IP and MAC address data to tag monitored hosts with indications of compromise obtained from your AMP for Endpoints deployment. If two different hosts monitored by different AMP solutions have the same IP and MAC address, the system can incorrectly tag monitored hosts with AMP for Endpoints IOCs.

The following table summarizes the differences between the event data generated by Firepower when using a Malware license, and event data generated by AMP for Endpoints.

Table 302: Summary of Data Differences Between AMP Products

Feature	AMP for Networks	AMP for Endpoints
Events generated	File events, captured files, malware events, and retrospective malware events	Malware events
Information in malware events	Basic malware event information, plus connection data (IP address, port, and application protocol)	In-depth malware event information; no connection data
Network file trajectory	FMC-based	FMC and the AMP for Endpoints management console each have a network file trajectory. Both are useful.

Related Topics

Integrate Firepower and AMP for Endpoints, on page 1437

Using File and Malware Event Workflows

Use this procedure to view file and malware events in a table and to manipulate the event view depending on the information relevant to your analysis. The page you see when you access events differs depending on the workflow, which is simply a series of pages you can use to evaluate events by moving from a broad to a more focused view. You can also create a custom workflow that displays only the information that matches your specific needs.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

You must be an Admin or Security Analyst user to perform this task.

Procedure

Choose one of the following:

- **Analysis** > **Files** > **File Events**
- **Analysis** > **Files** > **Malware Events**

Tip In the table view of events, several fields are hidden by default. To show a hidden field in an event view, expand the search constraints, then click the field name under **Disabled Columns**.

Tip To quickly view the connections where specific files were detected, choose the files using the check boxes in the table, then choose **Connections Events** from the **Jump to** drop-down list.

Tip Right-click an item in the table to see options. (Not every column offers options.)

Related Topics

File and Malware Event Fields, on page 2300
Predefined File Workflows, on page 2143
Predefined Malware Workflows, on page 2143
Configuring Event View Settings, on page 33

File and Malware Event Fields

File and malware events, which you can view and search using workflows, contain the fields listed in this section. Keep in mind that the information available for any individual event can vary depending on how and why it was generated.



Note Files identified as malware by AMP for Networks generate both a file event and a malware event. Malware events generated by AMP for Endpoints do not have corresponding file events, and file events do not have AMP for Endpoints-related fields.

Action

The action associated with file policy rule that detected the file, and any associated file rule action options.

AMP Cloud

The name of the AMP cloud where the AMP for Endpoints event originated.

Application File Name

The client application accessing the malware file when AMP for Endpoints detection occurred. These applications are **not** tied to network discovery or application control.

Application File SHA256

The SHA-256 hash value of the parent file accessing the AMP for Endpoints-detected or quarantined file when detection occurred.

Application Protocol

The application protocol used by the traffic in which a managed device detected the file.

Application Protocol Category or Tag

The criteria that characterize the application to help you understand the application's function.

Application Risk

The risk associated with the application traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of application detected in the connection has an associated risk; this field displays the highest of those.

Archive Depth

The level (if any) at which the file was nested in an archive file.

Archive Name

The name of the archive file (if any) which contained the malware file.

To view the contents of an archive file, go to any table under **Analysis > Files** that lists the archive file, right-click on the archive file's table row to open the context menu, then click **View Archive Contents**.

Archive SHA256

The SHA-256 hash value of the archive file (if any) which contains the malware file.

To view the contents of an archive file, go to any table under Analysis > Files that lists the archive file, right-click on that archive file's table row to open the context menu, then click **View Archive Contents**.

Business Relevance

The business relevance associated with the application traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of application detected in the connection has an associated business relevance; this field displays the lowest (least relevant) of those.

Category / File Type Category

The general categories of file type, for example: Office Documents, Archive, Multimedia, Executables, PDF files, Encoded, Graphics, or System Files.

Client

The client application that runs on one host and relies on a server to send a file.

Client Category or Tag

The criteria that characterize the application to help you understand the application's function.

Count

After you apply a constraint that creates two or more identical rows, the number of events that match the information in each row.

Detection Name

The name of the detected malware.

Detector

The AMP for Endpoints detector that identified the malware, such as ClamAV, Spero, or SHA.

Device

For file events and for malware events generated by Firepower devices, the name of the device that detected the file.

For malware events generated by AMP for Endpoints and for retrospective malware events generated by the AMP cloud, the name of the Firepower Management Center.

Disposition / File Disposition

The file's disposition:

Malware

Indicates that the AMP cloud categorized the file as malware, local malware analysis identified malware, or the file's threat score exceeded the malware threshold defined in the file policy.

Clean

Indicates that the AMP cloud categorized the file as clean, or that a user added the file to the clean list. Clean files appear in the malware table only if they were changed to clean.

Unknown

Indicates that the system queried the AMP cloud, but the file has not been assigned a disposition; in other words, the AMP cloud has not categorized the file.

Custom Detection

Indicates that a user added the file to the custom detection list.

Unavailable

Indicates that the system could not query the AMP cloud. You may see a small percentage of events with this disposition; this is expected behavior.

N/A

Indicates a Detect Files or Block Files rule handled the file and the Firepower Management Center did not query the AMP cloud.

File dispositions appear only for files for which the system queried the AMP cloud.

Domain

For file events and for malware events generated by Firepower devices, the domain of the device that detected the file. For malware events generated by AMP for Endpoints and for retrospective malware events generated by the AMP cloud, the domain associated with the AMP cloud connection that reported the event.

This field is only present if you have ever configured the Firepower Management Center for multitenancy.

Event Subtype

The AMP for Endpoints action that led to malware detection, for example, Create, Execute, Move, or Scan.

Event Type

The sub-type of malware event.

File Name

The name of the file.

File Path

The file path of the malware file detected by AMP for Endpoints, not including the file name.

File Policy

The file policy that detected the file.

File Storage / Stored (Search Only)

The storage status of the file associated with the event:

Stored

Returns all events where the associated file is currently stored.

Stored in connection

Returns all events where the system captured and stored the associated file, regardless of whether the associated file is currently stored.

Failed

Returns all events where the system failed to store the associated file.

File Timestamp

The time and date that AMP for Endpoints detected the malware file was created.

HTTP Response Code

The HTTP status code sent in response to a client's HTTP request when a file is transferred.

IOC

Whether the malware event triggered an indication of compromise (IOC) against a host involved in the connection. When AMP for Endpoints data triggers an IOC rule, a full malware event is generated, with the type AMP IOC.

Message

Additional information associated with a malware event. For file events and for malware events generated by Firepower devices, this field is populated only for files whose disposition has changed, that is, that have an associated retrospective event.

Receiving Continent

The continent of the host receiving the file.

Receiving Country

The country of the host receiving the file.

Receiving IP

For file events and for malware events generated by Firepower devices, the IP address of the host receiving the file.

See also A Note About Initiator/Responder, Source/Destination, and Sender/Receiver Fields, on page 2238.

For malware events generated by AMP for Endpoints, the IP address of the endpoint whose connector reported the event.

Receiving Port

The destination port used by the traffic where the file was detected.

Security Context

The metadata identifying the virtual firewall group through which the traffic passed. Note that the system only displays this field when managing at least one ASA FirePOWER device that is running in multiple context mode.

Sending Continent

The continent of the host sending the file.

Sending Country

The country of the host sending the file.

Sending IP

The IP address of the host sending the file.

See also A Note About Initiator/Responder, Source/Destination, and Sender/Receiver Fields, on page 2238.

Sending Port

The source port used by the traffic where the file was detected.

SHA256 / File SHA256

The SHA-256 hash value of the file.

To have a SHA256 value, the file must have been handled by one of:

- a Detect Files file rule with **Store files** enabled
- a Block Files file rule with **Store files** enabled
- a Malware Cloud Lookup file rule
- a Block Malware file rule
- AMP for Endpoints

This column also displays a network file trajectory icon that represents the most recently detected file event and file disposition, and that links to the network file trajectory.

Size (KB) / File Size (KB)

The size of the file, in kilobytes.

Note that if the system determines the file type of a file before the file is fully received, the file size may not be calculated. In this case, this field is blank.

SSL Actual Action (Search Only)

The action the system applied to encrypted traffic:

Block or Block with reset

Represents blocked encrypted connections.

Decrypt (Resign)

Represents an outgoing connection decrypted using a re-signed server certificate.

Decrypt (Replace Key)

Represents an outgoing connection decrypted using a self-signed server certificate with a substituted public key.

Decrypt (Known Key)

Represents an incoming connection decrypted using a known private key.

Default Action

Indicates the connection was handled by the default action.

Do not Decrypt

Represents a connection the system did not decrypt.

Field values are displayed in the **SSL Status** field on the search workflow pages.

SSL Certificate Information (Search Only)

The information stored on the public key certificate used to encrypt traffic, including:

- Subject/Issuer Common Name
- Subject/Issuer Organization
- Subject/Issuer Organization Unit
- Not Valid Before/After
- Serial Number, Certificate Fingerprint
- Public Key Fingerprint

SSL Failure Reason (Search Only)

The reason the system failed to decrypt encrypted traffic:

- Unknown
- No Match
- Success
- Uncached Session
- Unknown Cipher Suite
- Unsupported Cipher Suite
- Unsupported SSL Version
- SSL Compression Used
- Session Undecryptable in Passive Mode
- Handshake Error
- Decryption Error
- Pending Server Name Category Lookup
- Pending Common Name Category Lookup
- Internal Error
- Network Parameters Unavailable
- Invalid Server Certificate Handle
- Server Certificate Fingerprint Unavailable
- Cannot Cache Subject DN

- Cannot Cache Issuer DN
- Unknown SSL Version
- External Certificate List Unavailable
- External Certificate Fingerprint Unavailable
- Internal Certificate List Invalid
- Internal Certificate List Unavailable
- Internal Certificate Unavailable
- Internal Certificate Fingerprint Unavailable
- Server Certificate Validation Unavailable
- Server Certificate Validation Failure
- Invalid Action

Field values are displayed in the **SSL Status** field on the search workflow pages.

SSL Status

The action associated with the **SSL Actual Action** (SSL rule, default action, or undecryptable traffic action) that logged the encrypted connection. The **Lock icon** links to TLS/SSL certificate details. If the certificate is unavailable (for example, for connections blocked due to TLS/SSL handshake error), the lock icon is grayed out.

If the system fails to decrypt an encrypted connection, it displays the **SSL Actual Action** (undecryptable traffic action) taken, as well as the **SSL Failure Reason**. For example, if the system detects traffic encrypted with an unknown cipher suite and allows it without further inspection, this field displays `Do Not Decrypt (Unknown Cipher Suite)`.

When searching this field, type one or more of the **SSL Actual Action** and **SSL Failure Reason** values to view encrypted traffic the system handled or failed to decrypt.

SSL Subject/Issuer Country (Search Only)

The two-character ISO 3166-1 alpha-2 country code for the subject or issuer country associated with the encryption certificate.

Threat Name

The name of the detected malware.

Threat Score

The threat score most recently associated with this file. This is a value from 0 to 100 based on the potentially malicious behaviors observed during dynamic analysis.

The threat score icon links to the Dynamic Analysis Summary report.

Time

The date and time the event was generated. This field is not searchable.

Type / File Type

The type of file, for example, HTML or MSEXE.

URI / File URI

The URI of the connection associated with the file transaction, for example, the URL from which a user downloaded the file.

User

The username associated with the IP address that initiated the connection. If this IP address is external to your network, the associated username is typically unknown.

For file events and for malware events generated by Firepower devices, this field displays the username that was determined by an identity policy or authoritative logins. In absence of an identity policy, it displays *No Authentication Required*.

For malware events generated by AMP for Endpoints, AMP for Endpoints determines user names. These users **cannot** be tied to user discovery or control. They do not appear in the Users table, nor can you view details for these users.

Web Application

The application that represents the content or requested URL for HTTP traffic detected in the connection.

Web Application Category or Tag

Criteria that characterize the application to help you understand the application's function.

Malware Event Sub-Types

The following table lists the malware event subtypes, whether a malware event generated by AMP for Networks (a "network-based malware event") or AMP for Endpoints (an "endpoint-based malware event") can have that subtype, and whether the system uses that subtype to build network file trajectories.

Table 303: Malware Event Types

Malware Event Subtype/Search Value	AMP for Networks	AMP for Endpoints	File Trajectory
Threat Detected in Network File Transfer	yes	no	yes
Threat Detected in Network File Transfer (retrospective)	yes	no	yes
Threat Detected	no	yes	yes
Threat Detected in Exclusion	no	yes	yes
Threat Quarantined	no	yes	yes
AMP IOC (Indications of compromise)	no	yes	no
Blocked Execution	no	yes	no
Cloud Recall Quarantine	no	yes	no

Malware Event Subtype/Search Value	AMP for Networks	AMP for Endpoints	File Trajectory
Cloud Recall Quarantine Attempt Failed	no	yes	no
Cloud Recall Quarantine Started	no	yes	no
Cloud Recall Restore from Quarantine	no	yes	no
Cloud Recall Restore from Quarantine Failed	no	yes	no
Cloud Recall Restore from Quarantine Started	no	yes	no
Quarantine Failure	no	yes	no
Quarantined Item Restored	no	yes	no
Quarantine Restore Failed	no	yes	no
Quarantine Restore Started	no	yes	no
Scan Completed, No Detections	no	yes	no
Scan Completed With Detections	no	yes	no
Scan Failed	no	yes	no
Scan Started	no	yes	no

Information Available in File and Malware Event Fields

The following table lists whether the system displays information for each file and malware event field.

If your organization has deployed AMP for Endpoints and integrated that product with your Firepower deployment:

- Malware events and indications of compromise (IOCs) imported from your AMP for Endpoints deployment do not contain contextual connection information, but they do include information obtained at download or execution time, such as file path, invoking client application, and so on.
- File event table views do not display AMP for Endpoints-related fields.

Table 304: Information Available in File and Malware Event Fields

Field	File Event	Malware Events Detected by the Firepower System	Retrospective Events Detected by the Firepower System	Malware Events Detected by AMP for Endpoints
Action	yes	yes	yes	no
AMP Cloud	no	no	no	yes
Application File Name	no	no	no	yes

Field	File Event	Malware Events Detected by the Firepower System	Retrospective Events Detected by the Firepower System	Malware Events Detected by AMP for Endpoints
Application File SHA256	no	no	no	yes
Application Protocol	yes	yes	no	no
Application Protocol Category or Tag	yes	yes	yes	no
Application Risk	yes	yes	yes	no
Archive Depth	yes	yes	no	yes
Archive Name	yes	yes	no	yes
Archive SHA256	yes	yes	no	yes
Business Relevance	yes	yes	yes	no
Category / File Type Category	yes	yes	no	yes
Client	yes	yes	yes	no
Client Category or Tag	yes	yes	yes	no
Count	yes	yes	yes	yes
Detection Name	no	yes	no	no
Detector	no	no	no	yes
Device	yes	yes	yes	yes
Disposition / File Disposition	yes	yes	yes	no
Domain	yes	yes	yes	yes
Event Subtype	no	no	no	yes
Event Type	no	yes	yes	yes
File Name	yes	yes	no	yes
File Path	no	no	no	yes
File Policy	yes	no	no	no
File Timestamp	no	no	no	yes
HTTP Response Code	yes	yes	no	no
IOC (Indication of Compromise)	no	yes	yes	yes
Message	yes	yes	no	yes
Receiving Continent	yes	yes	yes	no

Field	File Event	Malware Events Detected by the Firepower System	Retrospective Events Detected by the Firepower System	Malware Events Detected by AMP for Endpoints
Receiving Country	yes	yes	no	no
Receiving IP	yes	yes	no	yes
Receiving Port	yes	yes	no	no
Security Context	yes	yes	yes	yes
Sending Continent	yes	yes	yes	no
Sending Country	yes	yes	no	no
Sending IP	yes	yes	no	no
Sending Port	yes	yes	no	no
SHA256 / File SHA256	yes	yes	yes	yes
Size (KB) / File Size (KB)	yes	yes	no	yes
SSL Actual Action (search only)	yes	yes	no	no
SSL Certificate Information (search only)	yes	yes	no	no
SSL Failure Reason (search only)	yes	yes	no	no
SSL Status	yes	yes	no	no
SSL Subject/Issuer Country (search only)	yes	yes	no	no
File Storage / Stored (search only)	yes	yes	no	no
Threat Name	no	yes	yes	yes
Threat Score	yes	yes	no	no
Time	yes	yes	yes	yes
Type / File Type	yes	yes	no	yes
URI / File URI	yes	yes	no	no
User	yes	yes	no	yes
Web Application	yes	yes	yes	no
Web Application Category or Tag	yes	yes	yes	no

View Details About Analyzed Files

File Composition Report

If you configure local malware analysis or dynamic analysis, the system generates a file composition report after analyzing a file. This report allows you to further analyze files and determine whether they may carry embedded malware.

The file composition report lists file properties, any objects embedded in the file, and any detected viruses. The file composition report may also list additional information specific to that file type. When the system prunes stored files, it also prunes the associated file composition report.

To view file composition information, see Using a Network File Trajectory, on page 2322.

View File Details in AMP Private Cloud

If you have deployed an AMP private cloud, you can view additional details about analyzed files in your private cloud.

For more information, see the documentation for your private cloud.

Procedure

Sign in directly to your AMP private cloud console.

Threat Scores and Dynamic Analysis Summary Reports

Threat Scores

Table 305: Threat Score Ratings

Threat Score	Numeric Score	Icon
`Low`	0-24	**Low**
`Medium`	25-69	**Medium**
`High`	70-94	**High**
`Very High`	95-100	**Very High**

The Firepower Management Center caches a file's threat score for the same amount of time as the file's disposition. If the system later detects these files, it displays the cached threat scores instead of re-querying the Cisco Threat Grid cloud or an Cisco Threat Grid on-premises appliance. You can automatically assign a malware file disposition to any file with a threat score that exceeds the defined malware threshold threat score.

Dynamic Analysis Summary

If a dynamic analysis summary is available, you can click the threat score icon to view it. If multiple reports exist, this summary is based on the most recent report matching the exact threat score. If none match the exact threat score, the report with the highest threat score is displayed. If more than one report exists, you can select a threat score to view each separate report.

The summary lists each component threat comprising the threat score. Each component threat is expandable to list the AMP cloud findings, as well as any processes related to this component threat.

The process tree shows the processes that started when the Cisco Threat Grid cloud attempted to run the file. This can help identify whether a file that contains malware is attempting to access processes and system resources beyond what is expected (for example, running a Word document opens Microsoft Word, then starts Internet Explorer, then runs the Java Runtime Environment).

Each listed process contains a process identifier you can use to verify the actual process. Child nodes in the process tree represent processes started as a result of parent processes.

From the dynamic analysis summary, you can click **View Full Report** to view the full Analysis Report, detailing the AMP cloud's full analysis, including general file information, a more in-depth review of all detected processes, a breakdown of the file analysis, and other relevant information.

Viewing Dynamic Analysis Results in the Cisco Threat Grid Public Cloud

Cisco Threat Grid offers more detailed reporting on analyzed files than is available in the Firepower Management Center. If your organization has an account in the Cisco Threat Grid public cloud, you can access the Cisco Threat Grid portal directly to view additional details about files sent for analysis from your managed devices.

Before you begin

- Associate your Firepower Management Center with your Cisco Threat Grid public cloud account. See Enabling Access to Dynamic Analysis Results in the Public Cloud, on page 1418.
- License requirement: Malware
- You must be in the global domain for this task.
- You must have one of the following user roles: Admin, Access Admin, Network Admin

Procedure

Step 1 Access the Cisco Threat Grid public cloud portal at the address provided in your Threat Grid documentation.

Step 2 Sign in with the account credentials that you used to create the association in the prerequisites to this task.

Step 3 View files submitted by your organization, or search for a particular file using its SHA.

If you have questions, see the Threat Grid documentation.

Using Captured File Workflows

When a managed device captures a file detected in network traffic, it logs an event.

Note If a device captures a file containing malware, the device generates two events: a file event when it detects the file, and a malware event when it identifies malware.

Use this procedure to view a list of captured files in a table and manipulate the event view depending on the information relevant to your analysis. The page you see when you access captured files differs depending on the workflow, which is simply a series of pages you can use to evaluate events by moving from a broad to a more focused view. You can also create a custom workflow that displays only the information that matches your specific needs.

If the system recaptures a file after a configuration change, such as an updated file policy, it updates existing information for that file.

For example, if you configure a file policy to capture files with a **Malware Cloud Lookup** action, the system stores the file disposition and threat score along with the file. Then, if you update your file policy, and the system recaptures the same file due to a new **Detect Files** action, the system updates the file's **Last Changed** value. However, the system does not remove the existing disposition and threat score, even though you did not perform another malware cloud lookup.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

You must be an Admin or Security Analyst user to perform this task.

Procedure

Choose **Analysis** > **Files** > **Captured Files**.

Tip In the table view of events, several fields are hidden by default. To show a hidden field in an event view, expand the search constraints, then click the field name under **Disabled Columns**.

Related Topics

Captured File Fields, on page 2314
Predefined Captured File Workflows, on page 2143
Configuring Event View Settings, on page 33

Captured File Fields

The table view of captured files, which is the final page in predefined captured file workflows, and which you can add to custom workflows, includes a column for each field in the captured files table.

When searching this table keep in mind that your search results depend on the available data in the events you are searching; depending on the available data, your search constraints may not apply. For example, if a file has never been submitted for dynamic analysis, it may not have an associated threat score.

Table 306: Captured File Fields

Field	Description
Archive Inspection Status	For archive files, the status of archive inspection: • Pending indicates that the system is still inspecting the archive file and its contents. If the file passes through your system again, complete information becomes available. • Extracted indicates that the system was able to extract and inspect the archive's contents. • Failed may, in rare cases, occur if the system is unable to process an extraction. • Depth Exceeded indicates that the archive contains further nested archive files beyond the maximum allowed depth. • Encrypted indicates that the archive file's contents are encrypted and could not be inspected. • Not Inspectable indicates that the system did not extract and inspect the archive's contents. Policy rule actions, policy configuration, and corrupted files are three major reasons for this status. To view the contents of an archive file, right-click on its row in the table to bring up the context menu, then choose **View Archive Contents**.
Category	The general categories of file type, for example: Office Documents, Archive, Multimedia, Executables, PDF files, Encoded, Graphics, or System Files.
Detection Name	The name of the detected malware.
Disposition	The file's AMP for Networks disposition: • Malware indicates that local malware analysis identified malware, the AMP cloud categorized the file as malware, or that the file's threat score exceeded the malware threshold defined in the file policy. • Clean indicates that the AMP cloud categorized the file as clean, or that a user added the file to the clean list. • Unknown indicates that the system queried the AMP cloud, but the file has not been assigned a disposition; in other words, the AMP cloud has not categorized the file. • Custom Detection indicates that a user added the file to the custom detection list. • Unavailable indicates that the system could not query the AMP cloud. You may see a small percentage of events with this disposition; this is expected behavior. • N/A indicates a Detect Files or Block Files rule handled the file and the Firepower Management Center did not query the AMP cloud.
Domain	The domain where the captured file was detected. This field is only present if you have ever configured the Firepower Management Center for multitenancy.

Field	Description
Dynamic Analysis Status	One or more of the following values indicating whether the file was submitted for dynamic analysis: • Analysis Complete — file submitted for dynamic analysis that received a threat score and dynamic analysis summary report • Capacity Handled — file stored because it could not be submitted currently • Capacity Handled (Network Issue) — file stored because it could not be submitted due to a network connectivity issue • Capacity Handled (Rate Limit) — file stored because it could not be submitted due to the maximum number of submissions reached • Device Not Activated — file not submitted because the device is not activated on the on-premises Cisco Threat Grid appliance. If you see this status, contact Support. • Failure (Analysis Timeout) — file submitted for which the AMP cloud has yet to return a result • Failure (Cannot Run File) — file submitted that the AMP cloud could not run in the test environment • Failure (Network Issue) — file that did not get submitted due to a network connectivity failure • Not Sent for Analysis — file not submitted • Not Suspicious (Not Sent For Analysis) — file pre-classified as non-malware • Previously Analyzed — file with a cached threat score, indicating that it has been previously sent • Sent for Analysis — file pre-classified as malware and queued for dynamic analysis
Dynamic Analysis Status Changed	The last time the file's dynamic analysis status changed.
File Name	The most recently detected file name associated with the file’s SHA-256 hash value.
Last Changed	The last time the information associated with this file was updated.
Last Sent	The time the file was most recently submitted to the AMP cloud for dynamic analysis.
Local Malware Analysis Status	One of the following values indicating whether the system performed local malware analysis on a file: • Analysis Complete — the system inspected the file using local malware analysis and pre-classified the file • Analysis Failed — the system attempted to inspect the file using local malware analysis and failed • Manual Request Submitted — a user submitted a file for local malware analysis • Not Analyzed — the system did not inspect the file with local malware analysis
SHA256	The SHA-256 hash value of the file, as well as a network file trajectory icon representing the most recently detected file event and file disposition. To view the network file trajectory, click the trajectory icon.

Field	Description
Storage Status	Indicates whether the file is stored on a managed device: • File Stored • Not Stored (Disposition Was Pending)
Threat Score	The threat score most recently associated with this file. To view the Dynamic Analysis Summary report, click the threat score icon.
Type	The type of file; for example, HTML or MSEXE.

Stored Files Download

Once a device stores a file, as long as the Firepower Management Center can communicate with that device and it has not deleted the file, you can download the file to a local host for long-term storage and analysis, and manually analyze the file. You can download a file from any associated file event, malware event, captured file view, or the file's trajectory.

Because malware is harmful, by default, you must confirm every file download. However, you can disable the confirmation in your User Preferences.

Because files with a disposition of Unknown may contain malware, when you download a file, the system first archives the file in a `.zip` package. The `.zip` file name contains the file disposition and file type, if available, and SHA-256 hash value. You can password-protect the `.zip` file to prevent accidental unpacking. You can edit or remove the default `.zip` file password in your User Preferences.



Caution Cisco strongly recommends you do **not** download malware, as it can cause adverse consequences. Exercise caution when downloading any file, as it may contain malware. Ensure you have taken any necessary precautions to secure the download destination before downloading files.

Manually Submit Files for Analysis

When you manually submit files for analysis, the system runs local analysis, then submits these files to the cloud for dynamic analysis. However, if local analysis is not enabled in a file policy, and you manually submit a file for analysis, the file will only be sent for dynamic analysis.

In addition to executable files, you can also submit file types not eligible for automatic submission, such as .swf, .jar, and others. This allows you to more quickly analyze a broad range of files, regardless of disposition, and pinpoint the exact causes of an incident.



Note The system checks the AMP cloud for updates (no more than once a day) to the list of file types eligible for dynamic analysis and the minimum and maximum file sizes you can submit.

Depending on the situation, there are two ways to submit files for analysis:

Before you begin

In order to manually submit captured files for analysis, one or more file rules must be configured to store files. For information, see File Policies and Malware Protection, on page 1401.

Procedure

Step 1 To submit a single file for analysis:

a) Select one of the following:

- **Analysis > Files > File Events**
- **Analysis > Files > Malware Events**
- **Analysis > Files > Captured Files**

b) Click **Table View of <Event type or files>**.

c) Right-click a file in the table and select **Analyze File**.

Step 2 To submit multiple captured files for analysis (up to 25 at a time):

a) Select **Analysis > Files > Captured Files**

b) Select the checkbox beside each file to analyze.

c) Click **Analyze**.

Network File Trajectory

The network file trajectory feature maps how hosts transferred files, including malware files, across your network. A trajectory charts file transfer data, the disposition of the file, and if a file transfer was blocked or the file was quarantined. You can determine which hosts and users may have transferred malware, which hosts are at risk, and observe file transfer trends.

You can track the transmission of any file with a AMP cloud-assigned disposition. The system can use information related to detecting and blocking malware from both AMP for Networks and AMP for Endpoints to build the a trajectory.

Recently Detected Malware and Analyzed Trajectories

The Network File Trajectory List page displays the malware most recently detected on your network, as well as the files whose trajectory maps you have most recently viewed. From these lists, you can view when each file was most recently seen on the network, the file's SHA-256 hash value, name, type, current file disposition, contents (for archive files), and the number of events associated with the file.

The page also contains a search box that lets you locate files, either based on SHA-256 hash value or file name, or by the IP address of the host that transferred or received a file. After you locate a file, you can click the **File SHA256** value to view the detailed trajectory map.

Network File Trajectory Detailed View

You can trace a file through the network by viewing the detailed network file trajectory. Search for a file's SHA 256 value or click a **File SHA 256** link in the Network File Trajectory list to view details about that file.

The network file trajectory details page has three parts:

- Summary Information — A file's trajectory page displays summary information about the file, including file identification information; when the file was first seen and most recently seen on the network, and by what user; the number of related events and hosts associated with the file; and the file's current disposition. From this section, if the managed device stored the file, you can download it locally, submit the file for dynamic analysis, or add the file to a file list.
- Trajectory Map — A file's trajectory map visually tracks a file from the first detection on your network to the most recent. The map shows when hosts transferred or received the file, how often they transferred the file, and when the file was blocked or quarantined. Vertical lines between data points represent file transfers between hosts. Horizontal lines connecting the data points show a host's file activity over time.

  The map also shows how often file events occurred for the file and when the system assigned the file a disposition or retrospective disposition. You can select a data point in the map and highlight a path that traces back to the first instance the host transferred that file; this path also intersects with every occurrence involving the host as either sender or receiver of the file, and identifies the user involved.
- Related Events — The Events table lists event information for each data point in the map. Using the table and the map, you can pinpoint specific file events, hosts and users on the network that transferred or received this file, related events in the map, and other related events in a table constrained on selected values.

Network File Trajectory Summary Information

The following summary information appears at the top of the details page for a file that appears in the Network File Trajectory list.

Tip To view related file events, click a field value link. The first page in the File Events default workflow opens in a new window, displaying all file events that also contain the selected value.

Table 307: Network File Trajectory Summary Information Fields

Name	Description
Archive Contents	For inspected archive files, the number of files the archive contains.

Name	Description
Current Disposition	One of the following AMP for Networks file dispositions: • `Malware` indicates that the AMP cloud categorized the file as malware, local malware analysis identified malware, or the file's threat score exceeded the malware threshold defined in the file policy. • `Clean` indicates that the AMP cloud categorized the file as clean, or that a user added the file to the clean list. • `Unknown` indicates that the system queried the AMP cloud, but the file has not been assigned a disposition; in other words, the AMP cloud has not categorized the file. • `Custom Detection` indicates that a user added the file to the custom detection list. • `Unavailable` indicates that the system could not query the AMP cloud. You may see a small percentage of events with this disposition; this is expected behavior. • `N/A` indicates a Detect Files or Block Files rule handled the file and the Firepower Management Center did not query the AMP cloud.
Detection Name	Name of the malware detected by local malware analysis.
Event Count	The number of events seen on the network associated with the file, and the number of events displayed in the map if there are more than 250 detected events.
File Category	The general categories of file type, for example, `Office Documents` or `System Files`.
File Names	The names of the file associated with the event, as seen on the network. If multiple file names are associated with a SHA-256 hash value, the most recent detected file name is listed. You can expand this to view the remaining file names by clicking `more`.
File SHA256	The SHA-256 hash value of the file. The hash is displayed by default in a condensed format. To view the full hash value, hover your pointer over it. If multiple SHA-256 hash values are associated with a file name, hover your pointer over the link to view all of the hash values.
File Size (KB)	The size of the file, in kilobytes.
File Type	The file type of the file, for example, `HTML` or `MSEXE`.
First Seen	The first time AMP for Networks or AMP for Endpoints detected the file, as well as the IP address of the host that first uploaded the file and identifying information for the user involved.
Last Seen	The most recent time AMP for Networks or AMP for Endpoints detected the file, as well as the IP address of the host that last downloaded the file and identifying information for the user involved.
Parent Application	The client application accessing the malware file when detection occurred by AMP for Endpoints. These applications are **not** tied to network discovery or application control.
Seen On	The number of hosts that either sent or received the file. Because one host can upload and download a file at different times, the total number of hosts may not match the total number of senders plus the total number of receivers in the `Seen On Breakdown` field.

Name	Description
Seen On Breakdown	The number of hosts that sent the file, followed by the number of hosts that received the file.
Threat Name	Name of the threat associated with the detected malware by AMP for Endpoints.
Threat Score	The file's threat score.

Network File Trajectory Map and Related Events List

The file trajectory map's y-axis contains a list of all host IP addresses that have interacted with the file. The IP addresses are listed in descending order based on when the system first detected the file on that host. Each row contains all events associated with that IP address, whether a single file event, file transfer, or retrospective event. The x-axis contains the date and time the system detected each event. The timestamps are listed in chronological order. If multiple events occurred within a minute, all are listed within the same column. You can scroll the map horizontally and vertically to view additional events and IP addresses.

The map displays up to 250 events associated with the file SHA-256 hash. If there are more than 250 events, the map displays the first 10, then truncates extra events with an **Arrow**. The map then displays the remaining 240 events.

The first page of the File Events default workflow appears in a new window with all the extra events constrained based on the file type. If malware events generated by AMP for Endpoints are not displayed, you must switch to the Malware Events table to view these.

Each data point represents an event plus the file disposition, as described in the legend below the map. For example, a Malware Block event icon combines the Malicious Disposition icon and the Block Event icon.

Malware events generated by AMP for Endpoints ("endpoint-based malware events") include one icon. A retrospective event displays an icon in the column for each host on which the file is detected. File transfer events always include two icons, one file send icon and one file receive icon, connected by a vertical line. Arrows indicate the file transfer direction from sender to receiver.

To track a file's progress through the network, you can click any data point to highlight a path that includes all data points related to the selected data point. This includes data points associated with the following types of events:

- any file transfers in which the associated IP address was either sender or receiver
- any malware events generated by AMP for Endpoints ("endpoint-based malware events") involving the associated IP address
- if another IP address was involved, all file transfers in which that associated IP address was either sender or receiver
- if another IP address was involved, any malware events generated by AMP for Endpoints ("endpoint-based malware events") involving the other IP address

All IP addresses and timestamps associated with any highlighted data point are also highlighted. The corresponding event in the Events table is also highlighted. If a path includes truncated events, the path itself is highlighted with a dotted line. Truncated events might intersect the path, but are not displayed in the map.

Using a Network File Trajectory

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Tip If your organization has deployed AMP for Endpoints, that product also has a network file trajectory feature. To pivot from FMC to AMP for Endpoints, see Work with Event Data in the AMP for Endpoints Console, on page 2323. For details about the file trajectory feature in AMP for Endpoints, see the AMP for Endpoints documentation.

Before you begin

If you are using AMP for Networks, you need the Malware license.

You must be an Admin or Security Analyst user to perform this task.

Procedure

Step 1 Choose **Analysis** > **Files** > **Network File Trajectory**.

Tip You can also access a file's trajectory from the Context Explorer, dashboard, or event views with file information.

Step 2 Click a **File SHA 256** link in the list.

Step 3 Optionally, enter a complete SHA-256 hash value, the host IP address, or the name of a file you want to track into the search field, and press Enter.

Tip If only one result matches, the Network File Trajectory page for that file appears.

Step 4 In the Summary Information section, you can:

- Add a file to a file list — To add a file or remove a file from the clean list or custom detection list, click **Edit** ( ).
- Download a file — To download a file, click **Download** ( ), and if prompted, confirm you want to download the file. If the file is unavailable for download, this download file is dimmed.
- Report — Click threat score to view the Dynamic Analysis Summary report.
- Submit for dynamic analysis — Click **AMP Cloud** to submit the file for dynamic analysis. If the file is unavailable for submission or you cannot connect to the AMP cloud, this AMP cloud is dimmed.
- View archive contents — To view information about an archive file's contents, click **View** ( ).
- View file composition — To view a file's composition, click **File List**. If the system has not generated a file composition report, this file list is dimmed.
- View captured files with same threat score — Click the threat score link to view all captured files with that threat score.

Note Cisco strongly recommends you do **not** download malware, as it can cause adverse consequences. Exercise caution when downloading any file, as it may contain malware. Ensure you have taken any necessary precautions to secure the download destination before downloading files.

Step 5 On the trajectory map, you can:

- Locate the first instance — Click an IP address to locate the first time a file event occurred involving an IP address. This highlights a path to that data point, as well as any intervening file events and IP addresses related to the first file event. The corresponding event in the Events table is also highlighted. The map scrolls to that data point if not currently visible.
- Track — Click any data point to highlight a path that includes all data points related to the selected data point, tracking a file's progress through the network.
- View hidden events — Click arrow to view all events not displayed in the File Summary event view.
- View matching file events — Hover your pointer over the **Matching File Event** to view summary information for the event. If you click any event summary information link, the first page of the File Events default workflow appears in a new window with all the extra events constrained based on the file type. The File Summary event view opens in a new window, displaying all file events that match on the criteria value you clicked.

Step 6 In the Events table, you can:

- Highlight — Choose a table row to highlight a data point in the map. The map scrolls to display the selected file event if not currently visible.
- Sort — Click the column headers to sort events in ascending or descending order.

Work with Event Data in the AMP for Endpoints Console

If your organization has deployed AMP for Endpoints, you can view malware event data in the AMP for Endpoints console, and use that application's global network file trajectory tool.

Tip For information about using AMP for Endpoints and its console, see the online help in the console or other documentation available from https://www.cisco.com/c/en/us/support/security/fireamp-endpoints/tsd-products-support-series-home.html

To access the AMP for Endpoints console from the Firepower Management Center:

Before you begin

- The connection to AMP for Endpoints must be configured (see Integrate Firepower and AMP for Endpoints, on page 1437) and Firepower Management Center must be able to connect to the AMP cloud.
- You will need your AMP for Endpoints credentials.
- You must be an Admin user to perform this task.
-

Procedure

Step 1 Choose **AMP** > **AMP Management**.

Step 2 Click the cloud name in the table.



CHAPTER 115

Using Host Profiles

The following topics describe how to use host profiles:

- Requirements and Prerequisites for Host Profiles, on page 2325
- Host Profiles, on page 2326
- Basic Host Information in the Host Profile, on page 2327
- Operating Systems in the Host Profile, on page 2329
- Servers in the Host Profile, on page 2333
- Web Applications in the Host Profile, on page 2338
- Host Protocols in the Host Profile, on page 2339
- Indications of Compromise in the Host Profile, on page 2340
- VLAN Tags in the Host Profile, on page 2340
- User History in the Host Profile, on page 2341
- Host Attributes in the Host Profile, on page 2341
- White List Violations in the Host Profile, on page 2345
- Malware Detections in the Host Profile, on page 2346
- Vulnerabilities in the Host Profile, on page 2347
- Scan Results in the Host Profile, on page 2349

Requirements and Prerequisites for Host Profiles

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Security Analyst

Host Profiles

A host profile provides a complete view of all the information the system has gathered about a single host. To access a host profile:

- navigate from any network map view.
- navigate from any event view that includes the IP addresses of hosts on monitored networks.

Host profiles provide basic information about detected hosts or devices, such as the host name or MAC addresses. Depending on your licenses and system configuration, host profiles can also provide you with the following information:

- the operating system running on a host
- the servers running on a host
- the clients and web applications running on a host
- the protocols running on a host
- the indications of compromise (IOC) tags on a host
- the VLAN tags on a host
- the last twenty-four hours of user activity on your network
- the compliance white violations associated with a host
- the most recent malware events for a host
- the vulnerabilities associated with a host
- the Nmap scan results for a host

Host attributes are also listed in the profile. You can use host attributes to classify hosts in ways that are important to your network environment. For example, you can:

- assign a host attribute that indicates the building where the host is located
- use the *host criticality* attribute to designate the business criticality of a given host and tailor correlation policies and alerts based on host criticality

From a host profile, you can view the existing host attributes applied to that host and modify the host attribute values.

If you use adaptive profile updates as part of a passive intrusion prevention deployment, you can tailor the way the system processes traffic so it best fits the type of operating system on the host and the servers and clients the host is running.

Optionally, you can perform an Nmap scan from the host profile to augment the server and operating system information in your host profile. The Nmap scanner actively probes the host to obtain information about the operating system and servers running on the host. The results of the scan are added to the list of operating system and server identities for the host.

Related Topics

Viewing Host Profiles, on page 2327

Host Profile Limitations

Unavailable Hosts

A host profile may not be available for every host on your network. Possible reasons include:

- The host was deleted from the network map because it timed out.
- You have reached your host limit.
- The host resides in a network segment that is not monitored by the network discovery policy.

Unavailable Information

The information displayed in a host profile may vary according to the type of host and the information available about the host.

For example:

- If your system detects a host using a non-IP-based protocol like STP, SNAP, or IPX, the host is added to the network map as a MAC host and much less information is available than for an IP host.
- The system can add hosts to the network map from exported NetFlow records, but the available information for these hosts is limited; see Differences between NetFlow and Managed Device Data, on page 1824.

Viewing Host Profiles

Procedure

You have two choices:

- On any network map, drill down to the IP address of the host whose profile you want to view.
- On any event view, click **Host Profile** or **Compromised Host** next to the IP address of the host whose profile you want to view.

Basic Host Information in the Host Profile

Each host profile provides basic information about a detected host or other device.

Descriptions of each of the basic host profile fields follow.

Domain

The domain associated with the host.

IP Addresses

All IP addresses (both IPv4 and IPv6) associated with the host. The system detects IP addresses associated with hosts and, where supported, groups multiple IP addresses used by the same host. IPv6 hosts often have

at least two IPv6 addresses (local-only and globally routable), and may also have IPv4 addresses. IPv4-only hosts may have multiple IPv4 addresses.

The host profile lists all detected IP addresses associated with that host. Where available, routable host IP addresses also include a flag icon and country code indicating the geolocation data associated with that address.

Note that only the first three addresses are shown by default. Click **show all** to show all addresses for a host.

Hostname

The fully qualified domain name of the host, if known.

NetBIOS Name

The NetBIOS name of the host, if available. Microsoft Windows hosts, as well as Macintosh, Linux, or other platforms configured to use NetBIOS, can have a NetBIOS name. For example, Linux hosts configured as Samba servers have NetBIOS names.

Device (Hops)

Either:

- the reporting device for the network where the host resides, as defined in the network discovery policy, or
- the device that processed the NetFlow data that added the host to the network map

The number of network hops between the device that detected the host and the host itself follows the device name, in parentheses. If multiple devices can see the host, the reporting device is displayed in bold.

If this field is blank, either:

- the host was added to the network map by a device that is not explicitly monitoring the network where the host resides, as defined in the network discovery policy, or
- the host was added using the host input feature and has not also been detected by the Firepower System.

MAC Addresses (TTL)

The host's detected MAC address or addresses and associated NIC vendors, with the NIC's hardware vendor and current time-to-live (TTL) value in parentheses.

If multiple devices detected the host, the Firepower Management Center displays all MAC addresses and TTL values associated with the host, regardless of which device reported them.

If the MAC address is displayed in bold font, the MAC address is the actual/true/primary MAC address of the host, definitively tied to the IP address by detection through ARP and DHCP traffic.

MAC addresses that are not displayed in bold font are secondary addresses, which cannot be definitively associated with the IP address of the host. For example, since the Firepower device can obtain MAC addresses only for hosts on its own network segments, if traffic originates from a network segment to which the Firepower device is not directly connected, the observed MAC address (i.e. the router MAC address) will be displayed as a secondary MAC address for the host.

Host Type

The type of device that the system detected: host, mobile device, jailbroken mobile device, router, bridge, NAT device, or load balancer.

The methods the system uses to distinguish network devices include:

- the analysis of Cisco Discovery Protocol (CDP) messages, which can identify network devices and their type (Cisco devices only)
- the detection of the Spanning Tree Protocol (STP), which identifies a device as a switch or bridge
- the detection of multiple hosts using the same MAC address, which identifies the MAC address as belonging to a router
- the detection of TTL value changes from the client side, or TTL values that change more frequently than a typical boot time, which identify NAT devices and load balancers
- The methods the system uses to distinguish mobile devices include:
- analysis of User-Agent strings in HTTP traffic from the mobile device's mobile browser
- monitoring of HTTP traffic of specific mobile applications

If a device is not identified as a network device or a mobile device, it is categorized as a host.

Last Seen

The date and time that any of a host's IP addresses was last detected.

Current User

The user most recently logged into this host.

Note that a non-authoritative user logging into a host only registers as the current user on the host if the existing current user is not an authoritative user.

View

Links to views of connection, discovery, malware, and intrusion event data, using the default workflow for that event type and constrained to show events related to the host; where possible, these events include all IP addresses associated with the host.

Operating Systems in the Host Profile

The system passively detects the identity of the operating system running on a host by analyzing the network and application stack in traffic generated by the host or by analyzing host data reported by the User Agent. The system also collates operating system information from other sources, such as the Nmap scanner or application data imported through the host input feature. The system considers the priority assigned to each identity source when determining which identity to use. By default, user input has the highest priority, followed by application or scanner sources, followed by the discovered identity.

Sometimes the system supplies a general operating system definition rather than a specific one because the traffic and other identity sources do not provide sufficient information for a more focused identity. The system collates information from the sources to use the most detailed definition possible.

Because the operating system affects the vulnerabilities list for the host and the event impact correlation for events targeting the host, you may want to manually supply more specific operating system information. In addition, you can indicate that fixes have been applied to the operating system, such as service packs and updates, and invalidate any vulnerabilities addressed by the fixes.

For example, if the system identifies a host's operating system as Microsoft Windows 2003, but you know that the host is actually running Microsoft Windows XP Professional with Service Pack 2, you can set the operating system identity accordingly. Setting a more specific operating system identity refines the list of vulnerabilities for the host, so your impact correlation for that host is more focused and accurate.

If the system detects operating system information for a host and that information conflicts with a current operating system identity that was supplied by an active source, an identity conflict occurs. When an identity conflict is in effect, the system uses both identities for vulnerabilities and impact correlation.

You can configure the network discovery policy to add discovery data to the network map for hosts monitored by NetFlow exporters. However, there is no operating system data available for these hosts, unless you set the use the host input feature to set the operating system identity.

If a host is running an operating system that violates a compliance white list in an activated network discovery policy, the Firepower Management Center marks the operating system information with the white list **Violation**. In addition, if a jailbroken mobile device violates an active white list, the icon appears next to the operating system for the device.

You can set a custom display string for the host's operating system identity. That display string is then used in the host profile.



Note Changing the operating system information for a host may change its compliance with a compliance white list.

In the host profile for a network device, the label for the Operating Systems section changes to Systems and an additional Hardware column appears. If a value for a hardware platform is listed under Systems, that system represents a mobile device or devices detected behind the network device. Note that mobile devices may or may not have hardware platform information, but hardware platform information is never detected for systems that are not mobile devices.

Descriptions of the operating system information fields displayed in the host profile follow.

Hardware

The hardware platform for a mobile device.

OS Vendor/Vendor

The operating system vendor.

OS Product/Product

One of the following values:

- the operating system determined most likely to be running on the host, based on the identity data collected from all sources
- `Pending` if the system has not yet identified an operating system and no other identity data is available

- `unknown` if the system cannot identify the operating system and no other identity data is available for the operating system

Note If the host's operating system is not one the system is capable of detecting, see Identifying Host Operating Systems, on page 1836:

OS Version/Version

The operating system version. If a host is a jailbroken mobile device, `Jailbroken` is indicated in parentheses after the version.

Source

One of the following values:

- User: user_name
- Application: app_name
- Scanner: scanner_type (Nmap or other scanner)
- Firepower

The system may reconcile data from multiple sources to determine the identity of an operating system.

Viewing Operating System Identities

You can view the specific operating system identities discovered or added for a host. The system uses source prioritization to determine the current identity for the host. In the list of identities, the current identity is highlighted by boldface text.

Note that the **View** is only available if multiple operating system identities exist for the host.

Procedure

Step 1 Click **View** in the **Operating System** or **Operating System Conflicts** section of a host profile.

Step 2 View the information described in Operating Systems in the Host Profile, on page 2329.

Step 3 Optionally, click **Delete** ( ) next to any operating system identity.

/firepower/fmc/fmc_config_guide/discovery-host-profiles/t_editing_server_identities.xml

Note You cannot delete Cisco-detected operating system identities.

This system removes the identity from the Operating System Identity Information pop-up window and, if applicable, updates the current identity for the operating system in the host profile.

Setting the Current Operating System Identity

You can set the current operating system identity for a host using the Firepower System web interface. Setting the identity through the web interface overrides all other identity sources so that identity is used for vulnerability assessment and impact correlation. However, if the system detects a conflicting operating system identity for the host after you edit the operating system, an operating system conflict occurs. Both operating systems are then considered current until you resolve the conflict.

Procedure

Step 1 Click **Edit** in the **Operating System** section of a host profile.

Step 2 You have several options:

- Choose **Current Definition** from the **OS Definition** drop-down list to confirm the current operating system identity through host input, then skip to step 6.
- Choose a variation on the current operating system identity from the **OS Definition** drop-down list, then skip to step 6.
- Choose **User-Defined** from the **OS Definition** drop-down list, then continue with step 3.

Step 3 Optionally, choose **Use Custom Display String** and modify the custom strings you want to display in the **Vendor String**, **Product String**, and **Version String** fields.

Step 4 Optionally, to change to an operating system from a different vendor, choose from the **Vendor** and **Product** drop-down lists.

Step 5 Optionally, to configure the operating system product release level, choose from the **Major**, **Minor**, **Revision**, **Build**, **Patch**, and **Extension** drop-down lists.

Step 6 Optionally, if you want to indicate that fixes for the operating system have been applied, click **Configure Fixes**.

Step 7 Choose the applicable fixes in the drop-down list, and click **Add**.

Step 8 Optionally, add the relevant patches and extensions using the **Patch** and **Extension** drop-down lists.

Step 9 Click **Finish**.

Related Topics

Operating System Identity Conflicts, on page 2332

Operating System Identity Conflicts

An operating system identity conflict occurs when a new identity detected by the system conflicts with the current identity, if that identity was provided by an active source, such as a scanner, application, or user.

The list of operating system identities in conflict displays in bold in the host profile.

You can resolve an identity conflict and set the current operating system identity for a host through the system web interface. Setting the identity through the web interface overrides all other identity sources so that identity is used for vulnerability assessment and impact correlation.

Related Topics

Configuring Network Discovery Identity Conflict Resolution, on page 1931

Making a Conflicting Operating System Identity Current

Procedure

Step 1 Navigate to the **Operating System** section of a host profile.

Step 2 You have two choices:

- Click **Make Current** next to the operating system identity you want to set as the operating system for the host.
- If the identity that you *do not* want as the current identity came from an active source, delete the unwanted identity.

Resolving an Operating System Identity Conflict

Procedure

Step 1 Click **Resolve** in the **Operating System Conflicts** section of a host profile.

Step 2 You have the following choices:

- Choose **Current Definition** from the **OS Definition** drop-down list to confirm the current operating system identity through host input, then skip to step 6.
- Choose a variation on one of the conflicting operating system identities from the **OS Definition** drop-down list, then skip to step 6.
- Choose **User-Defined** from the **OS Definition** drop-down list, then continue with step 3.

Step 3 Optionally, choose **Use Custom Display String** and enter the custom strings you want to display in the **Vendor String**, **Product String**, and **Version String** fields.

Step 4 Optionally, to change to an operating system from a different vendor, choose from the **Vendor** and **Product** drop-down lists.

Step 5 Optionally, to configure the operating system product release level, choose from the **Major**, **Minor**, **Revision**, **Build, Patch,** and **Extension** drop-down lists.

Step 6 Optionally, if you want to indicate that fixes for the operating system have been applied, click **Configure Fixes**.

Step 7 Add the fixes you have applied to the fixes list.

Step 8 Click **Finish**.

Related Topics

Configuring Network Discovery Identity Conflict Resolution, on page 1931

Servers in the Host Profile

The Servers Section of the host profile lists servers either detected on hosts on your monitored network, added from exported NetFlow records, or added through an active source like a scanner or the host input feature.

The list can include up to 100 servers per host. After that limit is reached, new server information from any source, whether active or passive, is discarded until you delete a server from the host or a server times out.

If you scan a host using Nmap, Nmap adds the results of previously undetected servers running on open TCP ports to the Servers list. If you perform an Nmap scan or import Nmap results, an expandable Scan Results section also appears in the host profile, listing the server information detected on the host by the Nmap scan. In addition, if the host is deleted from the network map, the Nmap scan results for that server for the host are discarded.

Note The system can add hosts to the network map from exported NetFlow records, but the available information for these hosts is limited; see Differences between NetFlow and Managed Device Data, on page 1824.

The process for working with servers in the host profile differs depending on how you access the profile:

- If you access the host profile by drilling down through the network map, the details for that server appear with the server name highlighted in bold. If you want to view the details for any other server on the host, click **View** ( ) next to that server name.
- If you access the host profile in any other way, expand the Servers section and click **View** ( ) next to the server whose details you want to see.

Note If the host is running a server that violates a compliance white list in an activated correlation policy, the Firepower Management Center marks the non-compliant server with the white list **Violation**.

Descriptions of the columns in the Servers list follow.

Protocol

The name of the protocol the server uses.

Port

The port where the server runs.

Application Protocol

One of:

- the name of the application protocol
- `pending` if the system cannot positively or negatively identify the application protocol for one of several reasons
- `unknown` if the system cannot identify the application protocol based on known application protocol fingerprints, or if the server was added through host input by adding a vulnerability with port information without adding a corresponding server

When you hover the mouse on an application protocol name, the tags display.

Vendor and Version

The vendor and version identified by the Firepower System, Nmap, or another active source, or acquired via the host input feature. The field is blank if none of the available sources provides an identification.

Related Topics

Host Limits and Discovery Event Logging, on page 1876
Differences between NetFlow and Managed Device Data, on page 1824
Application Detector Fundamentals, on page 1874

Server Details in the Host Profile

The Firepower Management Center lists up to 16 passively detected identities per server. Passive detection sources include network discovery data and NetFlow records. A server can have multiple passive identities if the system detects multiple vendors or versions of that server. For example, a load balancer between your managed device and your web server farm may cause your system to identify multiple passive identities for HTTP if your web servers are not running the same version of the server software. Note that the Firepower Management Center does not limit the number of server identities from active sources such as user input, scanners, or other applications.

The Firepower Management Center displays the current identity in bold. The system uses the current identity of a server for multiple purposes, including assigning vulnerabilities to a host, impact assessment, evaluating correlation rules written against host profile qualifications and compliance white lists, and so on.

The server detail may also display updated sub-server information known about the selected server.

The server detail may also display the server banner, which appears below the server details when you view a server from the host profile. Server banners provide additional information about a server that may help you identify the server. The system cannot identify or detect a misidentified server when an attacker purposely alters the server banner string. The server banner displays the first 256 bytes of the first packet detected for the server. It is collected only once, the first time the server is detected by the system. Banner content is listed in two columns, with a hexadecimal representation on the left and a corresponding ASCII representation on the right.



Note To view server banners, you must enable the **Capture Banners** check box in the network discovery policy. This option is disabled by default.

The server details section of the host profile includes the following information:

Protocol

The name of the protocol the server uses.

Port

The port where the server runs.

Hits

The number of times the server was detected by a Firepower System managed device or an Nmap scanner. The number of hits is 0 for servers imported through host input, unless the system detects traffic for that server.

Last Used

The time and date the server was last detected. The last used time for host input data reflects the initial data import time unless the system detects new traffic for that server. Scanner and application data imported through the host input feature times out according to settings in the Firepower Management Center configuration, but user input through the FMC web interface does not time out.

Application Protocol

The name of the application protocol used by the server, if known.

Vendor

The server vendor. This field does not appear if the vendor is unknown.

Version

The server version. This field does not appear if the version is unknown.

Source

One of the following values:

- User: user_name
- Application: app_name
- Scanner: scanner_type (Nmap or other scanner)
- `Firepower`, `Firepower Port Match`, or `Firepower Pattern Match` for applications detected by the Firepower System
- `NetFlow` for servers added to the network map from NetFlow records

The system may reconcile data from multiple sources to determine the identity of a server.

Related Topics

Current Identities for Applications and Operating Systems, on page 1821

Viewing Server Details

Procedure

In a host profile, click **View** ( ) next to a server in the **Servers** section.

Editing Server Identities

You can manually update the identity settings for a server on a host and configure any fixes that you have applied to the host to remove the vulnerabilities addressed by the fixes. You can also delete server identities.

Deleting an identity does not delete the server, even if you delete the only identity. Deleting an identity does remove the identity from the Server Detail pop-up window and, if applicable, updates the current identity for the server in the host profile.

You cannot edit or delete server identities added by a Cisco-managed device.

Procedure

Step 1 Navigate to the **Servers** section of a host profile.

Step 2 Click **View** to open the Server Detail pop-up window.

Step 3 To delete a server identity, click **Delete** () next to the server identity you want to remove.

Step 4 To modify a server identity, click **Edit** () next to the server in the servers list.

Step 5 You have two choices:

- Choose the current definition from the **Select Server Type** drop-down list.
- Choose the type of server from the **Select Server Type** drop-down list.

Step 6 Optionally, to only list vendors and products for that server type, choose the **Restrict by Server Type** check box.

Step 7 Optionally, to customize the name and version of the server, choose the **Use Custom Display String**, and enter a **Vendor String** and **Version String**.

Step 8 In the **Product Mappings** section, choose the operating system, product, and versions you want to use.

Example:

For example, if you want the server to map to Red Hat Linux 9, choose **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the version.

Step 9 If you want to indicate that fixes for the server have been applied, click **Configure Fixes**, and add the patches you want to apply for that server to the fixes list.

Step 10 Click **Finish**.

Resolving Server Identity Conflicts

A server identity conflict occurs when an active source, such as an application or scanner, adds identity data for a server to a host, after which the system detects traffic for that port that indicates a conflicting server identity.

Procedure

Step 1 In a host profile, navigate to the **Servers** section.

Step 2 Click resolve next to a server.

Step 3 Choose the type of server from the **Select Server Type** drop-down list.

Step 4 Optionally, to only list vendors and products for that server type, choose the **Restrict by Server Type** check box.

Step 5 Optionally, to customize the name and version of the server, choose **Use Custom Display String**, and enter a **Vendor String** and **Version String**.

Step 6 In the **Product Mappings** section, choose the operating system, product, and versions you want to use.

Example:

For example, if you want the server to map to Red Hat Linux 9, choose **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the version.

Step 7 If you want to indicate that fixes for the server have been applied, click **Configure Fixes**, and add the patches you want to apply for that server to the fixes list.

Step 8 Click **Finish**.

Related Topics

Configuring Network Discovery Identity Conflict Resolution, on page 1931

Web Applications in the Host Profile

The Web Application section of the host profile displays the clients and web applications that the system identifies as running on the hosts on your network. The system can identify client and web application information from both passive and active detection sources, although the information for hosts added from NetFlow records is limited.

Details in this section include the product and version of the detected applications on a host, any available client or web application information, and the time that the application was last detected in use.

The section lists up to 16 clients running on the host. After that limit is reached, new client information from any source, whether active or passive, is discarded until you delete a client application from the host or the system deletes the client from the host profile due to inactivity (the client times out).

Additionally, for each detected web browser, the system displays the first 100 web applications accessed. After that limit is reached, new web applications associated with that browser from any source, whether active or passive, are discarded until either:

- the web browser client application times out, or
- you delete application information associated with a web application from the host profile

If the host is running an application that violates a compliance white list in an activated correlation policy, the Firepower Management Center marks the non-compliant application with the white list **Violation**.

Tip To analyze the connection events associated with a particular application on the host, click **Logging** ( ) next to the application. The first page of your preferred workflow for connection events appears, showing connection events constrained by the type, product, and version of the application, as well as the IP address(es) of the host. If you do not have a preferred workflow for connection events, you must select one.

Descriptions of the application information that appears in a host profile follow.

Application Protocol

Displays the application protocol used by the application (HTTP browser, DNS client, and so on).

Client

Client information derived from payload if identified by the Firepower System, captured by Nmap, or acquired via the host input feature. The field is blank if none of the available sources provides an identification.

Version

Displays the version of the client.

Web Application

For web browsers, the content detected by the system in the http traffic. Web application information indicates the specific type of content (for example, WMV or QuickTime) identified by the Firepower System, captured by Nmap, or acquired via the host input feature. The field is blank if none of the available sources provides an identification.

Deleting Web Applications from the Host Profile

You can delete an application from a host profile to remove applications that you know are not running on the host. Note that deleting an application from a host may bring the host into compliance with a compliance white list.

Note If the system detects the application again, it re-adds it to the network map and the host profile.

Procedure

Step 1 In a host profile, navigate to the **Applications** section.

Step 2 Click **Delete** ( ) next to the application you want to delete.

Host Protocols in the Host Profile

Each host profile contains information about the protocols detected in the network traffic associated with the host. This information includes:

Protocol

The name of a protocol used by the host.

Layer

The network layer where the protocol runs (`Network` or `Transport`).

If a protocol displaying in the host profile violates a compliance white list in an activated correlation policy, the Firepower Management Center marks the non-compliant protocol with the white list **violation**.

If the host profile lists protocols that you know are not running on the host, you can delete those protocols. Deleting a protocol from a host may bring the host into compliance with a compliance white list.

Note If the system detects the protocol again, it re-adds it to the network map and the host profile.

Deleting a Protocol From the Host Profile

Procedure

Step 1 Navigate to the **Protocols** section of a host profile.

Step 2 Click **Delete** ( ) next to the protocol you want to delete.

Indications of Compromise in the Host Profile

The Firepower System correlates various types of data (intrusion events, Security Intelligence, connection events, and file or malware events) to determine whether a host on your monitored network is likely to be compromised by malicious means. Certain combinations and frequencies of event data trigger indications of compromise (IOC) tags on affected hosts.

The Indications of Compromise section of the host profile displays all indication of compromise tags for a host.

To configure the system to tag indications of compromise, see Enabling Indications of Compromise Rules, on page 1933.

For more information about working with indications of compromise, see Indications of Compromise Data, on page 2373 and the subtopics under that topic.

Related Topics

Indications of Compromise, on page 1932

VLAN Tags in the Host Profile

The VLAN Tag section of the host profile appears if the host is a member of a Virtual LAN (VLAN).

Physical network equipment often uses VLANs to create logical network segments from different network blocks. The system detects 802.1q VLAN tags and displays the following information for each:

- **VLAN ID** identifies the VLAN where the host is a member. This can be any integer between zero and 4095 for 802.1q VLANs.
- **Type** identifies the encapsulated packet containing the VLAN tag, which can be either Ethernet or Token Ring.
- **Priority** identifies the priority in the VLAN tag, which can be any integer from zero to 7, where 7 is the highest priority.

If VLAN tags are nested within the packet, the system processes and the Firepower Management Center displays the innermost VLAN tag. The system collects and displays VLAN tag information only for MAC addresses that it identifies through ARP and DHCP traffic.

VLAN tag information can be useful, for example, if you have a VLAN composed entirely of printers and the system detects a Microsoft Windows 2000 operating system in that VLAN. VLAN information also helps the system generate more accurate network maps.

User History in the Host Profile

The user history portion of the host profile provides a graphic representation of the last twenty-four hours of user activity. A typical user logs off in the evening and may share the host resource with another user. Periodic login requests, such as those made to check email, are indicated by short regular bars. A list of user identities is provided with bar graphs to indicate when the user login was detected. Note that for non-authoritative logins, the bar graph is gray.

Note that the system does associate a non-authoritative user login to a host with an IP address of that host, so the user does appear in the host's user history. However, if an authoritative user login is detected for the same host, the user associated with the authoritative user login takes over the association with the host IP address, and new non-authoritative user logins do not disrupt that user association with the host IP address. If you configure capture of failed logins in the network discovery policy, the list includes users that failed to log into the host.

Host Attributes in the Host Profile

You can use *host attributes* to classify hosts in ways that are important to your network environment. Three types of attributes are present in the Firepower System:

- *predefined host attributes*
- *compliance white list host attributes*
- *user-defined host attributes*

After you set a predefined host attribute or create a user-defined host attribute, you must assign a host attribute value.



Note Host attributes can be defined at any domain level. You can assign host attributes created in current and ancestor domains.

Predefined Host Attributes

The Firepower Management Center provides two predefined host attributes:

Host Criticality

Use this attribute to designate the business criticality of a given host and to tailor correlation responses to host criticality. For example, if you consider your organization's mail servers more critical to your business than a typical user workstation, you can assign a value of High to your mail servers and other

business-critical devices and Medium or Low to other hosts. You can then create a correlation policy that launches different alerts based on the criticality of an affected host.

Notes

Use this host-specific attribute to record information about the host that you want other analysts to view. For example, if you have a computer on the network that has an older, unpatched version of an operating system that you use for testing, you can use the Notes feature to indicate that the system is intentionally unpatched.

White List Host Attributes

Each compliance white list that you create automatically creates a host attribute with the same name as the white list. Possible values for white list host attributes are:

- Compliant — Identifies hosts that are compliant with the white list.
- Non-Compliant — Identifies hosts that violate the white list.
- Not Evaluated — Identifies hosts that are not valid targets of the white list or have not been evaluated for any reason.

You cannot edit the value of a white list host attribute or delete a white list host attribute.

User-Defined Host Attributes

If you want to identify hosts using criteria that differs from those used in the predefined host attributes or compliance white list host attributes, you can create user-defined host attributes. For example, you can:

- Assign physical location identifiers to hosts, such as a facility code, city, or room number.
- Assign a Responsible Party Identifier that indicates which system administrator is responsible for a given host. You can then craft correlation rules and policies to send alerts to the correct system administrator when problems related to a host are detected.
- Automatically assign values to hosts from a predefined list based on the hosts' IP addresses. This feature can be useful to assign values to new hosts when they appear on your network for the first time.

User-defined host attributes appear in the host profile page, where you can assign values on a per-host basis. You can also:

- Use the attributes in correlation policies and searches.
- View the attributes on the host attribute table view of events and generate reports based on them.

User-defined host attributes can be one of the following types:

Text

Allows you to manually assign a text string to a host.

Integer

Allows you to specify the first and last number of a range of positive integers, then manually assign one of these numbers to a host.

List

Allows you to create a list of string values, then manually assign one of the values to a host. You can also automatically assign values to hosts based on the host's IP addresses.

If you auto-assign values based on one IP address of a host with multiple IP addresses, those values will apply across all addresses associated with that host. Keep this in mind when you view the Host Attributes table.

When automatically assigning list values, consider using network objects rather than literal IP addresses. This approach can improve maintainability, particularly in a multidomain deployment where using override-enabled objects allows descendant domain administrators to tailor ancestor configurations to their local environments. In a multidomain deployment, be careful when defining auto-assigned lists at ancestor domain levels to avoid matching unintended hosts when the descendant domains use overlapping IP addresses.

URL

Allows you to manually assign a URL value to a host.

Deleting a user-defined host attribute removes it from every host profile where it is used.

Creating Text- or URL-Based Host Attributes

Procedure

Step 1 Choose **Analysis** > **Hosts** > **Host Attributes**.

Step 2 Click **Host Attribute Management**.

Step 3 Click **Create Attribute**.

Step 4 Enter a **Name**.

Step 5 Choose the **Type** of attribute that you want to create as described in User-Defined Host Attributes, on page 2342

Step 6 Click **Save**.

Creating Integer-Based Host Attributes

When you define an integer-based host attribute, you must specify the range of numbers that the attribute accepts.

Procedure

Step 1 Choose **Analysis** > **Hosts** > **Host Attributes**.

Step 2 Click **Host Attribute Management**.

Step 3 Click **Create Attribute**.

Step 4 Enter a **Name**.

Step 5 Choose the **Type** of attribute that you want to create as described in User-Defined Host Attributes, on page 2342.

Step 6 In the **Min** field, enter the minimum integer value that can be assigned to a host.

Step 7 In the **Max** field, enter the maximum integer value that can be assigned to a host.

Step 8 Click **Save**.

Creating List-Based Host Attributes

When you define a list-based host attribute, you must supply each of the values for the list. These values can contain alphanumeric characters, spaces, and symbols.

Procedure

Step 1 Choose **Analysis** > **Hosts** > **Host Attributes**.

Step 2 Click **Host Attribute Management**.

Step 3 Click **Create Attribute**.

Step 4 Enter a **Name**.

Step 5 Choose the **Type** of attribute that you want to create as described in User-Defined Host Attributes, on page 2342.

Step 6 To add a value to the list, click **Add Value**.

Step 7 In the **Name** field, enter the first value you want to add.

Step 8 Optionally, to auto-assign the attribute value you just added to your hosts, click **Add Networks**.

Step 9 Choose the value you added from the **Value** drop-down list.

Step 10 In the **IP Address** and **Netmask** fields, enter the IP address and network mask (IPv4) that represent the IP address block where you want to auto-assign this value.

Step 11 Repeat steps 6 through 10 to add additional values to the list and assign them automatically to new hosts that fall within an IP address block.

Step 12 Click **Save**.

Setting Host Attribute Values

You can set values for predefined and user-defined host attributes. You cannot set values for compliance white list host attributes generated by the system.

Procedure

Step 1 Open the host profile you want to modify.

Step 2 In the **Attributes** section, click **Edit Attributes**.

Step 3 Update attribute as desired.

Step 4 Click **Save**.

White List Violations in the Host Profile

A *compliance white list* (or *white list*) is a set of criteria that allows you to specify the operating systems, application protocols, clients, web applications, and protocols that are allowed to run on a specific subnet.

If you add a white list to an active correlation policy, when the system detects that a host is violating the white list, the Firepower Management Center logs a white list event—which is a special kind of correlation event—to the database. Each of these white list events is associated with a *white list violation*, which indicates how and why a particular host is violating the white list. If a host violates one or more white lists, you can view these violations in its host profile in two ways.

First, the host profile lists all of the individual white list violations associated with the host.

Descriptions of the white list violation information in the host profile follow.

Type

The type of the violation, that is, whether the violation occurred as a result of a non-compliant operating system, application, server, or protocol.

Reason

The specific reason for the violation. For example, if you have a white list that allows only Microsoft Windows hosts, the host profile displays the current operating system running on the host (such as `Linux Linux 2.4, 2.6`)

White List

The name of the white list associated with the violation.

Second, in the sections associated with operating systems, applications, protocols, and servers, the Firepower Management Center marks non-compliant elements with the white list **Violation**. For example, for a white list that allows only Microsoft Windows hosts, the host profile displays the white list violation icon next to the operating system information for that host.



Note You can use a host's profile to create a shared host profile for compliance white lists.

Creating Shared White List Host Profiles

Shared host profiles for compliance white lists specify which operating systems, application protocols, clients, web applications, and protocols are allowed to run on target hosts across multiple white lists. That is, if you create multiple white lists but want to use the same host profile to evaluate hosts running a particular operating system across the white lists, use a shared host profile.

You can use a host profile of any host with a known IP address to create a shared host profile that your compliance white lists can use. However, note that you cannot create a shared host profile based on an individual host's host profile if the system has not yet identified the operating system of the host.

Procedure

Step 1 In a host profile, click **Generate White List Profile**.

Step 2 Modify and save the shared host profile according to your specific needs.

Related Topics

Building White List Host Profiles, on page 1978

Malware Detections in the Host Profile

The Most Recent Malware Detections section lists the most recent malware events where the host sent or received a malware file, up to 100 events. The host profile lists both network-based malware events (those generated by AMP for Networks) and endpoint-based malware events (those generated by AMP for Endpoints).

If the host is involved in a file event where the file is then retrospectively identified as malware, the original events where the file was transmitted appear in the malware detections list after the malware identification occurs. When a file identified as malware is retrospectively determined not to be malware, the malware events related to that file no longer appear in the list. For example, if a file has a disposition of `Malware` and that disposition changes to `Clean`, the event for that file is removed from the malware detections list on the host profile.

When viewing malware detections in the host profile, you can view malware events for that host by clicking the **Malware**.

Description of the columns in the Most Recent Malware Detections sections of the host profile follow.

Time

The date and time the event was generated.

For an event where the file was retrospectively identified as malware, note that this is the time of the original event, not the time when the malware was identified.

Host Role

The host's role in the transmission of detected malware, either sender or receiver. Note that for malware events generated by AMP for Endpoints ("endpoint-based malware events"), the host is always the receiver.

Threat Name

The name of the detected malware.

File Name

The name of the malware file.

File Type

The type of file; for example, `PDF` or `MSEXE`.

Vulnerabilities in the Host Profile

The Vulnerabilities sections of the host profile list the vulnerabilities that affect that host. These vulnerabilities are based on the operating system, servers, and applications that the system detected on the host.

If there is an identity conflict for either the identity of the host's operating system or one of the application protocols on the host, the system lists vulnerabilities for both identities until the conflict is resolved.

Because no operating system information is available for hosts added to the network map from NetFlow data, the system cannot assign Vulnerable (impact level 1: red) impact levels for intrusion events involving those hosts. In such cases, use the host input feature to manually set the operating system identity for the hosts.

Server vendor and version information is often not included in traffic. By default, the system does not map the associated vulnerabilities for the sending and receiving hosts of such traffic. However, you can configure the system to map vulnerabilities for specific application protocols that do not have vendor or version information.

If you use the host input feature to add third-party vulnerability information for the hosts on your network, additional Vulnerabilities sections appear. For example, if you import vulnerabilities from a QualysGuard Scanner, host profiles on your include a QualysGuard Vulnerabilities section. For third-party vulnerabilities, the information in the corresponding Vulnerabilities section in the host profile is limited to the information that you provided when you imported the vulnerability data using the host input feature.

You can associate third-party vulnerabilities with operating systems and application protocols, but not clients. For information on importing third-party vulnerabilities, see the *Firepower System Host Input API Guide*.

Descriptions of the columns in the Vulnerabilities sections of the host profile follow.

Name

The name of the vulnerability.

Remote

Indicates whether the vulnerability can be remotely exploited. If this column is blank, the vulnerability definition does not include this information.

Component

The name of the operating system, application protocol, or client associated with the vulnerability.

Port

A port number, if the vulnerability is associated with an application protocol running on a specific port.

Related Topics

Vulnerability Data Fields, on page 2386
Vulnerability Deactivation, on page 2387

Downloading Patches for Vulnerabilities

You can download patches to mitigate the vulnerabilities discovered on the hosts on your network.

Procedure

Step 1 Access the host profile of a host for which you want to download a patch.
Step 2 Expand the **Vulnerabilities** section.
Step 3 Click the name of the vulnerability you want to patch.
Step 4 Expand the **Fixes** section to display the list of patches for the vulnerability.
Step 5 Click **Download** next to the patch you want to download.
Step 6 Download the patch and apply it to your affected systems.

Deactivating Vulnerabilities for Individual Hosts

You can use the host vulnerability editor to deactivate vulnerabilities on a host-by-host basis. When you deactivate a vulnerability for a host, it is still used for impact correlations for that host, but the impact level is automatically reduced one level.

Procedure

Step 1 Navigate to the **Vulnerabilities** section of a host profile.
Step 2 Click **Edit Vulnerabilities**.
Step 3 Choose the vulnerability from the **Valid Vulnerabilities** list, and click the down arrow to move it to the **Invalid Vulnerabilities** list.

Tip You can click and drag to choose multiple adjacent vulnerabilities; you can also double-click any vulnerability to move it from list to list.

Step 4 Click **Save**.

What to do next

- Optionally, activate the vulnerability for the host by moving it from the **Invalid Vulnerabilities** list to the **Valid Vulnerabilities** list.

Related Topics

Deactivating Individual Vulnerabilities, on page 2348
Deactivating Multiple Vulnerabilities, on page 2389

Deactivating Individual Vulnerabilities

If you deactivate a vulnerability in a host profile, it deactivates it for all hosts in your network map. However, you can reactivate it at any time.

In a multidomain deployment, deactivating a vulnerability in an ancestor domain deactivates it in all descendant domains. Leaf domains can activate or deactivate a vulnerability for their devices if the vulnerability is activated in the ancestor domain.

Procedure

Step 1 Access the vulnerability detail:

- In an affected host profile, expand the **Vulnerabilities** section, and click the name of the vulnerability you want to enable or disable.
- In the predefined workflow, choose **Analysis** > **Vulnerabilities** > **Vulnerabilities**, and click **View** ( ) next to the vulnerability you want to enable or disable.

Step 2 Choose **Disabled** from the **Impact Qualification** drop-down list.

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Confirm that you want to change the **Impact Qualification** value for all hosts on the network map.

Step 4 Click **Done**.

What to do next

- Optionally, activate the vulnerability by choosing **Enabled** from the **Impact Qualification** drop-down list while performing the steps above.

Related Topics

Deactivating Vulnerabilities for Individual Hosts, on page 2348
Deactivating Multiple Vulnerabilities, on page 2389
Operating System Identity Conflicts, on page 2332

Scan Results in the Host Profile

When you scan a host using Nmap, or when you import results from an Nmap scan, those results appear in the host profile for any hosts included in the scan.

The information that Nmap collects about the host operating system and any servers running on open unfiltered ports is added directly into the Operating System and Servers sections of the host profile, respectively. In addition, Nmap adds a list of the scan results for that host in the Scan Results section. Note that the scan must find open ports on the host for Scan Results section to appear in the profile.

Each result indicates the source of the information, the number and type of the scanned port, the name of the server running on the port, and any additional information detected by Nmap, such as the state of the port or the vendor name for the server. If you scan for UDP ports, servers detected on those ports only appear in the Scan Results section.

Note that you can run an Nmap scan from the host profile.

Scanning a Host from the Host Profile

You can perform a Nmap scan against a host from the host profile. After the scan completes, server and operating system information for that host are updated in the host profile. Any additional scan results are added to the Scan Results section of the host profile.

Caution Nmap-supplied server and operating system data remains static until you run another Nmap scan or override it with higher priority host input. If you plan to scan a host using Nmap, regularly schedule scans.

Before you begin

- Add an Nmap scan instance; see Adding an Nmap Scan Instance, on page 1863.

Procedure

Step 1 In the host profile, click **Scan Host**.

Step 2 Click **Scan** next to the scan remediation you want to use to scan the host.

The system scans the host and adds the results to the host profile.

Related Topics

Nmap Scan Automation, on page 184



CHAPTER 116

Working with Discovery Events

The following topics describe how to work with discovery events:

- Requirements and Prerequisites for Discovery Events, on page 2351
- Discovery and Identity Data in Discovery Events, on page 2351
- Viewing Discovery Event Statistics, on page 2352
- Viewing Discovery Performance Graphs, on page 2355
- Using Discovery and Identity Workflows, on page 2356
- History for Working with Discovery Events, on page 2402

Requirements and Prerequisites for Discovery Events

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Security Analyst

Discovery and Identity Data in Discovery Events

The system generates tables of events that represent the changes detected in your monitored network. You can use these tables to review the user activity on your network and determine how to respond. The *network discovery* and *identity* policies specify the kinds of data you want to collect, the network segments you want to monitor, and the specific hardware interfaces you want to use to do it.

You can use discovery and identity event tables to identify threats associated with hosts, applications, and users on your network. The system provides a set of predefined workflows that you can use to analyze the

events that your system generates. You can also create custom workflows that display only the information that matches your specific needs.

To collect and store network discovery and identity data for analysis, you must configure network discovery and identity policies. After you configure an identity policy, you must invoke it in your access control policy and deploy it to the devices you want to use to monitor traffic.

Your network discovery policy provides host, application, and non-authoritative user data. Your identity policy provides authoritative user data.

The following discovery event tables are located under the Analysis > Hosts, Analysis > Users, and Analysis > Vulnerabilities menus.

Discovery Event Table	Populated With Discovery Data?	Populated With Identity Data?
Hosts	Yes	No
Host Indications of Compromise	Yes	No
Applications	Yes	No
Application Details	Yes	No
Servers	Yes	No
Host Attributes	Yes	No
Discovery Events	Yes	Yes
User Indications of Compromise	Yes	Yes
User Activity	Yes	Yes
Users	Yes	Yes
Vulnerabilities	Yes	No
Third-Party Vulnerabilities	Yes	No

Viewing Discovery Event Statistics

The Discovery Statistics page displays a summary of the hosts, events, protocols, application protocols, and operating systems detected by the system.

The page lists statistics for the last hour and the total accumulated statistics. You can choose to view statistics for a particular device, or all devices. You can also view events that match the entries on the page by clicking the event, server, operating system, or operating system vendor listed within the summary.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Overview** > **Summary** > **Discovery Statistics**.

Step 2 From the **Select Device** list, choose the device whose statistics you want to view. Optionally, choose **All** to view statistics for all devices managed by the Firepower Management Center.

Step 3 You have the following options:

- In the Statistics Summary, view general statistics as described in The Statistics Summary Section, on page 2353.
- In the Event Breakdown, click the type of event you want to view. If no events appear, you may need to adjust the time range as described in Changing the Time Window, on page 2168.
- In the Protocol Breakdown, view the protocols currently in use by detected hosts.
- In the Application Protocol Breakdown, click the name of the application protocol you want to view.
- In the OS Breakdown, click the **OS Name** or **OS Vendor**.

Related Topics

The Event Breakdown Section, on page 2354
The Protocol Breakdown Section, on page 2354
The Application Protocol Breakdown Section, on page 2355
The OS Breakdown Section, on page 2355

The Statistics Summary Section

Descriptions of the rows of the Statistics Summary section follow.

Total Events

Total number of discovery events stored on the Firepower Management Center.

Total Events Last Hour

Total number of discovery events generated in the last hour.

Total Events Last Day

Total number of discovery events generated in the last day.

Total Application Protocols

Total number of application protocols from servers running on detected hosts.

Total IP Hosts

Total number of detected hosts identified by unique IP address.

Total MAC Hosts

Total number of detected hosts not identified by IP address.

Note that the Total MAC Hosts statistic remains the same whether you are viewing discovery statistics for all devices or for a specific device. This is so because managed devices discover hosts based on their IP addresses. This statistic gives the total of all hosts that are identified by other means and is independent of a given managed device.

Total Routers

Total number of detected nodes identified as routers.

Total Bridges

Total number of detected nodes identified as bridges.

Host Limit Usage

Total percentage of the host limit currently in use. The host limit is defined by the model of your Firepower Management Center. Note that the host limit usage only appears if you are viewing statistics for all managed devices.



Note If the host limit is reached and a host is deleted, the host will not reappear on the network map you purge discovery data.

Last Event Received

The date and time that the most recent discovery event occurred.

Last Connection Received

The date and time that the most recent connection was completed.

The Event Breakdown Section

The Event Breakdown section lists a count of each type of discovery event and host input event that occurred within the last hour, as well as a count of the total number of each event type stored in the database.

You can also use the Event Breakdown section to view details on discovery and host input events.

Related Topics

Discovery and Host Input Events, on page 2358

The Protocol Breakdown Section

The Protocol Breakdown section lists the protocols currently in use by detected hosts. It displays each detected protocol name, its “layer” in the protocol stack, and the total number of hosts that communicate using the protocol.

The Application Protocol Breakdown Section

The Application Protocol Breakdown section lists the application protocols that are currently in use by detected hosts. It lists the protocol name, the total number of hosts running the application protocol in the past hour, and the total number of hosts that have been detected running the protocol at any point.

You can also use the Application Protocol Breakdown section to view details on servers using the detected protocols.

Related Topics

Server Data, on page 2378

The OS Breakdown Section

The OS Breakdown section lists the operating systems currently running on the monitored network, along with their vendors and the total number of hosts running each operating system.

A value of `unknown` for the operating system name or version means that the operating system or its version does not match any of the system's fingerprints. A value of `pending` means that the system has not yet gathered enough information to identify the operating system or its version.

You can use the OS Breakdown section to view details on the detected operating systems.

Related Topics

Host Data, on page 2365

Viewing Discovery Performance Graphs

You can generate graphs that display performance statistics for managed devices with discovery events.

New data is accumulated for statistics graphs every five minutes. Therefore, if you reload a graph quickly, the data may not change until the next five-minute increment occurs.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

Edit the applicable network discovery policy to include applications, hosts, and users. (This may impact system performance.) See Configuring Network Discovery Rules, on page 1920 and Actions and Discovered Assets, on page 1921.

You must be an Admin or Maintenance user to perform this task.

Procedure

Step 1 Choose **Overview** > **Summary** > **Discovery Performance**.

Step 2 From the **Select Device** list, choose the Firepower Management Center or managed devices you want to include.

Step 3 From the **Select Graph(s)** list, choose the type of graph you want to create as described in Discovery Performance Graph Types, on page 2356.

Step 4 From the **Select Time Range** list, choose the time range you would like to use for the graph.

Step 5 Click **Graph** to graph the selected statistics.

Discovery Performance Graph Types

Descriptions of the available graph types follow.

Processed Events/Sec

Displays a graph that represents the number of events that the Data Correlator processes per second

Processed Connections/Sec

Displays a graph that represents the number of connections that the Data Correlator processes per second

Generated Events/Sec

Displays a graph that represents the number of events that the system generates per second

Mbits/Sec

Displays a graph that represents the number of megabits of traffic that are analyzed by the discovery process per second

Avg Bytes/Packet

Displays a graph that represents the average number of bytes included in each packet analyzed by the discovery process

K Packets/Sec

Displays a graph that represents the number of packets analyzed by the discovery process per second, in thousands

Using Discovery and Identity Workflows

The Firepower Management Center provides a set of event workflows that you can use to analyze the discovery and identity data that is generated for your network. The workflows are, along with the network map, a key source of information about your network assets.

The Firepower Management Center provides predefined workflows for discovery and identity data, detected hosts and their host attributes, servers, applications, application details, vulnerabilities, user activities, and users. You can also create custom workflows.

Procedure

Step 1 To access a predefined workflow:

- Discovery and Host Input Data — See Viewing Discovery and Host Input Events, on page 2364.

- Host Data — See Viewing Host Data, on page 2366.
- Host Attributes Data — See Viewing Host Attributes, on page 2372.
- Host or User Indications of Compromise Data — See View and Work with Indications of Compromise Data, on page 2374.
- Server Data — See Viewing Server Data, on page 2378.
- Application Data — See Viewing Application Data, on page 2381.
- Application Detail Data — See Viewing Application Detail Data, on page 2383.
- User Data — See Viewing User Data, on page 2397.
- User Activity Data — See Viewing User Activity Data, on page 2400.
- Network Map — See Viewing Network Maps, on page 2122.

Step 2 To access a custom workflow, choose **Analysis** > **Custom** > **Custom Workflows**.

Step 3 To access a workflow based on a custom table, choose **Analysis** > **Custom** > **Custom Tables**.

Step 4 Perform any of the following actions, which are common to all of the pages accessed in the network discovery workflows:

- Constrain Columns — To constrain the columns that display, click **Close** ( ) in the column heading that you want to hide. In the pop-up window that appears, click **Apply**.

  Tip To hide or show other columns, check or clear the appropriate check boxes before you click **Apply**. To add a disabled column back to the view, click the expand arrow to expand the search constraints, then click the column name under Disabled Columns.

- Delete — To delete some or all items in the current constrained view, check the check boxes next to items you want to delete and click **Delete**, or click **Delete All**. These items remain deleted until the system's discovery function is restarted, when they may be detected again.

  Caution Before you delete a session on the Analysis > Users > Users page, verify that the session is actually closed. After you delete the active session, an applicable policy will not be able to detect the session on the device, and therefore the session will not be monitored or blocked even if the policy was configured to perform those actions.

  Note You **cannot** delete Cisco (as opposed to third-party) vulnerabilities; you can, however, mark them reviewed.

- Drill Down — To drill down to the next page in the workflow, see Using Drill-Down Pages, on page 2158.
- Navigate Current Page — To navigate within the current workflow page, see Workflow Page Navigation Tools, on page 2155.
- Navigate within a Workflow — To navigate between pages in the current workflow, keeping the current constraints, click the appropriate page link at the top left of the workflow page.
- Navigate to Other Workflows — To navigate to other event views to examine associated events, see Inter-Workflow Navigation, on page 2173.

- Sort Data — To sort data in a workflow, click the column title. Click the column title again to reverse the sort order.
- View Host Profile — To view the host profile for an IP address, click **Host Profile** or, for hosts with active indications of compromise (IOC) tags, the **Compromised Host** that appears next to the IP address.
- View User Profile — To view user identity information, click the user icon that appears next to the **User Identity**, or for users associated with IOCs,**Red User**.

Related Topics

Using Workflows, on page 2150
Purging Data from the FMC Database, on page 198

Discovery and Host Input Events

The system generates discovery events that communicate the details of changes in your monitored network segments. *New* events are generated for newly discovered network features, and change events are generated for any change in previously identified network assets.

During its initial network discovery phase, the system generates new events for each host and any TCP or UDP servers discovered running on each host. Optionally, you can configure the system to use exported NetFlow records to generate these new host and server events.

In addition, the system generates new events for each network, transport, and application protocol running on each discovered host. You can disable detection of application protocols in discovery rules configured to monitor NetFlow exporters, but not in discovery rules configured to monitor Firepower System managed devices. If you enable host or user discovery in non-NetFlow discovery rules, applications are automatically discovered.

After the initial network mapping is complete, the system continuously records network changes by generating change events. Change events are generated whenever the configuration of a previously discovered asset changes.

When a discovery event is generated, it is logged to the database. You can use the Firepower Management Center web interface to view, search, and delete discovery events. You can also use discovery events in correlation rules. Based on the type of discovery event generated as well as other criteria that you specify, you can build correlation rules that, when used in a correlation policy, launch remediations and syslog, SNMP, and email alert responses when network traffic meets your criteria.

You can add data to the network map using the host input feature. You can add, modify, or delete operating system information, which causes the system to stop updating that information for that host. You can also manually add, modify, or delete application protocols, clients, servers, and host attributes or modify vulnerability information. When you do this, the system generates host input events.

Discovery Event Types

You can configure the types of discovery events the system logs in your network discovery policy. When you view the discovery events table, the event type is listed in the **Event** column. Descriptions of the discovery event types follow.

Additional MAC Detected for Host

This event is generated when the system detects a new MAC address for a previously discovered host.

This event is often generated when the system detects hosts passing traffic through a router. While each host has a different IP address, they all appear to have the MAC address associated with the router. When the system detects the actual MAC address associated with the IP address, it displays the MAC address in bold text within the host profile and displays an "ARP/DHCP detected" message within the event description in the event view.

Client Timeout

This event is generated when the system drops a client from the database due to inactivity.

Client Update

This event is generated when the system detects a payload (that is, a specific type of content, such as audio, video, or webmail) in HTTP traffic.

DHCP: IP Address Changed

This event is generated when the system detects that a host IP address has changed due to DHCP address assignment.

DHCP: IP Address Reassigned

This event is generated when a host is reusing an IP address; that is, when a host obtains an IP address formerly used by another physical host due to DHCP IP address assignment.

Hops Change

This event is generated when the system detects a change in the number of network hops between a host and the device that detects the host. This may happen if:

- The device sees host traffic through different routers and is able to make a better determination of the host's location.
- The device detects an ARP transmission from the host, indicating that the host is on a local segment.

Host Deleted: Host Limit Reached

This event is generated when the host limit on the Firepower Management Center is exceeded and a monitored host is deleted from the network map.

Host Dropped: Host Limit Reached

This event is generated when the host limit on the Firepower Management Center is reached and a new host is dropped. Compare this with the previous event where old hosts are deleted from the network map when the host limit is reached.

To drop new hosts when the host limit is reached, go to **Policies** > **Network Discovery** > **Advanced** and set **When Host Limit Reached** to **Drop hosts**.

Host IOC Set

This event is generated when an IOC (Indications of Compromise) is set for a host and generates an alert.

Host Timeout

This event is generated when a host is dropped from the network map because the host has not produced traffic within the interval defined in the network discovery policy. Note that individual host IP addresses and MAC addresses time out individually; a host does not disappear from the network map unless all of its associated addresses have timed out.

If you change the networks you want to monitor in your network discovery policy, you may want to manually delete old hosts from the network map so that they do not count against your host limit.

Host Type Changed to Network Device

This event is generated when the system detects that a detected host is actually a network device.

Identity Conflict

This event is generated when the system detects a new server or operating system identity that conflicts with a current active identity for that server or operating system.

If you want to resolve identity conflicts by rescanning the host to obtain newer active identity data, you can use Identity Conflict events to trigger an Nmap remediation.

Identity Timeout

This event is generated when server or operating system identity data from an active source times out.

If you want to refresh identity data by rescanning the host to obtain newer active identity data, you can use Identity Conflict events to trigger an Nmap remediation.

MAC Information Change

This event is generated when the system detects a change in the information associated with a specific MAC address or TTL value.

This event often occurs when the system detects hosts passing traffic through a router. While each host hnetas a different IP address, they will all appear to have the MAC address associated with the router. When the system detects the actual MAC address associated with the IP address, it displays the MAC address in bold text within the host profile and displays an “ARP/DHCP detected” message within the event description in the event view. The TTL may change because the traffic may pass through different routers or if the system detects the actual MAC address of the host.

NETBIOS Name Change

This event is generated when the system detects a change to a host’s NetBIOS name. This event will only be generated for hosts using the NetBIOS protocol.

New Client

This event is generated when the system detects a new client.



Note To collect and store client data for analysis, make sure that you enable application detection in your discovery rules in the network discovery policy.

New Host

This event is generated when the system detects a new host running on the network.

This event can also be generated when a device processes NetFlow data that involves a new host. To generate an event in this case, configure the network discovery rule that manages NetFlow data to discover hosts.

New Network Protocol

This event is generated when the system detects that a host is communicating with a new network protocol (IP, ARP, and so on).

New OS

This event is generated when the system either detects a new operating system for a host, or a change in a host's operating system.

New TCP Port

This event is generated when the system detects a new TCP server port (for example, a port used by SMTP or web services) active on a host. This event is not used to identify the application protocol or the server associated with it; that information is transmitted in the TCP Server Information Update event.

This event can also be generated when a device processes NetFlow data involving a server on your monitored networks that does not already exist in the network map. To generate an event in this case, configure the network discovery rule that manages NetFlow data to discover applications.

New Transport Protocol

This event is generated when the system detects that a host is communicating with a new transport protocol, such as TCP or UDP.

New UDP Port

This event is generated when the system detects a new UDP server port running on a host.

This event can also be generated when a device processes NetFlow data involving a server on your monitored networks that does not already exist in the network map. To generate an event in this case, configure the network discovery rule that manages NetFlow data to discover applications.

TCP Port Closed

This event is generated when the system detects that a TCP port has closed on a host.

TCP Port Timeout

This event is generated when the system has not detected activity from a TCP port within the interval defined in the system's network discovery policy.

TCP Server Information Update

This event is generated when the system detects a change in a discovered TCP server running on a host.

This event may be generated if a TCP server is upgraded.

UDP Port Closed

This event is generated when the system detects that a UDP port has closed on a host.

UDP Port Timeout

This event is generated when the system has not detected activity from a UDP port within the interval defined in the network discovery policy.

UDP Server Information Update

This event is generated when the system detects a change in a discovered UDP server running on a host.

This event may be generated if a UDP server is upgraded.

VLAN Tag Information Update

This event is generated when the system detects a change in the VLAN tag attributed to a host.

Related Topics

Host Input Event Types, on page 2362
Network Discovery Data Storage Settings, on page 1934
Application and Operating System Identity Conflicts, on page 1822
Network Discovery Identity Conflict Settings, on page 1930

Host Input Event Types

When you view a table of discovery events, the event type is listed in the **Event** column.

Contrast host input events, which are generated when a user takes a specific action (such as manually adding a host), with discovery events, which are generated when the system itself detects a change in your monitored network (such as detecting traffic from a previously undetected host).

You can configure the types of host input events that the system logs by modifying your network discovery policy.

If you understand the information the different types of host input events provide, you can more effectively determine which events you want to log and alert on, and how to use these alerts in correlation policies. In addition, knowing the names of the event types can help you craft more effective event searches. Descriptions of the different types of host input events follow.

Add Client

This event is generated when a user adds a client.

Add Host

This event is generated when a user adds a host.

Add Protocol

This event is generated when a user adds a protocol.

Add Scan Result

This event is generated when the system adds the results of an Nmap scan to a host.

Add Port

This event is generated when a user adds a server port.

Delete Client

This event is generated when a user deletes a client from the system.

Delete Host/Network

This event is generated when a user deletes an IP address or subnet from the system.

Delete Protocol

This event is generated when a user deletes a protocol from the system.

Delete Port

This event is generated when a user deletes a server port or group of server ports from the system.

Host Attribute Add

This event is generated when a user creates a new host attribute.

Host Attribute Delete

This event is generated when a user deletes a user-defined host attribute.

Host Attribute Delete Value

This event is generated when a user deletes a value assigned to a host attribute.

Host Attribute Set Value

This event is generated when a user sets a host attribute value for a host.

Host Attribute Update

This event is generated when a user changes the definition of a user-defined host attribute.

Set Host Criticality

This event is generated when a user sets or modifies the host criticality value for a host.

Set Operating System Definition

This event is generated when a user sets the operating system for a host.

Set Server Definition

This event is generated when a user sets the vendor and version definitions for a server.

Set Vulnerability Impact Qualification

This event is generated when a vulnerability impact qualification is set.

When a vulnerability is disabled at a global level from being used for impact qualifications, or when a vulnerability is enabled at a global level, this event is generated.

Vulnerability Set Invalid

This event is generated when a user invalidates (or reviews) a vulnerability or vulnerabilities.

Vulnerability Set Valid

This event is generated when a user validates a vulnerability that was previously marked as invalid.

Related Topics

Discovery Event Types, on page 2358

Viewing Discovery and Host Input Events

Discovery events workflows allow you to view data from both discovery events and host input events. You can manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access events differs depending on the workflow you use. You can use the predefined workflow, which includes the table view of discovery events and a terminating host view page. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1 Choose **Analysis** > **Hosts** > **Discovery Events**.

Step 2 You have the following options:

- Adjust the time range as described in Changing the Time Window, on page 2168.

  Note Events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; see Using Discovery and Identity Workflows, on page 2356.
- Learn more about the contents of the columns in the table; see Discovery Event Fields, on page 2364.

Related Topics

Using Discovery and Identity Workflows, on page 2356

Discovery Event Fields

Descriptions of the fields that can be viewed and searched in the discovery events table follow.

Time

The time that the system generated the event.

Event

The discovery event type or host input event type.

IP Address

The IP address associated with the host involved in the event.

User

The last user to log into the host involved in the event before the event was generated. If only non-authoritative users log in after an authoritative user, the authoritative user remains the current user for the host unless another authoritative user logs in.

MAC Address

The MAC address of the NIC used by the network traffic that triggered the discovery event. This MAC address can be either the actual MAC address of the host involved in the event, or the MAC address of a network device that the traffic passed through.

MAC Vendor

The MAC hardware vendor of the NIC used by the network traffic that triggered the discovery event.

When searching this field, enter **`virtual_mac_vendor`** to match events that involve virtual hosts.

Port

The port used by the traffic that triggered the event, if applicable.

Description

The text description of the event.

Domain

The domain of the device that discovered the host. This field is only present if you have ever configured the Firepower Management Center for multitenancy.

Device

The name of the managed device that generated the event. For new host and new server events based on NetFlow data, this is the managed device that processed the data.

Related Topics

Event Searches, on page 2177

Host Data

The system generates an event when it detects a host and collects information about it to build the host profile. You can use the Firepower Management Center web interface to view, search, and delete hosts.

While viewing hosts, you can create traffic profiles and compliance white lists based on selected hosts. You can also assign host attributes, including host criticality values (which designate business criticality) to groups

of hosts. You can then use these criticality values, white lists, and traffic profiles within correlation rules and policies.

The system can add hosts to the network map from exported NetFlow records, but the available information for these hosts is limited; see Differences between NetFlow and Managed Device Data, on page 1824.

Related Topics

Differences between NetFlow and Managed Device Data, on page 1824

Viewing Host Data

You can use the Firepower Management Center to view a table of hosts that the system has detected. Then, you can manipulate the view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access hosts differs depending on the workflow you use. Both predefined workflows terminate in a host view, which contains a host profile for every host that meets your constraints. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1 Access the host data:

- If you are using the predefined workflow, choose **Analysis** > **Hosts** > **Hosts**.
- If you are using a custom workflow that does not include the table view of hosts, click **(switch workflow)**, then choose **Hosts**.

Step 2 You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; see Using Discovery and Identity Workflows, on page 2356.
- Learn more about the contents of the columns in the table; see Host Data Fields, on page 2366.
- Right-click an item in the table to see options. (Not every column offers options.)
- Assign a host attribute to specific hosts; see Setting Host Attributes for Selected Hosts, on page 2373.
- Create traffic profiles for specific hosts, see Creating a Traffic Profile for Selected Hosts, on page 2370.
- Create a compliance white list based on specific hosts, see Creating a Compliance White List Based on Selected Hosts, on page 2371.

Host Data Fields

When the system discovers a host, it collects data about that host. That data can include the host's IP addresses, the operating system it is running, and more. You can view some of that information in the table view of hosts.

Descriptions of the fields that can be viewed and searched in the hosts table follow below.

Last Seen

The date and time any of the host's IP addresses was last detected by the system. The Last Seen value is updated at least as often as the update interval you configured in the network discovery policy, as well as when the system generates a new host event for any of the host's IP addresses.

For hosts with operating system data updated using the host input feature, the Last Seen value indicates the date and time when the data was originally added.

IP Address

The IP addresses associated with the host.

MAC Address

The host's detected MAC address of the NIC.

The MAC Address field appears in the Table View of Hosts, which you can find in the Hosts workflow. You can also add the MAC Address field to:

- custom tables that include fields from the Hosts table
- drill-down pages in custom workflows based on the Hosts table

MAC Vendor

The host's detected MAC hardware vendor of the NIC.

The MAC Vendor field appears in the Table View of Hosts, which you can find in the Hosts workflow. You can also add the MAC Vendor field to:

- custom tables that include fields from the Hosts table
- drill-down pages in custom workflows based on the Hosts table

When searching this field, enter `virtual_mac_vendor` to match events that involve virtual hosts.

Current User

The user identity (username) of the currently logged in user on the host.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Host Criticality

The user-specified criticality value assigned to the host.

NetBIOS Name

The NetBIOS name of the host. Only hosts running the NetBIOS protocol will have a NetBIOS name.

VLAN ID

VLAN ID used by the host.

Hops

The number of network hops from the device that detected the host to the host.

Host Type

The type of host. Can be any of the following: host, mobile device, jailbroken mobile device, router, bridge, NAT device, and load balancer.

The methods the system uses to distinguish network devices include:

- the analysis of Cisco Discovery Protocol (CDP) messages, which can identify network devices and their type (Cisco devices only)
- the detection of the Spanning Tree Protocol (STP), which identifies a device as a switch or bridge
- the detection of multiple hosts using the same MAC address, which identifies the MAC address as belonging to a router
- the detection of TTL value changes from the client side, or TTL values that change more frequently than a typical boot time, which identify NAT devices and load balancers

If a device is not identified as a network device, it is categorized as a host.

When searching this field, enter `!host` to search for all network devices.

Hardware

The hardware platform for a mobile device.

OS

One of the following:

- The operating system (name, vendor, and version) either detected on the host or updated using Nmap or the host input feature
- `unknown` if the operating system does not match any known fingerprint
- `pending` if the system has not yet gathered enough information to identify the operating system

If the system detects multiple identities, it displays those identities in a comma-separated list.

This field appears when you invoke the hosts event view from the Custom Analysis widget on the dashboard. It is also a field option in custom tables based on the Hosts table.

When searching this field, enter `n/a` to include hosts where the operating system has not yet been identified.

OS Conflict

This field is search only.

OS Vendor

One of the following:

- The vendor of the operating system detected on the host or updated using Nmap or the host input feature
- `unknown` if the operating system does not match any known fingerprint
- `pending` if the system has not yet gathered enough information to identify the operating system

If the system detects multiple vendors, it displays those vendors in a comma-separated list.

When searching this field, enter `n/a` to include hosts where the operating system has not yet been identified.

OS Name

One of the following:

- The operating system detected on the host or updated using Nmap or the host input feature
- `unknown` if the operating system does not match any known fingerprint
- `pending` if the system has not yet gathered enough information to identify the operating system

If the system detects multiple names, it displays those names in a comma-separated list.

When searching this field, enter `n/a` to include hosts where the operating system has not yet been identified.

OS Version

One of the following:

- The version of the operating system detected on the host or updated using Nmap or the host input feature
- `unknown` if the operating system does not match any known fingerprint
- `pending` if the system has not yet gathered enough information to identify the operating system

If the system detects multiple versions, it displays those versions in a comma-separated list.

When searching this field, enter `n/a` to include hosts where the operating system has not yet been identified.

Source Type

The type of source used to establish the host's operating system identity:

- User: user_name
- Application: app_name
- Scanner: scanner_type (Nmap or scanner added through network discovery configuration)
- `Firepower` for operating systems detected by the system

The system may reconcile data from multiple sources to determine the identity of an operating system.

Confidence

One of the following:

- the percentage of confidence that the system has in the identity of the operating system running on the host, for hosts detected by the system
- 100%, for operating systems identified by an active source, such as the host input feature or Nmap scanner
- `unknown`, for hosts for which the system cannot determine an operating system identity, and for hosts added to the network map based on NetFlow data

When searching this field, enter `n/a` to include hosts added to the network map based on NetFlow data.

Notes

The user-defined content of the Notes host attribute.

Domain

The domain associated with the host. This field is only present if you have ever configured the Firepower Management Center for multitenancy.

Device

Either the managed device that detected the traffic or the device that processed NetFlow or host input data.

If this field is blank, either of the following conditions is true:

- The host was added to the network map by a device that is not explicitly monitoring the network where the host resides, as defined in the network discovery policy.
- The host was added using the host input feature and has not also been detected by the system.

Count

The number of events that match the information that appears in each row. This field appears only after you apply a constraint that creates two or more identical rows.

Related Topics

Event Searches, on page 2177
Operating System Identity Conflicts, on page 2332

Creating a Traffic Profile for Selected Hosts

A traffic profile is a profile of the traffic on your network, based on connection data collected over a timespan that you specify. After you create a traffic profile, you can detect abnormal network traffic by evaluating new traffic against your profile, which presumably represents normal network traffic.

You can use the Hosts page to create a traffic profile for a group of hosts that you specify. The traffic profile will be based on connections detected where one of the hosts you specify is the initiating host. Use the sort and search features to isolate the hosts for which you want to create a profile.

Before you begin

You must be an Admin user to perform this task.

Procedure

Step 1 On a table view in the hosts workflow, check the check boxes next to the hosts for which you want to create a traffic profile.

Step 2 At the bottom of the page, click **Create Traffic Profile**.

Step 3 Modify and save the traffic profile according to your specific needs.

Related Topics

Introduction to Traffic Profiles, on page 2021

Creating a Compliance White List Based on Selected Hosts

Compliance white lists allow you to specify which operating systems, clients, and network, transport, or application protocols are allowed on your network.

You can use the Hosts page to create a compliance white list based on the host profiles of a group of hosts that you specify. Use the sort and search features to isolate the hosts that you want to use to create a white list.

Before you begin

You must be an Admin user to perform this task.

Procedure

Step 1 On a table view in the hosts workflow, check the check boxes next to the hosts for which you want to create a white list.

Step 2 At the bottom of the page, click **CreateWhite List**.

Step 3 Modify and save the white list according to your specific needs.

Related Topics

Introduction to Compliance White Lists, on page 1971

Host Attribute Data

The Firepower System collects information about the hosts it detects and uses that information to build host profiles. However, there may be additional information about the hosts on your network that you want to provide to your analysts. You can add notes to a host profile, set the business criticality of a host, or provide any other information that you choose. Each piece of information is called a *host attribute*.

You can use host attributes in host profile qualifications, which constrain the data you collect while building a traffic profile, and also can limit the conditions under which you want to trigger a correlation rule. You can also set attribute values in response to a correlation rule.

Related Topics

Viewing Host Attributes, on page 2372
Configuring Set Attribute Remediations, on page 2042

Viewing Host Attributes

You can use the Firepower Management Center to view a table of hosts detected by the system, along with their host attributes. Then, you can manipulate the view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access host attributes differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of host attributes that lists all detected hosts and their attributes, and terminates in a host view page, which contains a host profile for every host that meets your constraints.

You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1 Access the host attributes data:

- If you are using the predefined workflow, choose **Analysis** > **Hosts** > **Host Attributes**.
- If you are using a custom workflow that does not include the table view of host attributes, click **(switch workflow)**, then choose **Attributes**.

Step 2 You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; see Using Discovery and Identity Workflows, on page 2356.
- Learn more about the contents of the columns in the table; see Host Attribute Data Fields, on page 2372.
- Assign a host attribute to specific hosts; see Setting Host Attributes for Selected Hosts, on page 2373.

Host Attribute Data Fields

Note that the host attributes table does not display hosts identified only by MAC addresses.

Descriptions of the fields that can be viewed and searched in the host attributes table follow.

IP Address

The IP addresses associated with a host.

Current User

The user identity (username) of the currently logged in user on the host.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Host Criticality

The user-assigned importance of a host to your enterprise. You can use the host criticality in correlation rules and policies to tailor policy violations and their responses to the importance of a host involved in an event. You can assign a host criticality of low, medium, high, or none.

Notes

Information about the host that you want other analysts to view.

Any user-defined host attribute, including those for compliance white lists

The value of the user-defined host attribute. The host attributes table contains a field for each user-defined host attribute.

Domain

The domain associated with the host. This field is only present if you have ever configured the Firepower Management Center for multitenancy.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Related Topics

Event Searches, on page 2177

Setting Host Attributes for Selected Hosts

You can configure predefined and user-defined host attributes from a host workflow.

Procedure

Step 1 In a host workflow, check the check boxes next to the hosts to which you want to add a host attribute.

Tip Use the sort and search features to isolate the hosts to which you want to assign particular attributes.

Step 2 At the bottom of the page, click **Set Attributes**.

Step 3 Optionally, set the host criticality for the hosts you selected. You can choose **None**, **Low**, **Medium**, or **High**.

Step 4 Optionally, add notes to the host profiles of the hosts you selected in the text box.

Step 5 Optionally, set any user-defined host attributes you have configured.

Step 6 Click **Save**.

Indications of Compromise Data

The Firepower System correlates various types of data (intrusion events, Security Intelligence, connection events, and file or malware events) to determine whether a host on your monitored network is likely to be compromised by malicious means. Certain combinations and frequencies of event data trigger indications of

compromise (IOC) tags on affected hosts. The IP addresses of these hosts appear in event views with a **Red Compromised Host icon** .

When a host is identified as potentially compromised, the user associated with that compromise is also tagged. These users appear in event views with a **Red User icon**.

If a file containing malware is seen again within 300 seconds of being tagged as an IOC, another IOC is not generated. If the same file is seen more than 300 seconds later, a new IOC will be generated.

To configure the system to tag events as indications of compromise, see Enabling Indications of Compromise Rules, on page 1933.

Related Topics

Enabling Indications of Compromise Rules, on page 1933

View and Work with Indications of Compromise Data

You can use the Firepower Management Center to view tables showing Indications of Compromise (IOC). Manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see depends on the workflow you use. The predefined IOC workflows terminate in a profile view, which contains a host or user profile for every host or user that meets your constraints. You can also create a custom workflow that displays only the information that matches your specific needs.

Before you begin

- For your system to detect and tag indications of compromise (IOC), you must activate the IOC feature in the network discovery policy and enable at least one IOC rule. See Enabling Indications of Compromise Rules, on page 1933.
- Users must be identified in an active Identity policy.

Procedure

Step 1 Determine which location in the web interface presents information that meets your needs.

You can use the following locations to view or work with Indication of Compromise data:

- Event Viewer (under the Analysis menu) — Connection, Security Intelligence, intrusion, malware, and IOC discovery event views indicate whether an event triggered an IOC. Note that malware events generated by AMP for Endpoints that trigger IOC rules have the event type `AMP IOC` and appear with an event subtype that specifies the compromise.
- Dashboard — In the dashboard, Threats of the Summary Dashboard displays, by default, IOC tags by host and by user. The Custom Analysis widget offers presets based on IOC data.
- Context Explorer — The Indications of Compromise section of the Context Explorer displays graphs of hosts by IOC category and IOC categories by host.
- Network Map page — The Indications of Compromise under Analysis > Hosts > Network Map groups potentially compromised hosts on your network by type of compromise and IP address.

- Network File Trajectory details page — The details pages for files listed under Analysis > Files > Network File Trajectory let you track indications of compromise on your network.
- Host Indications of Compromise page — The Host Indications of Compromise page under the Analysis > Hosts menu lists monitored hosts, grouped by IOC tag. Use the workflows on this page to drill down into your data.
- User Indications of Compromise page — The User Indications of Compromise page under the Analysis > Users menu lists users associated with potential IOC events, grouped by IOC tag. Use the workflows on this page to drill down into your data.
- Host Profile page — The host profile for a potentially compromised host displays all IOC tags associated with that host, and lets you resolve IOC tags and configure IOC rule states.
- User Profile page — The user profile for a user associated with a potential IOC event displays all IOC tags associated with that user, and lets you resolve IOC tags and configure IOC rule states. (The user profile is labeled "User Identity" in the Firepower Management Center web interface.)

Step 2 If applicable, do one of the following and use the rest of the steps in this procedure:

Option	Description
To research IOCs on hosts:	• If you are using the predefined workflow, choose **Analysis** > **Hosts** > **Indications of Compromise**. • If you are using a custom workflow that does not include the Host IOC table view, click **(switch workflow)**, then choose **Host Indications of Compromise** .
To research IOCs associated with users:	• If you are using the predefined workflow, choose **Analysis > Users > Indications of Compromise**. • If you are using a custom workflow that does not include the User IOC table view, click **(switch workflow)**, then choose **User Indications of Compromise**.

Step 3 You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; see Using Discovery and Identity Workflows, on page 2356.
- Learn more about the contents of the columns in the table; see Indications of Compromise Data Fields, on page 2376.
- On a Host Indications of Compromise page: View the host profile for a compromised host by clicking **Compromised Host** in the **IP Address** column.
- On a User Indications of Compromise page: View the user profile associated with a compromise by clicking **Red User** in the **User** column.
- Mark IOC events resolved so they no longer appear in the list. To do so, check the check boxes next to the IOC events you want to modify, then click **Mark Resolved**.
- View details of events that triggered the IOC by clicking **View** ( ) in the **First Seen** or **Last Seen** columns.

- See more options: Right-click a value in the table.

Indications of Compromise Data Fields

The following are the fields in Host or User IOC (indication of compromise) tables. Not every IOC-related table includes all fields.

IP Address (When viewing Host IOC data)

The IP address associated with the host that triggered the IOC.

User (When viewing User IOC data)

The username, realm, and authentication source of the user associated with the event that triggered the IOC.

Category

Brief description of the type of compromise indicated, such as `Malware Executed` or `Impact 1 Attack`.

Event Type

Identifier associated with a specific IOC, referring to the event that triggered it.

Description

Description of the impact on the potentially compromised host, such as `This host may be under remote control` or `Malware has been executed on this host`.

First Seen/Last Seen

The first/most recent date and time that events triggering the IOC occurred.

Domain

The domain of the host that triggered the IOC. This field is only present if you have ever configured the Firepower Management Center for multitenancy.

Related Topics

Event Searches, on page 2177

Editing Indication of Compromise Rule States for a Single Host or User

When enabled in a network discovery policy, indication of compromise rules apply to all hosts in the monitored network and to authoritative users that are associated with IOC events on that network. You can disable a rule for an individual host or user to avoid unhelpful IOC tags (for example, you may not want to see IOC tags for a DNS server.) If a rule is disabled in the applicable network discovery policy, it cannot be enabled for a specific host or user. Disabling a rule for a particular host does not affect tagging for the user involved in the same event, and vice-versa.

Procedure

Step 1 Navigate to the **Indications of Compromise** section of a host or user profile.

Step 2 Click **Edit Rule States**.

Step 3 In the **Enabled** column for a rule, click the slider to enable or disable it.

Step 4 Click **Save**.

Viewing Source Events for Indication of Compromise Tags

You can use the Indications of Compromise section of the host profile and the user profile to navigate quickly to the events that triggered the IOC tags. Analyzing these events can give you the information you need to determine what, and whether, action is required to address threats of compromise.

Clicking **View** ( ) next to the timestamp of an IOC tag navigates to the table view of events for the relevant event type, constrained to show only the event that triggered the IOC tag.

Only the first instance of a User IOC is displayed in the Firepower Management Center. Subsequent instances are caught by the DNS Server."

Procedure

Step 1 In a host or user profile, navigate to the **Indications of Compromise** section.

Step 2 Click **View** ( ) in the **First Seen** or **Last Seen** column for the IOC tag you want to investigate.

Resolving Indication of Compromise Tags

After you have analyzed and addressed the threats indicated by an indication of compromise (IOC) tag, or if you determine that an IOC tag represents a false positive, you can mark an event resolved. Marking an event resolved removes it from the host profile and the user profile; when all active IOC tags on a profile are resolved, the **Compromised Host** or a user is associated with an indication of compromise **Red User icon** no longer appears. You can still view the IOC-triggering events for the resolved IOC.

If the events that triggered the IOC tag recur, the tag is set again unless you have disabled the IOC rule for the host or user.

Procedure

Step 1 In a host or user profile, navigate to the **Indications of Compromise** section.

Step 2 You have two choices:

- To mark an individual IOC tag resolved, click **Delete** ( ) to the right of the tag you want to resolve.
- To mark all IOC tags on the profile resolved, click **Mark All Resolved**.

Server Data

The Firepower System collects information about all servers running on hosts on monitored network segments. This information includes:

- the name of the server
- the application and network protocols used by the server
- the vendor and version of the server
- the IP address associated with the host running a server
- the port on which the server communicates

When the system detects a server, it generates a discovery event unless the associated host has already reached its maximum number of servers. You can use the Firepower Management Center web interface to view, search, and delete server events.

You can also base correlation rules on server events. For example, you could trigger a correlation rule when the system detects a chat server, such as ircd, running on one of your hosts.

The system can add hosts to the network map from exported NetFlow records, but the available information for these hosts is limited; see Differences between NetFlow and Managed Device Data, on page 1824.

Related Topics

Host Limits and Discovery Event Logging, on page 1876
Differences between NetFlow and Managed Device Data, on page 1824

Viewing Server Data

You can use the Firepower Management Center to view a table of detected servers. Then, you can manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access servers differs depending on the workflow you use. All the predefined workflows terminate in a host view, which contains a host profile for every host that meets your constraints. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1 Access the server data:

- If you are using the predefined workflow, choose **Analysis** > **Hosts** > **Servers**.
- If you are using a custom workflow that does not include the table view of servers, click **(switch workflow)**, then choose **Servers**.

Step 2 You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; see Using Discovery and Identity Workflows, on page 2356.
- Learn more about the contents of the columns in the table; see Server Data Fields, on page 2379.

- Edit server identities by checking the check boxes next to the events for servers you want to edit, then clicking **Set Server Identity**.
- Right-click an item in the table to see options. (Not every column offers options.)

Related Topics

Editing Server Identities, on page 2336

Server Data Fields

Descriptions of the fields that can be viewed and searched in the servers table follow below.

Last Used

The date and time the server was last used on the network or the date and time that the server was originally updated using the host input feature. The Last Used value is updated at least as often as the update interval you configured in the network discovery policy, as well as when the system detects a server information update.

IP Address

The IP address associated with the host running the server.

Port

The port where the server is running.

Protocol

The network or transport protocol used by the server.

Application Protocol

One of the following:

- the name of the application protocol for the server
- `pending` if the system cannot positively or negatively identify the server for one of several reasons
- `unknown` if the system cannot identify the server based on known server fingerprints or if the server was added through host input and did not include the application protocol

Category, Tags, Risk, or Business Relevance for Application Protocols

The categories, tags, risk level, and business relevance assigned to the application protocol. These filters can be used to focus on a specific set of data.

Vendor

One of the following:

- the server vendor as identified by the system, Nmap or another active source, or that you specified using the host input feature

- blank, if the system cannot identify its vendor based on known server fingerprints, or if the server was added to the network map using NetFlow data

Version

One of the following:

- the server version as identified by the system, Nmap or another active source, or that you specified using the host input feature
- blank, if the system cannot identify its version based on known server fingerprints, or if the server was added to the network map using NetFlow data

Web Application

The web application based on the payload content detected by the system in the HTTP traffic. Note that if the system detects an application protocol of `HTTP` but cannot detect a specific web application, the system supplies a generic web browsing designation.

Category, Tags, Risk, or Business Relevance for Web Applications

The categories, tags, risk level, and business relevance assigned to the web application. These filters can be used to focus on a specific set of data.

Hits

The number of times the server was accessed. For servers added using the host input feature, this value is always 0.

Source Type

One of the following values:

- User: user_name
- Application: app_name
- Scanner: scanner_type (Nmap or scanner added through network discovery configuration)
- `Firepower`, `Firepower Port Match`, or `Firepower Pattern Match` for servers detected by the Firepower System
- `NetFlow` for servers added using NetFlow data

Domain

The domain of the host running the server. This field is only present if you have ever configured the Firepower Management Center for multitenancy.

Device

Either the managed device that detected the traffic or the device that processed NetFlow or host input data.

Current User

The user identity (username) of the currently logged in user on the host.

When a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Count

The number of events that match the information that appears in each row. This field appears only after you apply a constraint that creates two or more identical rows.

Related Topics

Event Searches, on page 2177
Network Discovery Data Storage Settings, on page 1934

Application and Application Details Data

When a monitored host connects to another host, the system can, in many cases, determine what application was used. The Firepower System detects the use of many email, instant messaging, peer-to-peer, web applications, as well as other types of applications.

For each detected application, the system logs the IP address that used the application, the product, the version, and the number of times its use was detected. You can use the web interface to view, search, and delete application events. You can also update application data on a host or hosts using the host input feature.

If you know which applications are running on which hosts, you can use that knowledge to create host profile qualifications, which constrain the data you collect while building a traffic profile, and also can limit the conditions under which you want to trigger a correlation rule. You can also base correlation rules on the detection of application. For example, if you want your employees to use a specific mail client, you could trigger a correlation rule when the system detects a different mail client running on one of your hosts.

You can obtain the latest information about Firepower's application detectors by carefully reading both the release notes for each Firepower System update and advisories for each VDB update.

To collect and store application data for analysis, make sure that you enable application detection in your network discovery policy.

Viewing Application Data

You can use the Firepower Management Center to view a table of detected applications. Then, you can manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access applications differs depending on the workflow you use. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1 Access the application data:

- If you are using the predefined workflow, choose **Analysis** > **Hosts** > **Application Details**.
- If you are using a custom workflow that does not include the table view of application details, click **(switch workflow)**, then choose **Clients**.

Step 2 You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; see Using Discovery and Identity Workflows, on page 2356.
- Learn more about the contents of the columns in the table; see Application Data Fields, on page 2382.
- Open the Application Detail View for a specific application by clicking **Application Detail View** next to a client, application protocol, or web application.

Application Data Fields

When the system detects traffic for a known client, application protocol, or web application, it logs information about the application and the host running it.

Descriptions of the fields that can be viewed and searched in the applications table follow.

Application

The name of the detected application.

IP Address

The IP address associated with the host using the application.

Type

The type of application:

Application Protocols

Represents communications between hosts.

Client Applications

Represents software running on a host.

Web Applications

Represents the content or requested URL for HTTP traffic.

Category

A general classification for the application that describes its most essential function. Each application belongs to at least one category.

Tag

Additional information about the application. Applications can have any number of tags, including none.

Risk

How likely the application is to be used for purposes that might be against your organization's security policy. An application's risk can range from Very Low to Very High.

Of Application Protocol Risk, Client Risk, and Web Application Risk, the highest of the three detected, when available, in the traffic that triggered the intrusion event.

Business Relevance

The likelihood that the application is used within the context of your organization's business operations, as opposed to recreationally. An application's business relevance can range from Very Low to Very High.

Of Application Protocol Business Relevance, Client Business Relevance, and Web Application Business Relevance, the lowest of the three detected, when available, in the traffic that triggered the intrusion event.

Current User

The user identity (username) of the currently logged in user on the host.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Domain

The domain of the host using the application. This field is only present if you have ever configured the Firepower Management Center for multitenancy.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Related Topics

Event Searches, on page 2177

Viewing Application Detail Data

You can use the Firepower Management Center to view a table of detected application details. Then, you can manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access application details differs depending on the workflow you use. There are two predefined workflows. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1 Access the application details data:

- If you are using the predefined workflow, choose **Analysis** > **Hosts** > **Application Details**.
- If you are using a custom workflow that does not include the table view of application details, click **(switch workflow)**, then select **Clients**.

Step 2 You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; see Using Discovery and Identity Workflows, on page 2356.
- Learn more about the contents of the columns in the table; see Application Detail Data Fields, on page 2384.
- Open the Application Detail View for a specific application by clicking **Application Detail View** next to a client.

Application Detail Data Fields

When the system detects traffic for a known client, application protocol, or web application, it logs information about the application and the host running it.

Descriptions of the fields that can be viewed and searched in the application details table follow.

Last Used

The time that the application was last used or the time that the application data was updated using the host input feature. The Last Used value is updated at least as often as the update interval you configured in the network discovery policy, as well as when the system detects an application information update.

IP Address

The IP address associated with the host using the application.

Client

The name of the application. Note that if the system detected an application protocol but could not detect a specific client, `client` is appended to the application protocol name to provide a generic name.

Version

The version of the application.

Category, Tags, Risk, or Business Relevance for Clients, Application Protocols, and Web Applications

The categories, tags, risk level, and business relevance assigned to the application. These filters can be used to focus on a specific set of data.

Application Protocol

The application protocol used by the application. Note that if the system detected an application protocol but could not detect a specific client, `client` is appended to the application protocol name to provide a generic name.

Web Application

The web application based on the payload content or URL detected by the system in the HTTP traffic. Note that if the system detects an application protocol of HTTP but cannot detect a specific web application, the system supplies a generic web browsing designation here.

Hits

The number of times the system detected the application in use. For applications added using the host input feature, this value is always `0`.

Domain

The domain of the host using the application. This field is only present if you have ever configured the Firepower Management Center for multitenancy.

Device

The device that generated the discovery event containing the application detail.

Current User

The user identity (username) of the currently logged in user on the host.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Related Topics

Event Searches, on page 2177
Network Discovery Data Storage Settings, on page 1934

Vulnerability Data

The Firepower System includes its own vulnerability tracking database which is used, in conjunction with the system's fingerprinting capability, to identify the vulnerabilities associated with the hosts on your network. The operating systems, servers, and clients running on your hosts have different sets of associated vulnerabilities.

You can use the Firepower Management Center to:

- Track and review the vulnerabilities for each host.

- Deactivate vulnerabilities for a host after you patch the host or otherwise judge it immune to a vulnerability.

Vulnerabilities for vendorless and versionless servers are not mapped unless the applications protocols used by the servers are mapped in the Firepower Management Center configuration. Vulnerabilities for vendorless and versionless clients cannot be mapped.

Related Topics

Mapping Vulnerabilities for Servers, on page 918

Vulnerability Data Fields

Except as noted, these fields appear on all pages under **Analysis > Hosts > Vulnerabilities**.

Additional Information

This section is on the Vulnerability Details page. It is no longer in use.

Available Exploits

This information is no longer available and the field is blank.

Bugtraq ID

The third-party Bugtraq database is no longer available, so this field is blank.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

CVE ID

The identification number associated with the vulnerability in MITRE's Common Vulnerabilities and Exposures (CVE) database (https://cve.mitre.org/).

This field is available on the Vulnerability Details page. The CVE ID also appears at the beginning of the Title column in vulnerabilities tables.

Date Published

The date the vulnerability was published.

Description

A brief description of the vulnerability, from the National Vulnerability Database (NVD).

For the complete description, look up the CVE ID in the NVD.

Impact Qualification

This field is available only on the Vulnerability Details page.

Use the drop-down list to enable or disable a vulnerability. The Firepower Management Center ignores disabled vulnerabilities in its impact correlations.

The setting you specify here determines how the vulnerability is treated on a system-wide basis and is not limited to the host profile where you select the value.

Remote

Indicates whether the vulnerability is remotely exploitable (`TRUE`/`FALSE`).

Snort ID

The identification number associated with the vulnerability in the Snort ID (SID) database. That is, if an intrusion rule can detect network traffic that exploits a particular vulnerability, that vulnerability is associated with the intrusion rule's SID.

Note that a vulnerability can be associated with more than one SID (or no SIDs at all). If a vulnerability is associated with more than one SID, the vulnerabilities table includes a row for each SID.

Solution

This information is no longer available and the field is blank.

SVID

The vulnerability identification number that the Firepower system uses to track vulnerabilities.

To view details for this vulnerability, click **View** ().

Technical Description

The base score and Common Vulnerability Scoring System score (CVSS) from the National Vulnerability Database (NVD).

Title

The CVE ID of the vulnerability followed by its description.

Vulnerability Impact

The severity of the vulnerability on a scale of 0 to 10, with 10 being the most severe.

Related Topics

Event Searches, on page 2177

Vulnerability Deactivation

Deactivating a vulnerability prevents the system from using that vulnerability to evaluate intrusion impact correlations. You can deactivate a vulnerability after you patch the hosts on your network or otherwise judge them immune. Note that if the system discovers a new host that is affected by that vulnerability, the vulnerability is considered valid (and is not automatically deactivated) for that host.

Deactivating a vulnerability within a vulnerabilities workflow that is **not** constrained by IP addresses deactivates the vulnerability for *all* detected hosts on your network. You can deactivate vulnerabilities within the vulnerabilities workflow only on:

- the second page of the default vulnerabilities workflow, **Vulnerabilities on the Network**, which shows only the vulnerabilities that apply to the hosts on your network

- a page in a vulnerabilities workflow, custom or predefined, that you constrained based on IP address using a search.

You can deactivate a vulnerability for a single host using the network map, using the host's host profile, or by constraining the vulnerabilities workflow based on the IP addresses of the host or hosts for which you want to deactivate vulnerabilities. For hosts with multiple associated IP addresses, this function applies only to the single, selected IP address of that host.

In a multidomain deployment, deactivating a vulnerability in an ancestor domain deactivates it in all descendant domains. Leaf domains can activate or deactivate a vulnerability for their devices if the vulnerability is activated in the ancestor domain.

Related Topics

Deactivating Vulnerabilities for Individual Hosts, on page 2348
Deactivating Individual Vulnerabilities, on page 2348
Deactivating Multiple Vulnerabilities, on page 2389

Viewing Vulnerability Data

You can use the Firepower Management Center to view a table of vulnerabilities. Then, you can manipulate the event view depending on the information you are looking for.

The page you see when you access vulnerabilities differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of vulnerabilities. The table view contains a row for each vulnerability in the database, regardless of whether any of your detected hosts exhibit the vulnerabilities. The second page of the predefined workflow contains a row for each vulnerability (that you have not deactivated) that applies to detected hosts on your network. The predefined workflow terminates in a vulnerability detail view, which contains a detailed description for every vulnerability that meets your constraints.

Tip If you want to see the vulnerabilities that apply to a single host or set of hosts, perform a search for vulnerabilities, specifying an IP address or range of IP addresses for the hosts.

You can also create a custom workflow that displays only the information that matches your specific needs.

The table of vulnerabilities is not restricted by domain in a multidomain deployment.

Procedure

Step 1 Access the table of vulnerabilities:

- If you are using the predefined vulnerabilities workflow, choose **Analysis** > **Vulnerabilities** > **Vulnerabilities**.
- If you are using a custom workflow that does not include the table view of vulnerabilities, click **(switch workflow)**, then choose **Vulnerabilities**.

Step 2 You have the following options:

- Perform basic workflow actions; see Using Discovery and Identity Workflows, on page 2356.
- Deactivate vulnerabilities so they are no longer used for intrusion impact correlation for currently vulnerable hosts; see Deactivating Multiple Vulnerabilities, on page 2389.

- View the details for a vulnerability by clicking **View** ( ) in the SVID column. Alternatively, constrain on the vulnerability ID and drill down to the vulnerability details page. See options for viewing additional details at Viewing Vulnerability Details, on page 2389.
- View the full text of a vulnerability title by right-clicking the title and choosing **Show Full Text**.

Viewing Vulnerability Details

Procedure

You can view vulnerability details in any of the following ways:

- Choose **Analysis** > **Vulnerabilities** > **Vulnerabilities**, and click **View** ( ) next to the SVID.
- Choose **Analysis** > **Vulnerabilities** > **Third-Party Vulnerabilities** and click **View** ( ) next to the SVID.
- Choose **Analysis** > **Hosts** > **Network Map**, and click **Vulnerabilities**.
- View the profile of a host affected by the vulnerability (**Analysis** > **Hosts** > **Network Map**, click **Hosts**, then drill down and click the host you are investigating), and expand the **Vulnerabilities** section of the profile.

Deactivating Multiple Vulnerabilities

Deactivating a vulnerability within a vulnerabilities workflow that is **not** constrained by IP addresses deactivates the vulnerability for *all* detected hosts on your network.

In a multidomain deployment, deactivating a vulnerability in an ancestor domain deactivates it in all descendant domains. Leaf domains can activate or deactivate a vulnerability for their devices so long as the vulnerability is activated in the ancestor domain.

Procedure

Step 1 Access the table of vulnerabilities:

- If you are using the predefined vulnerabilities workflow, choose **Analysis** > **Vulnerabilities** > **Vulnerabilities**.
- If you are using a custom workflow that does not include the table view of vulnerabilities, click **(switch workflow)**, then choose **Vulnerabilities**.

Step 2 Click **Vulnerabilities on the Network**.

Step 3 Check the check boxes next to vulnerabilities you want to deactivate.

Step 4 Click **Review** at the bottom of the page.

Related Topics

Deactivating Vulnerabilities for Individual Hosts, on page 2348
Deactivating Individual Vulnerabilities, on page 2348

Third-Party Vulnerability Data

The Firepower System includes its own vulnerability tracking database which is used, in conjunction with the system's fingerprinting capability, to identify the vulnerabilities associated with the hosts on your network.

You can augment the system's vulnerability data with imported network map data from third-party applications. To do so, your organization must be able to write scripts or create command line import files to import the data. For more information, see the *Firepower System Host Input API Guide*.

To include imported data in impact correlations, you must map third-party vulnerability information to the operating system and application definitions in the database. You cannot map third-party vulnerability information to client definitions.

Viewing Third-Party Vulnerability Data

After you use the host input feature to import third-party vulnerability data, you can use the Firepower Management Center to view a table of third-party vulnerabilities. Then, you can manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access third-party vulnerabilities differs depending on the workflow you use. There are two predefined workflows. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1 Access the third-party vulnerabilities data:

- If you are using the predefined workflow, choose **Analysis** > **Vulnerabilities** > **Third-Party Vulnerabilities**.
- If you are using a custom workflow that does not include the table view of third-party vulnerabilities, click **(switch workflow)**, then choose **Vulnerabilities by Source** or **Vulnerabilities by IP Address**.

Step 2 You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; see Using Discovery and Identity Workflows, on page 2356.
- Learn more about the contents of the columns in the table; see Third-Party Vulnerability Data Fields, on page 2390.
- View the vulnerability details for a third-party vulnerability by clicking **View** () in the SVID column. Alternatively, constrain on the vulnerability ID and drill down to the vulnerability details page.

Third-Party Vulnerability Data Fields

Descriptions of the fields that can be viewed and searched in the third-party vulnerabilities table follow.

Vulnerability Source

The source of the third-party vulnerabilities, for example, QualysGuard or NeXpose.

Vulnerability ID

The ID number associated with the vulnerability for its source.

IP Address

The IP address associated with the host affected by the vulnerability.

Port

A port number, if the vulnerability is associated with a server running on a specific port.

Bugtraq ID

The identification number associated with the vulnerability in the Bugtraq database. (http://www.securityfocus.com/bid/)

CVE ID

The identification number associated with the vulnerability in MITRE's Common Vulnerabilities and Exposures (CVE) database (https://cve.mitre.org/).

SVID

The legacy vulnerability identification number that the system uses to track vulnerabilities

Click **View** () to access the vulnerability details for the SVID.

Snort ID

The identification number associated with the vulnerability in the Snort ID (SID) database. That is, if an intrusion rule can detect network traffic that exploits a particular vulnerability, that vulnerability is associated with the intrusion rule's SID.

Note that a vulnerability can be associated with more than one SID (or no SIDs at all). If a vulnerability is associated with more than one SID, the vulnerabilities table includes a row for each SID.

Title

The title of the vulnerability.

Description

A brief description of the vulnerability.

Domain

The domain of the host with the vulnerability. This field is only present if you have ever configured the Firepower Management Center for multitenancy.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Related Topics

Event Searches, on page 2177

Users and User Activity Data

User and user activity data are displayed in individual user-related workflows:

- Users — this workflow displays all users seen on your network. A single user occupies a single row in this table. For more information, see User Data, on page 2395.
- User Activity — this workflow displays all user activity seen on your network. A single user with more than one instance of user activity would occupy several rows in this table. For more information, see User Activity Data, on page 2398.

For more information about the identity sources that populate these workflows, see About User Identity Sources, on page 1891.

User-Related Fields

User-related data is displayed in the users and user activity tables.

Table 308: Users and User Activity Field Descriptions

Field	Description	Users Table	User Activity Table
Authentication Type	The type of authentication: `No Authentication`, `Passive Authentication`, `Active Authentication`, `Guest Authentication`, or `Failed Authentication`.	No	Yes
Count	**Note** The **Count** field appears only after you apply a constraint that creates two or more identical rows. The number of users or events that match the information that appears in a particular row.	Yes	Yes
Current IP	The IP address associated with the host that the user is logged into. This field is blank if another authoritative user logs into the host with the same IP address after the user's login, unless the user is an authoritative user and the new user is a non-authoritative user. (The system associates the IP address with the last authoritative user that logged in with the host.)	Yes	No
Department	The user's department, as obtained by a realm. If there is no department explicitly associated with the user on your servers, the department is listed as whatever default group the server assigns. For example, on Active Directory, this is `Users (ad)`. This field is blank if: • You have not configured a realm. • The Firepower Management Center cannot correlate the user in the FMC database with an LDAP record (for example, for users added to the database via an AIM, Oracle, or SIP login).	Yes	No

Field	Description	Users Table	User Activity Table
Description	More information, if available, about the user or user activity.	No	Yes
Device	For user activity detected by traffic-based detection, the name of the device that detected the user. For other types of user activity, the managing Firepower Management Center.	No	Yes
Domain	In the Users table, the domain associated with the user's realm. In the User Activity table, the domain where the user activity was detected. This field is only present if you have ever configured the Firepower Management Center for multitenancy.	Yes	Yes
E-Mail	The user's email address. This field is blank if: • The user was added to the database via an AIM login. • The user was added to the database via an LDAP login and there is no email address associated with the user on your LDAP servers.	Yes	No
End Port	If the user was reported by the TS Agent and their session is currently active, this field identifies the end value for the port range assigned to the user. This field is blank if the user's TS Agent session is inactive or if the user was reported by another identity source.	Yes	Yes
Endpoint Location	The IP address of the network device that used ISE to authenticate the user, as identified by ISE. If you do not configure ISE, this field is blank.	No	Yes
Endpoint Profile	The user's endpoint device type, as identified by Cisco ISE. If you do not configure ISE, this field is blank.	No	Yes
Event	The user activity event type.	No	Yes
First Name	The user's first name, as obtained by a realm. This field is blank if: • You have not configured a realm. • The Firepower Management Center cannot correlate the user in the FMC database with an LDAP record (for example, for users added to the database via an AIM, Oracle, or SIP login). • There is no first name associated with the user on your servers.	Yes	No

Field	Description	Users Table	User Activity Table
IP Address	For User Login activity, the IP address involved in the login, which can be an IP address of the user's host (for LDAP, POP3, IMAP, FTP, HTTP, MDNS, and AIM logins), the server (for SMTP and Oracle logins), or the session originator (for SIP logins). Note that an associated IP address does not mean the user is the current user for that IP address; when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. For other types of user activity, this field is blank.	No	Yes
Last Name	The user's last name, as obtained by a realm. This field is blank if: • You have not configured a realm. • The Firepower Management Center cannot correlate the user in the FMC database with an LDAP record (for example, for users added to the database via an AIM, Oracle, or SIP login). • There is no last name associated with the user on your servers.	Yes	No
Phone	The user's telephone number, as obtained by a realm. This field is blank if: • You have not configured a realm. • The Firepower Management Center cannot correlate the user in the FMC database with an LDAP record (for example, for users added to the database via an AIM, Oracle, or SIP login). • There is no telephone number associated with the user on your servers.	Yes	No
Realm	The identity realm associated with the user.	Yes	Yes
Security Group Tag	The Security Group Tag (SGT) attribute applied by Cisco TrustSec as the packet entered a trusted TrustSec network. If you do not configure ISE, this field is blank.	No	Yes
Start Port	If the user was reported by the TS Agent and their session is currently active, this field identifies the start value for the port range assigned to the user. This field is blank if the user's TS Agent session is inactive or if the user was reported by another identity source.	No	Yes
Time	The time that the system detected the user activity.	No	Yes

Field	Description	Users Table	User Activity Table
Type	The protocol used to detect the user. One of the following: ldap, pop3, imap, oracle, sip, http, ftp, mdns, and aim. Users are not added to the database based on SMTP logins, so smtp does not appear in this field.	Yes	Yes
User	At minimum, this field displays the user's realm and username. For example, `Lobby\jsmith`, where Lobby is the realm and jsmith is the username. If a realm downloads additional user data from an LDAP server and the system associates it with a user, this field also displays the user's first name, last name, and type. For example, `John Smith (Lobby\jsmith, LDAP)`, where John Smith is the user's name and LDAP is the type. **Note** Because traffic-based detection can record unsuccessful AIM logins, the Firepower Management Center may store invalid AIM users (for example, if a user misspelled his or her username).	Yes	No
Username	The username associated with the user.	Yes	Yes

User Data

When an identity source reports a user login for a user who is not already in the database, the user is added to the database, unless you have specifically restricted that login type.

The system updates the users database when one of the following occurs:

- A user on the Firepower Management Center manually deletes a non-authoritative user from the Users table.
- An identity source reports a logoff by that user.
- A realm ends the user session as specified by the realm's **User Session Timeout: Authenticated Users**, **User Session Timeout: Failed Authentication Users**, or **User Session Timeout: Guest Users** setting.

Note If you have ISE configured, you may see host data in the users table. Because host detection by ISE is not fully supported, you cannot perform user control using ISE-reported host data.

The type of user login that the system detected determines what information is stored about the new user.

Identity Source	Login Type	User Data Stored
ISE	Active Directory LDAP RADIUS RSA	• username • current IP address • Security Group Tag (SGT) • endpoint profile/device type • endpoint location/location IP • type (LDAP)
User Agent	Active Directory	• username • current IP address • type (LDAP)
TS Agent	Active Directory	• username • current IP address • start port • end port • type (LDAP)
captive portal	Active Directory LDAP	• username • current IP address • type (LDAP)
traffic-based detection	LDAP AIM Oracle SIP HTTP FTP MDNS	• username • current IP address • type (AD)
	POP3 IMAP	• username • current IP address • email address • type (`pop3` or `imap`)

If you configure a realm to automatically download users, the Firepower Management Center queries the servers based on the interval you specified. It may take five to ten minutes for the Firepower Management Center database to update with user metadata after the system detects a new user login. The Firepower Management Center obtains the following information and metadata about each user:

- username
- first and last names
- email address
- department
- telephone number
- current IP address
- Security Group Tag (SGT), if available
- endpoint profile, if available
- endpoint location, if available
- start port, if available
- end port, if available

The number of users the Firepower Management Center can store in its database depends on your Firepower Management Center model. When a non-authoritative user login to a host is detected, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user login is detected for that host, only another authoritative user login changes the current user.

Note that traffic-based detection of AIM, Oracle, and SIP logins create duplicate user records because they are not associated with any of the user metadata that the system obtains from LDAP servers. To prevent overuse of user count because of duplicate user records from these protocols, configure traffic-based detection to ignore those protocols.

You can search, view, and delete users from the database; you can also purge all users from the database.

For information about general user-related event troubleshooting, see Troubleshoot Realms and User Downloads, on page 1945.

Viewing User Data

You can view a table of users, and then manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access users differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of users that lists all detected users, and terminates in a user details page. The user details page provides information on every user that meets your constraints.

Procedure

Step 1 Access the users data:

- If you are using the predefined workflow, choose **Analysis** > **Users** > **Users**.
- If you are using a custom workflow that does not include the table view of users, click **(switch workflow)**, then choose **Users**.

Step 2 You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; see Using Discovery and Identity Workflows, on page 2356.
- Learn more about the contents of the columns in the table; see User-Related Fields, on page 2392.

User Activity Data

The Firepower System generates events that communicate the details of user activity on your network. When the system detects user activity, the user activity data is logged to the database. You can view, search, and delete user activity; you can also purge all user activity from the database.

The system logs a user activity event when a user is seen on your network for the first time. Subsequent appearances by that user do not log new user activity events. However, if the user's IP address changes, the system logs a new user activity event.

The Firepower System also correlates user activity with other types of events. For example, intrusion events can tell you the users who were logged into the source and destination hosts at the time of the event. This correlation can tell you who was logged into the host that was targeted by an attack, or who initiated an internal attack or portscan.

You can also use user activity in correlation rules. Based on the type of user activity as well as other criteria that you specify, you can build correlation rules that, when used in a correlation policy, launch remediations and alert responses when network traffic meets your criteria.



Note If you have ISE configured, you may see host data in the users table. Because host detection by ISE is not fully supported, you cannot perform user control using ISE-reported host data.

Descriptions of the four types of user activity data follow.

New User Identity

This type of event is generated when the system detects a login by an unknown user that is not in the database.

The system logs a user activity event when a user is seen on your network for the first time. Subsequent appearances by that user do not log new user activity events. However, if the user's IP address changes, the system logs a new user activity event.

User Login

This type of event is generated when any of the following occur:

- ISE, user agent, or TS Agent reports a successful user login.
- Captive portal performs a successful or failed user authentication.
- Traffic-based detection detects a successful or failed user login.



Note SMTP logins detected by traffic-based detection are not recorded unless there is already a user with a matching email address in the database.

When a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user.

If you are using captive portal or traffic-based detection, note the following about failed user login and failed user authentication data:

- Failed logins reported by traffic-based detection (LDAP, IMAP, FTP, and POP3 traffic) are displayed in the table view of user activity, but not in the table view of users. If a known user failed to log in, the system identifies them by their username. If an unknown user failed to log in, the system uses **Failed Authentication** as their username.
- Failed authentications reported by captive portal are displayed in both the table view of user activity and the table view of users. If a known user failed to authenticate, the system identifies them by their username. If an unknown user failed to authenticate, the system identifies them by the username they entered.

Delete User Identity

This type of event is generated when you manually delete a user from the database.

User Identity Dropped: User Limit Reached

This type of event is generated when the system detects a user that is not in the database, but cannot add the user because you have reached the maximum number of users in the database as determined by your Firepower Management Center model.

After you reach the user limit, in most cases the system stops adding new users to the database. To add new users, you must either manually delete old or inactive users from the database, or purge all users from the database.

However, the system favors authoritative users. If you have reached the limit and the system detects a login for a previously undetected authoritative user, the system deletes the non-authoritative user who has remained inactive for the longest time, and replaces it with the new authoritative user.

User Indications of Compromise Events

The following user IOC changes are logged in the user activity database:

- When indications of compromise are resolved.
- When indication of compromise rules are enabled or disabled for users.

For information about general user-related event troubleshooting, see Troubleshoot Realms and User Downloads, on page 1945.

Related Topics

The User Activity Database, on page 1830

Viewing User Activity Data

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

You can view a table of user activity, and then manipulate the event view depending on the information you are looking for. The page you see when you access user activity differs depending on the workflow you use. You can use the predefined workflow, which includes the table view of user activity and terminates in a user details page, which contains user details for every user that meets your constraints. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1 Access the user activity data:

- If you are using the predefined workflow, choose **Analysis** > **Users** > **User Activity**.
- If you are using a custom workflow that does not include the table view of user activity, click **(switch workflow)**, then choose **User Activity**.

Tip If no events appear, you may need to adjust the time range; see Changing the Time Window, on page 2168.

Step 2 You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; see Using Discovery and Identity Workflows, on page 2356.
- Learn more about the contents of the columns in the table; see User-Related Fields, on page 2392.

User Profile and Host History

You can learn more about a specific user by viewing the User pop-up window. The page that appears, called the "User Profile" in this document, is titled "User Identity" in the web interface.

You can display the window from:

- any event view that associates user data with other kinds of events
- the table view of users

User information also appears in the terminating page for users workflows.

The user data you see is the same as you would see in the table view of users.

Indications of Compromise Section

For information about this section, see:

- Indications of Compromise, on page 1932

- Indications of Compromise Data Fields, on page 2376
- Editing Indication of Compromise Rule States for a Single Host or User, on page 2376
- Resolving Indication of Compromise Tags, on page 2377
- Viewing Source Events for Indication of Compromise Tags, on page 2377

Host History Section

The host history provides a graphic representation of the last twenty-four hours of the user's activity. A list of IP addresses of the hosts that the user logged into and logged off of approximates login and logout times with bar graphs. A typical user might log on to and off of multiple hosts in the course of a day. For example, periodic automated logins to a mail server would display as multiple short sessions, while longer logins (such as during working hours) display longer sessions.

If you use traffic-based detection or captive portal to capture failed logins, the host history also includes hosts where the user failed to log in.

The data used to generate the host history is stored in the user history database, which by default stores 10 million user login events. If you do not see any data in the host history for a particular user, either that user is inactive, or you may need to increase the database limit.

Related Topics

User Data Fields

Viewing User Details and Host History

Procedure

You have two options:

- In any event view that lists users, click user that appears next to a user identity **User icon**, or, for users associated with an indication of compromise, **Red User icon**.
- In any users workflow, click the Users terminating page.

History for Working with Discovery Events

Table 309:

Feature	Version	Details
Vulnerability data changes	All	The Bugtraq resource for vulnerability data is no longer available. Where possible, vulnerability information is now updated from the National Vulnerability Database (NVD). However, Bugtraq ID, Solution, Available Exploits, and Additional Information will remain blank. Modified screens: • All pages under Analysis > Hosts > Vulnerabilities • Hosts and Vulnerabilities tabs on Analysis > Hosts > Network Map pages Supported Platforms: FMC



CHAPTER 117

Correlation and Compliance Events

The following topics describe how to view correlation and compliance events.

- Viewing Correlation Events, on page 2403
- Using Compliance White List Workflows, on page 2406
- Remediation Status Events, on page 2411

Viewing Correlation Events

When a correlation rule within an active correlation policy triggers, the system generates a correlation event and logs it to the database.

Note When a compliance white list within an active correlation policy triggers, the system generates a white list event.

You can view a table of correlation events, then manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access correlation events differs depending on the workflow you use. You can use the predefined workflow, which includes the table view of correlation events. You can also create a custom workflow that displays only the information that matches your specific needs.

Before you begin

You must be an Admin or Security Analyst user to perform this task.

Procedure

Step 1 Choose **Analysis** > **Correlation** > **Correlation Events** .

Optionally, to use a different workflow, including a custom workflow, click **(switch workflow)** by the workflow title.

Tip If you are using a custom workflow that does not include the table view of correlation events, click **(switch workflow)**, then choose **Correlation Events**.

Step 2 Optionally, adjust the time range as described in Changing the Time Window, on page 2168.

Step 3 Perform any of the following actions:

- To learn more about the columns that appear, see Correlation Event Fields, on page 2404.
- To view the host profile for an IP address, click host profile that appears next to the IP address.
- To view user identity information, click the user icon that appears next to the **User Identity**, or for users associated with IOCs,**Red User**.
- To sort and constrain events or to navigate within the current workflow page, see Using Workflows, on page 2150.
- To navigate between pages in the current workflow, keeping the current constraints, click the appropriate page link at the top left of the workflow page.
- To drill down to the next page in the Workflows, constraining on a specific value, see Using Drill-Down Pages, on page 2158.
- To delete some or all correlation events, check the check boxes next to the events you want to delete and click **Delete**, or click **Delete All** and confirm you want to delete all the events in the current constrained view.
- To navigate to other event views to view associated events, see Inter-Workflow Navigation, on page 2173.

Related Topics

Database Event Limits, on page 879
Workflow Pages, on page 2153

Correlation Event Fields

When a correlation rule triggers, the system generates a correlation event. The fields in the correlation events table that can be viewed and searched are described in the following table.

Table 310: Correlation Event Fields

Field	Description
Description	The description of the correlation event. The information in the description depends on how the rule was triggered. For example, if the rule was triggered by an operating system information update event, the new operating system name and confidence level appears.
Device	The name of the device that generated the event that triggered the policy violation.
Domain	The domain of the device whose monitored traffic triggered the policy violation. This field is only present if you have ever configured the Firepower Management Center for multitenancy.

Field	Description
Impact	The impact level assigned to the correlation event based on the correlation between intrusion data, discovery data, and vulnerability information. When searching this field, valid case-insensitive values are `Impact 0`, `Impact Level 0`, `Impact 1`, `Impact Level 1`, `Impact 2`, `Impact Level 2`, `Impact 3`, `Impact Level 3`, `Impact 4`, and `Impact Level 4`. Do not use impact icon colors or partial strings (for example, do not use `blue`, `level 1`, or `0`).
Ingress Interface or Egress Interface	The ingress or egress interface in the intrusion or connection event that triggered the policy violation.
Ingress Security Zone or Egress Security Zone	The ingress or egress security zone in the intrusion or connection event that triggered the policy violation.
Inline Result	One of: • a black down arrow, indicating that the system dropped the packet that triggered the intrusion rule • a gray down arrow, indicating that the system would have dropped the packet in an inline, switched, or routed deployment if you enabled the **Drop when Inline** intrusion policy option • blank, indicating that the triggered intrusion rule was not set to Drop and Generate Events When using this field to search for policy violations triggered by intrusion events, type either: • `dropped`, to specify whether the packet was dropped in an inline, switched, or routed deployment • `would have dropped`, to specify whether the packet would have dropped if the intrusion policy had been set to drop packets in an inline, switched, or routed deployment Note that the system does not drop packets in a passive deployment, including when an inline set is in tap mode, regardless of the rule state or the drop behavior of the intrusion policy.
Policy	The name of the policy that was violated.
Priority	The priority of the correlation event, which is determined by the priority of either the triggered rule or the violated correlation policy. When searching this field, enter `none` for no priority.
Rule	The name of the rule that triggered the policy violation.
Security Intelligence Category	The name of the object that represents or contains the blocked IP address in the event that triggered the policy violation. When searching this field, specify the Security Intelligence category associated with the correlation event that triggered the policy violation. The Security Intelligence category can be the name of a Security Intelligence object, the global Block list, a custom Security Intelligence list or feed, or one of the categories in the Intelligence Feed.
Source Continent or Destination Continent	The continent associated with the source or destination host IP addresses in the event that triggered the policy violation.

Field	Description
Source Country or Destination Country	The country associated with the source or destination IP address in the event that triggered the policy violation.
Source Host Criticality or Destination Host Criticality	The user-assigned host criticality of the source or destination host involved in the correlation event: `None`, `Low`, `Medium`, or `High`. Note that only correlation events generated by rules based on discovery events, host input events, or connection events contain a source host criticality.
Source IP or Destination IP	The IP address of the source or destination host in the event that triggered the policy violation.
Source Port/ICMP Type or Destination Port/ICMP Code	The source port or ICMP type for the source traffic or the destination port or ICMP code for destination traffic associated with the event that triggered the policy violation.
Source User or Destination User	The name of the user logged in to the source or destination host in the event that triggered the policy violation.
Time	The date and time that the correlation event was generated. This field is not searchable.
Count	The number of events that match the information that appears in each row. Note that the **Count** field appears only after you apply a constraint that creates two or more identical rows. This field is not searchable

Related Topics

Event Searches, on page 2177

Using Compliance White List Workflows

The Firepower Management Center provides a set of workflows that you can use to analyze the white list events and violations that are generated for your network. The workflows are, along with the network map and dashboard, a key source of information about the compliance of your network assets.

The system provides predefined workflows for white list events and violations. You can also create custom workflows. When you are using a compliance white list workflow, you can perform many common actions.

Before you begin

You must be an Admin, Security Analyst, or Discovery Admin user to perform this task.

Procedure

Step 1 Access a white list workflow using the **Analysis** > **Correlation** menu.

Step 2 You have the following options:

- Switch Workflow — To use a different workflow, including a custom workflow, click **(switch workflow)**.
- Time Range — To adjust the time range, which is useful if no events appear, see Changing the Time Window, on page 2168.

- Host Profile — To view the host profile for an IP address, click **Host Profile**() or, for hosts with active indications of compromise (IOC) tags, the **Compromised Host** that appears next to the IP address.
- User Profile (events only) — To view user identity information, click the user icon that appears next to the **User Identity**, or for users associated with IOCs,**Red User**.
- Constrain — To constrain the columns that appear, click **Close** ( ) in the column heading that you want to hide. In the pop-up window that appears, click **Apply**.

  Tip To hide or show other columns, select or clear the appropriate check boxes before you click **Apply**. To add a disabled column back to the view, expand the search constraints, then click the column name under Disabled Columns.

- Drill Down — See Using Drill-Down Pages, on page 2158.
- Sort — To sort data in a workflow, click the column title. Click the column title again to reverse the sort order.
- Navigate This Page — See Workflow Page Traversal Tools, on page 2155.
- Navigate Between Pages — To navigate between pages in the current workflow, keeping the current constraints, click the appropriate page link at the top left of the workflow page.
- Navigate Between Event Views — To navigate to other event views to view associated events, click **Jump to** and select the event view from the drop-down list.
- Delete Events (events only) — To delete some or all items in the current constrained view, select the check boxes next to items you want to delete and click **Delete** or click **Delete All**.

Related Topics

Workflow Pages, on page 2153
Configuring Event View Settings, on page 33

Viewing White List Events

After its initial evaluation, the system generates a *white list event* whenever a monitored host goes out of compliance with an active white list. White list events are a special kind of correlation event, and are logged to the FMC correlation event database.

You can use the Firepower Management Center to view a table of compliance white list events. Then, you can manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access white list events differs depending on the workflow you use. You can use a predefined workflow, which terminates in a table view of events. You can also create a custom workflow that displays only the information that matches your specific needs.

Before you begin

You must be an Admin, Security Analyst, or Discovery Admin user to perform this task.

Procedure

Step 1 Choose **Analysis** > **Correlation**> **White List Events**.

Step 2 You have the following options:

- To perform basic workflow actions, see Using Compliance White List Workflows, on page 2406.
- To learn more about the contents of the columns in the table, see White List Event Fields, on page 2408.
- To see more options, right-click values in the table.

White List Event Fields

White list events, which you can view and search using workflows, contain the following fields.

Device

The name of the managed device that detected the white list violation.

Description

A description of how the white list was violated. For example:

```
Client "AOL Instant Messenger" is not allowed.
```

Violations that involve an application protocol indicate the application protocol name and version, as well as the port and protocol (TCP or UDP) it is using. If you restrict prohibitions to a particular operating system, the description includes the operating system name. For example:

```
Server "ssh / 22 TCP (OpenSSH 3.6.1p2)" is not allowed on Operating System
"Linux Linux 2.4 or 2.6".
```

Domain

The domain of the host that has become non-compliant with the white list. This field is only present if you have ever configured the Firepower Management Center for multitenancy.

Host Criticality

The user-assigned host criticality of the source host that is out of compliance with the white list: None, Low, Medium, or High.

IP Address

The IP address of the host that has become non-compliant with the white list.

Policy

The name of the correlation policy that was violated, that is, the correlation policy that includes the white list.

Port

The port, if any, associated with the discovery event that triggered an application protocol white list violation (a violation that occurred as a result of a non-compliant application protocol). For other types of white list violations, this field is blank.

Priority

The priority specified by the policy or white list that triggered the policy violation. This is determined either by the priority of the white list in a correlation policy or by the priority of the correlation policy itself. Note that the white list priority overrides the priority of its policy. When searching this field, enter `none` for no priority.

Time

The date and time that the white list event was generated. This field is not searchable.

User

The identity of any known user logged in to the host that has become non-compliant with the white list.

White List

The name of the white list.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows. This field is not searchable.

Viewing White List Violations

The system keeps a record of the current *white list violations* on your network. Each violation represents something disallowed running on one of your hosts. If a host becomes compliant, the system removes the now-corrected violation from the database.

You can use the Firepower Management Center to view a table of white list violations for all active white lists. Then, you can manipulate the event view depending on the information you are looking for.

The page you see when you access white list violations differs depending on the workflow you use. The predefined workflows terminate in a host view, which contains a host profile for every host that meets your constraints. You can also create a custom workflow that displays only the information that matches your specific needs.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Correlation** > **White List Violations**.

Step 2 You have the following options:

- To perform basic workflow actions, see Using Compliance White List Workflows, on page 2406.

- To learn more about the contents of the columns in the table, see White List Violation Fields, on page 2410.
- To see more options, right-click values in the table.

White List Violation Fields

White list violations, which you can view and search using workflows, contain the following fields.

Domain

The domain where the non-compliant host resides. This field is only present if you have ever configured the Firepower Management Center for multitenancy.

Information

Any available vendor, product, or version information associated with the white list violation. For protocols that violate a white list, this field also indicates whether the violation is due to a network or transport protocol.

IP Address

The IP address of the non-compliant host.

Port

The port, if any, associated with the event that triggered an application protocol white list violation (a violation that occurred as a result of a non-compliant application protocol). For other types of white list violations, this field is blank.

Protocol

The protocol, if any, associated with the event that triggered an application protocol white list violation (a violation that occurred as a result of a non-compliant application protocol). For other types of white list violations, this field is blank.

Time

The date and time that the white list violation was detected.

Type

The type of white list violation, that is, whether the violation occurred as a result of a non-compliant:

- operating system (os) (When searching this field, enter **`os`** or **`operating system`**.)
- application protocol (server)
- client
- protocol
- web application (web) (When searching this field, enter **`web application`**.)

White List

The name of the white list that was violated.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows. This field is not searchable.

Remediation Status Events

When a remediation triggers, the system logs a remediation status event to the database. These events can be viewed on the Remediation Status page. You can search, view, and delete remediation status events.

Related Topics

Remediation Status Table Fields, on page 2412

Viewing Remediation Status Events

The page you see when you access remediation status events differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of remediations. The table view contains a row for each remediation status event. You can also create a custom workflow that displays only the information that matches your specific needs.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

You must be an Admin user to perform this task.

Procedure

Step 1 Choose **Analysis** > **Correlation** > **Status**.

Step 2 Optionally, adjust the time range as described in Changing the Time Window, on page 2168.

Step 3 Optionally, to use a different workflow, including a custom workflow, click **(switch workflow)** by the workflow title.

Tip If you are using a custom workflow that does not include the table view of remediations, click **(switch workflow)** menu by the workflow title, then choose **Remediation Status**.

Step 4 You have the following options:

- To learn more about the columns that appear, see Remediation Status Table Fields, on page 2412.
- To sort and constrain the events, see Using Workflows, on page 2150.
- To navigate to the correlation events view to see associated events, click **Correlation Events**.
- To bookmark the current page so that you can quickly return to it, click **Bookmark This Page**. To navigate to the bookmark management page, click **View Bookmarks**.

- To generate a report based on the data in the table view, click **Report Designer** as described in Creating a Report Template from an Event View, on page 2055.
- To drill down to the next page in the workflow, see Using Drill-Down Pages, on page 2158.
- To delete remediation status events from the system, check the check boxes next to events you want to delete and click **Delete** or click **Delete All** and confirm you want to delete all the events in the current constrained view.
- To search for remediation status events, click **Search**.

Related Topics

Using Workflows, on page 2150

Remediation Status Table Fields

The following table describes the fields in the remediation status table that can be viewed and searched.

Table 311: Remediation Status Fields

Field	Description
Domain	The domain of the device whose monitored traffic triggered the policy violation, that in turn triggered the remediation. This field is only present if you have ever configured the Firepower Management Center for multitenancy.
Policy	The name of the correlation policy that was violated and triggered the remediation.
Remediation Name	The name of the remediation that was launched.

Field	Description
Result Message	A message that describes what happened when the remediation was launched. Status messages include: • Successful completion of remediation • Error in the input provided to the remediation module • Error in the remediation module configuration • Error logging into the remote device or server • Unable to gain required privileges on remote device or server • Timeout logging into remote device or server • Timeout executing remote commands or servers • The remote device or server was unreachable • The remediation was attempted but failed • Failed to execute remediation program • Unknown/unexpected error If custom remediation modules are installed, you may see additional status messages that are implemented by the custom module.
Rule	The name of the correlation rule that triggered the remediation.
Time	The date and time that the Firepower Management Center launched the remediation
Count	The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows. This field is not searchable.

Related Topics

Event Searches, on page 2177

Using the Remediation Status Events Table

You can change the layout of the event view or constrain the events in the view by a field value.

When you disable a column, it is disabled for the duration of your session unless you add it back later. If you disable the first column, the Count column is added.

Clicking a value within a row in a table view constrains the table view and does not drill down to the next page.



Tip Table views always include “Table View” in the page name.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

You must be an Admin user to perform this task.

Procedure

Step 1 Choose **Analysis** > **Correlation** > **Status**.

Tip If you are using a custom workflow that does not include the table view of remediations, click **(switch workflow)** menu by the workflow title, then choose **Remediation Status**.

Step 2 You have the following options:

- To learn more about the columns that appear, see Remediation Status Table Fields, on page 2412.
- To sort and constrain the events, see Using Workflows, on page 2150.



CHAPTER 118

Auditing the System

The following topics describe how to audit activity on your system:

- The System Log, on page 2415
- About System Auditing, on page 2417

The System Log

The System Log (syslog) page provides you with system log information for the appliance.

You can audit activity on your system in two ways. The appliances that are part of the Firepower System generate an audit record for each user interaction with the web interface, and also record system status messages in the system log.

The system log displays each message generated by the system. The following items are listed in order:

- the date that the message was generated
- the time that the message was generated
- the host that generated the message
- the message itself

Viewing the System Log

System log information is local. For example, you **cannot** use the Firepower Management Center to view system status messages in the system logs on your managed devices.

You can filter messages using most syntax accepted by the UNIX file search utility Grep. This includes using Grep-compatible regular expressions for pattern matching.

Before you begin

You must be an Admin or Maintenace user and be in the Global domain to view system statistics.

Procedure

Step 1 Choose **System** > **Monitoring** > **Syslog**.

Step 2 To search for specific message content in the system log:

a) Enter a word or query in the filter field as described in Syntax for System Log Filters, on page 2416.

Only Grep-compatible search syntax is supported.

Examples:

To search for all log entries that contain the user name "Admin," use `Admin`.

To search for all log entries that are generated on November 27, use `Nov[[:space:]]*27` or `Nov.*27` (but not `Nov 27` or `Nov*27`).

To search for all log entries that contain authorization debugging information on November 5, use `Nov[[:space:]]*5.*AUTH.*DEBUG`.

b) To make your search case-sensitive, select **Case-sensitive**. (By default, filters are not case-sensitive.)

c) To search for all system log messages that do **not** meet the criteria you entered, select **Exclusion**.

d) Click **Go**.

Syntax for System Log Filters

The following table shows the regular expression syntax you can use in System Log filters:

Table 312: System Log Filter Syntax

Syntax Component	Description	Example
`.`	Matches any character or white space	`Admi.` matches `Admin`, `AdmiN`, `Admi1`, and `Admi`
`[[:alpha:]]`	Matches any alphabetic character	`[[:alpha:]]dmin` matches `Admin`, `bdmin`, and
`[[:upper:]]`	Matches any uppercase alphabetic character	`[[:upper:]]dmin` matches `Admin`, `Bdmin`, and
`[[:lower:]]`	Matches any lowercase alphabetic character	`[[:lower:]]dmin` matches `admin`, `bdmin`, and
`[[:digit:]]`	Matches any numeric character	`[[:digit:]]dmin` matches `0dmin`, `1dmin`, and
`[[:alnum:]]`	Matches any alphanumeric character	`[[:alnum:]]dmin` matches `1dmin`, `admin`, `2dmin`,
`[[:space:]]`	Matches any white space, including tabs	`Feb[[:space:]]29` matches logs from Februar
`*`	Matches zero or more instances of the character or expression it follows	`ab*` matches `a`, `ab`, `abb`, `ca`, `cab`, and `cabb` `[ab]*` matches anything
`?`	Matches zero or one instances	`ab?` matches `a` or `ab`
`\`	Allows you to search for a character typically interpreted as regular expression syntax	`alert\?` matches `alert?`

About System Auditing

The appliances that are part of the Firepower System generate an audit record for each user interaction with the web interface.

Related Topics

Standard Reports, on page 2051

Audit Records

Firepower Management Centers and 7000 and 8000 Series devices log read-only auditing information for user activity. Audit logs are presented in a standard event view that allows you to view, sort, and filter audit log messages based on any item in the audit view. You can easily delete and report on audit information and can view detailed reports of the changes that users make.

The audit log stores a maximum of 100,000 entries. When the number of audit log entries exceeds 100,000, the appliance prunes the oldest records from the database to reduce the number to 100,000.



Note If you reboot a 7000 or 8000 Series device, then log into the auxiliary CLI as soon as you are able, any commands you execute are not recorded in the audit log until the local web interface is available.

Viewing Audit Records

On a Firepower Management Center or 7000 and 8000 Series devices, you can view a table of audit records. The predefined audit workflow includes a single table view of events. You can manipulate the table view depending on the information you are looking for. You can also create a custom workflow that displays only the information that matches your specific needs.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

You must be an Admin user to perform this procedure.

Procedure

Step 1 Access the audit log workflow using **System** > **Monitoring** > **Audit**.

Step 2 If no events appear, you may need to adjust the time range. For more information, see Event Time Constraints, on page 2165.

Note Events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.

Step 3 You have the following choices:

- To learn more about the contents of the columns in the table, see The System Log, on page 2415.

- To sort and constrain events on the current workflow page, see Using Table View Pages, on page 2158.
- To navigate between pages in the current workflow, keeping the current constraints, click the appropriate page link at the top left of the workflow page. For more information, see Using Workflows, on page 2150.
- To drill down to the next page in the workflow, see Using Drill-Down Pages, on page 2158.
- To constrain on a specific value, click a value within a row. If you click a value on a drill-down page, you move to the next page and constrain on the value. Note that clicking a value within a row in a table view constrains the table view and does **not** drill down to the next page. See Event View Constraints, on page 2171 for more information.

Tip Table views always include "Table View" in the page name.

- To delete audit records, check the check boxes next to events you want to delete, then click **Delete**, or click **Delete All** to delete all events in the current constrained view.
- To bookmark the current page so you can quickly return to it, click **Bookmark This Page**. For more information, see Bookmarks, on page 2174.
- To navigate to the bookmark management page, click **View Bookmarks**. For more information, see Bookmarks, on page 2174.
- To generate a report based on the data in the current view, click **Report Designer**. For more information, see Creating a Report Template from an Event View, on page 2055.
- To view a summary of a change recorded in the audit log, click **Compare** next to applicable events in the **Message** column. For more information, see Using the Audit Log to Examine Changes, on page 2419.

Related Topics

Event View Constraints, on page 2171

Audit Log Workflow Fields

The following table describes the audit log fields that can be viewed and searched.

Table 313: Audit Log Fields

Field	Description
Time	Time and date that the appliance generated the audit record.
User	User name of the user that triggered the audit event.
Subsystem	The full menu path the user followed to generate the audit record. For example, **System** > **Monitoring** > **Audit** is the menu path to view the audit log. In a few cases where a menu path is not relevant, the Subsystem field displays only the event type. For example, **Login** classifies user login attempts.
Message	The action the user performed or the button the user clicked on the page. For example, `Page View` signifies that the user simply viewed the page indicated in the Subsystem, while `Save` means that the user clicked the **Save** button on the page. Changes made to the Firepower System appear with a **Compare icon** that you can click to see a summary of the changes.

Field	Description
Source IP	IP address associated with the host used by the user. Note: When searching this field you must type a specific IP address; you cannot use IP ranges when searching audit logs.
Domain	The current domain of the user when the audit event was triggered. This field is only present if you have ever configured the Firepower Management Center for multitenancy.
Configuration Change (search only)	Specifies whether to view audit records of configuration changes in the search results. (`yes` or `no`)
Count	The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows. This field is not searchable.

Related Topics

Event Searches, on page 2177

The Audit Events Table View

You can change the layout of the event view or constrain the events in the view by a field value. When disabling columns, after you click the **Close** ( ) in the column heading that you want to hide, in the pop-up window that appears, click **Apply**. When you disable a column, it is disabled for the duration of your session (unless you add it back later). Note that when you disable the first column, the Count column is added.

To hide or show other columns, or to add a disabled column back to the view, select or clear the appropriate check boxes before you click **Apply**.

Clicking a value within a row in a table view constrains the table view and does not drill down to the next page in the workflow.

Tip Table views always include “Table View” in the page name.

Related Topics

Using Workflows, on page 2150

Using the Audit Log to Examine Changes

You can use the audit log to view detailed reports of some of the changes to your system. These reports compare the current configuration of your system to its most recent configuration before a supported change was made.

The Compare Configurations page displays the differences between the system configuration before changes and the running configuration in a side-by-side format. The audit event type, time of last modification, and name of the user who made the change are displayed in the title bar above each configuration.

Differences between the two configurations are highlighted:

- Blue indicates that the highlighted setting is different in the two configurations, and the difference is noted in red text.

- Green indicates that the highlighted setting appears in one configuration but not the other.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

You must be an Admin user to perform this procedure.

Procedure

Step 1 Choose **System** > **Monitoring** > **Audit**.

Step 2 Click **Compare** next to an applicable audit log event in the **Message** column.

Tip You can navigate through changes individually by clicking **Previous** or **Next** above the title bar. If the change summary is more than one page long, you can also use the scroll bar on the right to view additional changes.

Suppressing Audit Records

If your auditing policy does not require that you audit specific types of user interactions with the Firepower System, you can prevent those interactions from generating audit records on a Firepower Management Center or 7000 and 8000 Series devices. For example, by default, each time a user views the online help, the Firepower System generates an audit record. If you do not need to keep a record of these interactions, you can automatically suppress them.

To configure audit event suppression, you must have access to an appliance's `admin` user account, and you must be able to either access the appliance's console or open a secure shell.

Caution Make sure that only authorized personnel have access to the appliance and to its `admin` account.

Before you begin

You must be an Admin user to perform this procedure.

Procedure

In the `/etc/sf` directory, create one or more `AuditBlock` files in the following form, where type is one of the types described in Audit Block Types, on page 2421:

`AuditBlock.`*type*

Note If you create an `AuditBlock.type` file for a specific type of audit message, but later decide that you no longer want to suppress them, you must delete the contents of the `AuditBlock.type` file but leave the file itself on the Firepower System.

Audit Block Types

The contents for each audit block type must be in a specific format, as described in the following table. Make sure you use the correct capitalization for the file names. Note also that the contents of the files are case sensitive.

Note that when you add an `AuditBlock` file, an audit record with a subsystem of `Audit` and a message of `Audit Filter type Changed` is added to the audit events. For security reasons, this audit record **cannot** be suppressed.

Table 314: Audit Block Types

Type	Description
Address	Create a file named `AuditBlock.address` and include, one per line, each IP address that you want to suppress from the audit log. You can use partial IP addresses provided that they map from the beginning of the address. For example, the partial address `10.1.1` matches addresses from `10.1.1.0` through `10.1.1.255`.
Message	Create a file named `AuditBlock.message` and include, one per line, the message substrings that you want to suppress. Note that substrings are matched so that if you include `backup` in your file, all messages that include the word `backup` are suppressed.
Subsystem	Create a file named `AuditBlock.subsystem` and include, one per line, each subsystem that you want to suppress. Note that substrings are **not** matched. You must use exact strings. See Audited Subsystems, on page 2421 for a list of subsystems that are audited.
User	Create a file named `AuditBlock.user` and include, one per line, each user account that you want to suppress. You can use partial string matching provided that they map from the beginning of the username. For example, the partial username `IPSAnalyst` matches the user names `IPSAnalyst1` and `IPSAnalyst2`.

Audited Subsystems

The following table lists audited subsystems.

Table 315: Subsystem Names

Name	Includes user interactions with...
Admin	Administrative features such as system and access configuration, time synchr backup and restore, device management, user account management, and sche
Alerting	Alerting functions such as email, SNMP, and syslog alerting

Name	Includes user interactions with...
Audit Log	Audit event views
Audit Log Search	Audit event searches
Command Line	Command line interface
Configuration	Email alerting
COOP	Continuity of operations feature
Date	Date and time range for event views
Default Subsystem	Options that do not have assigned subsystems
Detection & Prevention Policy	Menu options for intrusion policies
Error	System-level errors
eStreamer	eStreamer configuration
EULA	Reviewing the end user license agreement
Events	Intrusion and discovery event views
Events Clipboard	Intrusion event clipboard
Events Reviewed	Reviewed intrusion events
Events Search	Any event search
Failed to install rule update `rule_update_id`	Installing rule updates
Header	Initial presentation of the user interface after a user logs in
Health	Health monitoring
Health Events	Health monitoring event views
Help	Online help
High Availability	Establishing and managing Firepower Management Centers in high availability p
IDS Impact Flag	Impact flag configuration for intrusion events
IDS Policy	Intrusion policies
IDSRule sid:`sig_id` rev:`rev_num`	Intrusion rules by SID
Incidents	Intrusion incidents
Install	Installing updates
Intrusion Events	Intrusion events

Name	Includes user interactions with...
Login	Web interface login and logout functions
Logout	Web interface logout functions
Menu	Any menu option
Configuration export > `config_type` > `config_name`	Importing configurations of a specific type and name
Permission Escalation	User role escalation
Preferences	User preferences, such as the time zone for a user account and individual ever
Policy	Any policy, including intrusion policies
Register	Registering devices on a FMC
RemoteStorageDevice	Configuring remote storage devices
Reports	Report listing and report designer features
Rules	Intrusion rules, including the intrusion rules editor and the rule importation p
Rule Update Import Log	Viewing the rule update import log
Rule Update Install	Installing rule updates
Session Expiration	Web interface session timeouts
Status	Syslog, as well as host and performance statistics
System	Various system-wide settings
Task Queue	Viewing background process status
Users	Creating and modifying user accounts and roles

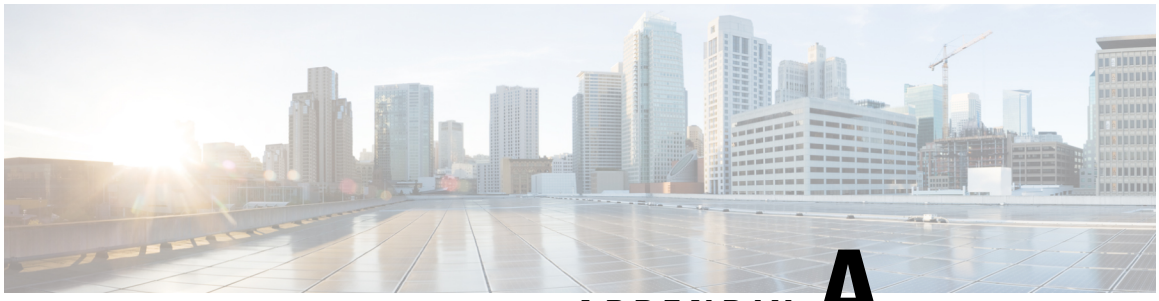
About Sending Audit Logs to an External Location

To send audit logs to an external location from the FMC, see:

- Audit Logs, on page 897
- Audit Log Certificate, on page 900

For Classic devices, see:

- Stream Audit Logs from Classic Devices, on page 934
- Require Valid Audit Log Server Certificates for Classic Devices, on page 935
- Obtain a Signed Client Certificate for Secure Audit Log Streaming on a 7000/8000 Series Device, on page 942



APPENDIX A

Security, Internet Access, and Communication Ports

The following topics provide information on system security, internet access, and communication ports:

- Security Requirements, on page 2425
- Cisco Clouds, on page 2425
- Internet Access Requirements, on page 2426
- Communication Port Requirements, on page 2427

Security Requirements

To safeguard the Firepower Management Center, you should install it on a protected internal network. Although the FMC is configured to have only the necessary services and ports available, you must make sure that attacks cannot reach it (or any managed devices) from outside the firewall.

If the FMC and its managed devices reside on the same network, you can connect the management interfaces on the devices to the same protected internal network as the FMC. This allows you to securely control the devices from the FMC. You can also configure multiple management interfaces to allow the FMC to manage and isolate traffic from devices on other networks.

Regardless of how you deploy your appliances, inter-appliance communication is encrypted. However, you must still take steps to ensure that communications between appliances cannot be interrupted, blocked, or tampered with; for example, with a distributed denial of service (DDoS) or man-in-the-middle attack.

Cisco Clouds

The Firepower System uses Cisco's Collective Security Intelligence (CSI) cloud to obtain the threat intelligence data it uses to assess risk for files and to obtain URL category and reputation. With the correct licenses, you can specify communications options for the AMP for Networks and URL Filtering features.

Additional information:

- **Advanced Malware Protection**

  The public cloud is configured by default; to make changes, see Change AMP Options, on page 1415.

- **URL filtering**

For information, see:

- URL Filtering Options, on page 1235
- Enable URL Filtering Using Category and Reputation, on page 1235

Internet Access Requirements

By default, the system is configured to connect to the internet on ports 443/tcp (HTTPS) and 80/tcp (HTTP). If you do not want your appliances to have direct access to the internet, you can configure a proxy server. For many features, your location can determine which resources the system access.

In most cases, it is the FMC that accesses the internet. Both FMCs in a high availability pair should have internet access. Depending on the feature, sometimes both peers access the internet, and sometimes only the active peer does.

Sometimes, managed devices also access the internet. For example, if your malware protection configuration uses dynamic analysis, managed devices submit files directly to the Threat Grid cloud. Or, you may synchronize a device to an external NTP server.

Table 316: Internet Access Requirements

Feature	Reason	FMC High Availability	Resource
AMP for Networks	Malware cloud lookups.	Both peers perform lookups.	See Required Server Addresses for Proper Cisco Secure Endpoint & Malware Analytics Operations.
	Download signature updates for file preclassification and local malware analysis.	Active peer downloads, syncs to standby.	updates.vrt.sourcefire.com amp.updates.vrt.sourcefire.com
	Submit files for dynamic analysis (managed devices). Query for dynamic analysis results (FMC).	Both peers query for dynamic analysis reports.	fmc.api.threatgrid.com fmc.api.threatgrid.eu
AMP for Endpoints integration	Receive malware events detected by AMP for Endpoints from the AMP cloud. Display malware events detected by the Firepower system in AMP for Endpoints. Use centralized file Block and Allow lists created in AMP for Endpoints to override dispositions from the AMP cloud.	Both peers receive events. You must also configure the cloud connection on both peers (configuration is not synced).	See Required Server Addresses for Proper Cisco Secure Endpoint & Malware Analytics Operations.
Security Intelligence	Download Security Intelligence feeds.	Active peer downloads, syncs to standby.	intelligence.sourcefire.com

Feature	Reason	FMC High Availability	Resource
URL filtering	Download URL category and reputation data. Manually query (look up) URL category and reputation data. Query for uncategorized URLs.	Active FMC downloads, syncs to standby.	database.brightcloud.com service.brightcloud.com
Cisco Smart Licensing	Communicate with the Cisco Smart Software Manager.	Active peer communicates.	tools.cisco.com:443 www.cisco.com
System updates	Download updates *directly* to the FMC: • System software • Intrusion rules • Vulnerability database (VDB) • Geolocation database (GeoDB)	Update intrusion rules, the VDB, and the GeoDB on the active peer, which then syncs to the standby. Upgrade the system software independently on each peer. See the Cisco Firepower Management Center Upgrade Guide, Version 6.0–7.0.	cisco.com sourcefire.com
Time synchronization	Synchronize time in your deployment. Not supported with a proxy server.	Any appliance using an external NTP server must have internet access.	0.sourcefire.pool.ntp.org 1.sourcefire.pool.ntp.org 2.sourcefire.pool.ntp.org 3.sourcefire.pool.ntp.org
RSS feeds	Display the Cisco Threat Research Blog on the dashboard.	Any appliance displaying RSS feeds must have internet access.	feeds.feedburner.com
Whois	Request whois information for an external host. Not supported with a proxy server.	Any appliance requesting whois information must have internet access.	The whois client tries to guess the right server to query. If it cannot guess, it uses: • NIC handles: whois.networksolutions.com • IPv4 addresses and network names: whois.arin.net

Communication Port Requirements

Firepower appliances communicate using a two-way, SSL-encrypted communication channel on port 8305/tcp. This port *must* remain open for basic intra-platform communication.

Other ports allow secure management, as well as access to external resources required by specific features. In general, feature-related ports remain closed until you enable or configure the associated feature. Do *not* change or close an open port until you understand how this action will affect your deployment.

Table 317: Firepower Communication Port Requirements

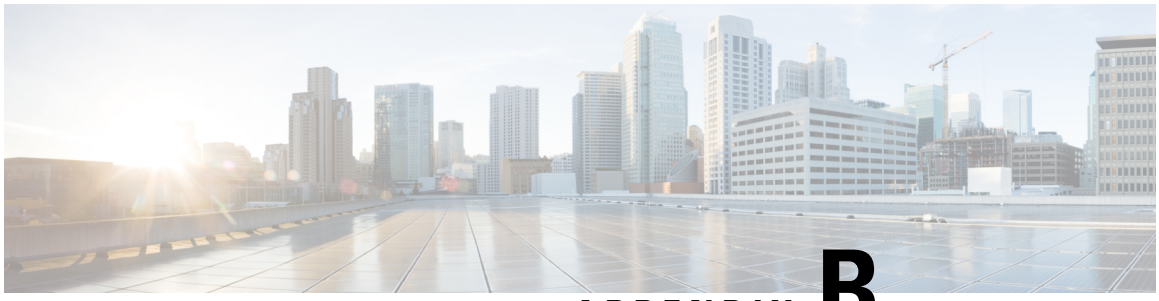
Port	Protocol/Feature	Platforms	Direction	Details
7/UDP	UDP/audit logging	FMC, classic	Outbound	Verify connectivity with the syslog server when configuring audit logging.
22/tcp	SSH	FMC Any device	Inbound	Secure remote connections to the appliance.
25/tcp	SMTP	FMC	Outbound	Send email notices and alerts.
53/tcp 53/udp	DNS	FMC Any device	Outbound	DNS
67/udp 68/udp	DHCP	FMC Any device	Outbound	DHCP
80/tcp	HTTP	FMC 7000/8000 series	Outbound	Display RSS feeds in the dashboard.
80/tcp	HTTP	FMC	Outbound	Download or query URL category and reputation data (port 443 also required).
80/tcp	HTTP	FMC	Outbound	Download custom Security Intelligence feeds over HTTP.
123/udp	NTP	FMC Any device	Outbound	Synchronize time.
161/udp	SNMP	FMC Any device	Inbound	Allow access to MIBs via SNMP polling.
162/udp	SNMP	FMC Any device	Outbound	Send SNMP alerts to a remote trap server.
389/tcp 636/tcp	LDAP	FMC FTD 7000/8000 series	Outbound	Communicate with an LDAP server for external authentication. Obtain metadata for detected LDAP users (FMC only). Configurable.
443/tcp	HTTPS	FMC 7000/8000 series	Inbound	Access the web interface.
443/tcp	HTTPS	FMC FTD	Inbound	Communicate with integrated and third-party products using the Firepower REST API, including Cisco Terminal Services (TS) Agent.

Port	Protocol/Feature	Platforms	Direction	Details
443/tcp	HTTPS	FMC Any device	Outbound	Send and receive data from the internet. For details, see Internet Access Requirements, on page 2426.
443	HTTPS	FMC	Outbound	Communicate with the AMP cloud (public or private) See also information for port 32137.
443	HTTPS	FMC	Inbound and Outbound	Integrate with AMP for Endpoints
514/udp	Syslog (alerts)	FMC Any device	Outbound	Send alerts to a remote syslog server.
623/udp	SOL/LOM	FMC 7000/8000 series	Inbound	Lights-Out Management (LOM) using a Serial Over LAN (SOL) connection.
885/tcp	Captive portal	Any device	Inbound	Communicate with a captive portal identity source.
1500/tcp 2000/tcp	Database access	FMC	Inbound	Allow read-only access to the event database by a third-party client.
1812/udp 1813/udp	RADIUS	FMC FTD 7000/8000 series	Outbound	Communicate with a RADIUS server for external authentication and accounting. Configurable.
3306/tcp	User Agent	FMC	Inbound	Communicate with User Agents.
5222/tcp	ISE	FMC	Outbound	Communicate with an ISE identity source.
6514/tcp	Syslog (audit events)	FMC 7000/8000 series NGIPSv ASA FirePOWER	Outbound	Send audit logs to a remote syslog server, when TLS is configured.
8302/tcp	eStreamer	FMC 7000/8000 series	Inbound	Communicate with an eStreamer client.
8305/tcp	Appliance communications	FMC Any device	Both	Securely communicate between appliances in a deployment. Configurable. If you change this port, you must change it for *all* appliances in the deployment. We recommend you keep the default.
8307/tcp	Host input client	FMC	Inbound	Communicate with a host input client.

Port	Protocol/Feature	Platforms	Direction	Details
32137/tcp	AMP for Networks	FMC	Outbound	Communicate with the Cisco AMP cloud. This is a legacy configuration. We recommend you use the default (443).

Related Topics

Identifying the LDAP Authentication Server, on page 82
Configuring RADIUS Connection Settings, on page 96



APPENDIX B

Classic Device Command Line Reference

The Classic device CLI reference applies to:

- 7000 and 8000 Series
- ASA FirePOWER
- NGIPSv

For other Firepower appliances:

- Firepower Threat Defense: See the Cisco Firepower Threat Defense Command Reference.
- Firepower Management Center: Not available. Instead, the FMC supports Linux shell access under Cisco Technical Assistance Center (TAC) supervision.

- About the Classic Device CLI, on page 2431
- Classic Device CLI Management Commands, on page 2432
- Classic Device CLI Show Commands, on page 2435
- Classic Device CLI Configuration Commands, on page 2464
- Classic Device CLI System Commands, on page 2483

About the Classic Device CLI

After you log into a Classic device (7000 and 8000 Series, ASA FirePOWER, NGIPSv) via the CLI (see Logging Into the CLI on 7000/8000 Series, ASA FirePOWER, and NGIPSv Devices, on page 27), you can use the commands described in this appendix to view, configure, and troubleshoot your device.

Note If you reboot a 7000 or 8000 Series device and then log in to the CLI as soon as you are able, any commands you execute are not recorded in the audit log until the web interface is available.

Note that CLI commands are case-insensitive with the exception of parameters whose text is not part of the CLI framework, such as user names and search filters.

Related Topics

Firepower System User Interfaces, on page 21

Classic Device CLI Modes

The CLI encompasses four modes. The default mode, CLI Management, includes commands for navigating within the CLI itself. The remaining modes contain commands addressing three different areas of classic device functionality; the commands within these modes begin with the mode name: `system`, `show`, or `configure`.

When you enter a mode, the CLI prompt changes to reflect the current mode. For example, to display version information about system components, you can enter the full command at the standard CLI prompt:

```
> show version
```

If you have previously entered `show` mode, you can enter the command without the `show` keyword at the show mode CLI prompt:

```
show> version
```

Classic Device CLI Access Levels

Within each mode, the commands available to a user depend on the user's CLI access. When you create a user account, you can assign it one of the following CLI access levels:

- Basic — The user has read-only access and cannot run commands that impact system performance.
- Configuration — The user has read-write access and can run commands that impact system performance.
- None — The user is unable to log into the CLI.

On 7000 and 8000 Series devices, you can assign command line permissions on the User Management page in the local web interface. On NGIPSv and ASA FirePOWER, you assign command line permissions using the CLI.

Classic Device CLI Management Commands

The CLI management commands provide the ability to interact with the CLI. These commands do not affect the operation of the device.

configure password

Allows the current user to change their password. After issuing the command, the CLI prompts the user for their current (or old) password, then prompts the user to enter the new password twice.

Access

Basic

Syntax

```
configure password
```

Example

```
> configure password
Enter current password:
Enter new password:
Confirm new password:
```

exit

Moves the CLI context up to the next highest CLI context level. Issuing this command from the default mode logs the user out of the current CLI session, and is equivalent to issuing the `logout` CLI command.

Access

Basic

Syntax

```
exit
```

Example

```
configure network ipv4> exit
configure network>
```

expert

Invokes the Linux shell.



Caution We strongly recommend that you do not use the Linux shell unless directed by Cisco TAC or explicit instructions in the user documentation. For more information, see Firepower System User Accounts, on page 19.

Access

Configuration

Syntax

```
expert
```

Example

```
> expert
```

history

Displays the command line history for the current session.

Access

Basic

Syntax

```
history limit
```

where `limit` sets the size of the history list. To set the size to unlimited, enter zero.

Example

```
history 25
```

logout

Logs the current user out of the current CLI console session.

Access

Basic

Syntax

```
logout
```

Example

```
> logout
```

? (question mark)

Displays context-sensitive help for CLI commands and parameters. Use the question mark (`?`) command as follows:

- To display help for the commands that are available within the current CLI context, enter a question mark (`?`) at the command prompt.
- To display a list of the available commands that start with a particular character set, enter the abbreviated command immediately followed by a question mark (`?`).
- To display help for a command's legal arguments, enter a question mark (`?`) in place of an argument at the command prompt.

Note that the question mark (`?`) is not echoed back to the console.

Access

Basic

Syntax

```
?
abbreviated_command ?
command [arguments] ?
```

Example

```
> ?
```

Classic Device CLI Show Commands

Show commands provide information about the state of the device. These commands do not change the operational mode of the device and running them has minimal impact on system operation. Most show commands are available to all CLI users; however, only users with configuration CLI access can issue the `show user` command.

access-control-config

Displays the currently deployed access control configurations, including:

- Security Intelligence settings
- Names of any subpolicies the access control policy invokes
- Intrusion variable set data
- Logging settings
- Other advanced settings, including policy-level performance, preprocessing, and general settings

Also displays policy-related connection information, such as source and destination port data (including type and code for ICMP entries) and the number of connections that matched each access control rule (hit counts).

Access

Basic

Syntax

```
show access-control-config
```

Example

```
> show access-control-config
```

alarms

Displays currently active (failed/down) hardware alarms on the device. This command is not available on NGIPSv and ASA FirePOWER devices.

Access

Basic

Syntax

```
show alarms
```

Example

```
> show alarms
```

arp-tables

Displays the Address Resolution Protocol tables applicable to your network. This command is not available on NGIPSv and ASA FirePOWER.

Access

Basic

Syntax

```
show arp-tables
```

Example

```
> show arp-tables
```

audit-log

Displays the audit log in reverse chronological order; the most recent audit log events are listed first.

Access

Basic

Syntax

```
show audit-log
```

Example

```
> show audit-log
```

audit_cert

Displays the current audit log client certificate.

Access

Basic

Syntax

```
show audit_cert
```

Example

```
> show audit_cert
```

bypass

On 7000 or 8000 Series devices, lists the inline sets in use and shows the bypass mode status of those sets as one of the following:

- `armed`—the interface pair is configured to go into hardware bypass if it fails (**Bypass Mode: Bypass**), or has been forced into fail-close with the **`configure bypass close`** command
- `engaged`—the interface pair has failed open or has been forced into hardware bypass with the **`configure bypass open`** command
- `off`—the interface pair is set to fail-close (**Bypass Mode: Non-Bypass**); packets are blocked if the interface pair fails

Access

Basic

Syntax

```
show bypass
```

Example

```
> show bypass
s1p1 ↔ s1p2: status 'armed'
s1p1 ↔ s1p2: status 'engaged'
```

high-availability Commands

Displays information about high-availability configuration, status, and member devices or stacks. This command is not available on NGIPSv and ASA FirePOWER devices.

Access

Basic

config

Displays the high-availability configuration on the device.

Syntax

```
show high-availability config
```

Example

```
> show high-availability config
```

high-availability ha-statistics

Displays state sharing statistics for a device in a high-availability pair.

Syntax

```
show high-availability ha-statistics
```

Example

```
> show high-availability ha-statistics
```

cpu

Displays the current CPU usage statistics appropriate for the platform for all CPUs on the device.

For 7000 and 8000 Series devices, the following values are displayed:

- CPU — Processor number.
- Load — The CPU utilization, represented as a number from 0 to 100. 0 is not loaded and 100 is completely loaded.

For NGIPSv and ASA FirePOWER, the following values are displayed:

- CPU — Processor number.
- %user — Percentage of CPU utilization that occurred while executing at the user level (application).

- %nice — Percentage of CPU utilization that occurred while executing at the user level with nice priority.
- %sys — Percentage of CPU utilization that occurred while executing at the system level (kernel). This does not include time spent servicing interrupts or softirqs. A softirq (software interrupt) is one of up to 32 enumerated software interrupts that can run on multiple CPUs at once.
- %iowait — Percentage of time that the CPUs were idle when the system had an outstanding disk I/O request.
- %irq — Percentage of time spent by the CPUs to service interrupts.
- %soft — Percentage of time spent by the CPUs to service softirqs.
- %steal — Percentage of time spent in involuntary wait by the virtual CPUs while the hypervisor was servicing another virtual processor.
- %guest — Percentage of time spent by the CPUs to run a virtual processor.
- %idle — Percentage of time that the CPUs were idle and the system did not have an outstanding disk I/O request.

Access

Basic

Syntax

```
show cpu [procnum]
```

where `procnum` is the number of the processor for which you want the utilization information displayed. Valid values are 0 to one less than the total number of processors on the system. If `procnum` is used for a 7000 or 8000 Series device, it is ignored because for that platform, utilization information can only be displayed for all processors.

```
> show cpu
```

database Commands

The `show database` commands configure the device's management interface.

Access

Basic

processes

Displays a list of running database queries.

Access

Basic

Syntax

```
show database processes
```

Example

```
> show database processes
```

slow-query-log

Displays the slow query log of the database.

Access

Basic

Syntax

```
show database slow-query-log
```

Example

```
> show database slow-query-log
```

device-settings

Displays information about application bypass settings specific to the current device.

Access

Basic

Syntax

```
show device-settings
```

Example

```
> show device-settings
```

disk

Displays the current disk usage.

Access

Basic

Syntax

```
show disk
```

Example

```
> show disk
```

disk-manager

Displays detailed disk usage information for each part of the system, including silos, low watermarks, and high watermarks.

Access

Basic

Syntax

```
show disk-manager
```

Example

```
> show disk-manager
```

dns

Displays the current DNS server addresses and search domains.

Access

Basic

Syntax

```
show dns
```

Example

```
> show dns
```

fan-status

Displays the current status of hardware fans. This command is not available on NGIPSv and ASA FirePOWER devices.

Access

Basic

Syntax

```
show fan-status
```

Example

```
> show fan-status
```

fastpath-rules

Displays the currently configured 8000 Series fastpath rules. This command is only available on 8000 Series devices.

Access

Basic

Syntax

```
show fastpath-rules
```

Example

```
> show fastpath-rules
```

gui

Displays the current state of the web interface. This command is not available on NGIPSv and ASA FirePOWER.

Access

Basic

Syntax

```
show gui
```

Example

```
> show gui
```

hostname

Displays the device's host name and appliance UUID. If you edit the host name of a device using the CLI, confirm that the changes are reflected on the managing Firepower Management Center. In some cases, you may need to edit the device management settings manually.

Access

Basic

Syntax

```
show hostname
```

Example

```
> show hostname
```

hosts

Displays the contents of an ASA FirePOWER module's /etc/hosts file.

Access

Basic

Syntax

```
show hosts
```

Example

```
> show hosts
```

hyperthreading

Displays whether hyperthreading is enabled or disabled. This command is not available on ASA FirePOWER.

Access

Basic

Syntax

```
show hyperthreading
```

Example

```
> show hyperthreading
```

inline-sets

Displays configuration data for all inline security zones and associated interfaces. This command is not available on ASA FirePOWER.

Access

Basic

Syntax

```
show inline-sets
```

Example

```
> show inline-sets
```

interfaces

If no parameters are specified, displays a list of all configured interfaces. If a parameter is specified, displays detailed information about the specified interface.

Access

Basic

Syntax

```
show interfaces interface
```

where *interface* is the specific interface for which you want the detailed information.

Example

```
> show interfaces
```

ifconfig

Displays the interface configuration for an ASA FirePOWER module.

Access

Basic

Syntax

```
show ifconfig
```

Example

```
> show ifconfig
```

lcd

Displays whether the LCD hardware display is enabled or disabled. This command is not available on NGIPSv and ASA FirePOWER.

Access

Basic

Syntax

```
show lcd
```

Example

```
> show lcd
```

link-aggregation Commands

The `show link-aggregation` commands display configuration and statistics information for link aggregation groups (LAGs). This command is not available on NGIPSv and ASA FirePOWER devices.

Access

Basic

configuration

Displays configuration details for each configured LAG, including LAG ID, number of interfaces, configuration mode, load-balancing mode, LACP information, and physical interface type.

Access

Basic

Syntax

```
show link-aggregation configuration
```

Example

```
> show link-aggregation configuration
```

statistics

Displays statistics, per interface, for each configured LAG, including status, link state and speed, configuration mode, counters for received and transmitted packets, and counters for received and transmitted bytes.

Access

Basic

Syntax

```
show link-aggregation statistics
```

Example

```
> show link-aggregation statistics
```

link-state

Displays type, link, and speed, duplex state, and bypass mode of the ports on the device. This command is not available on ASA FirePOWER devices.

Access

Basic

Syntax

```
show link-state
```

Example

```
> show link-state
```

log-ips-connection

Displays whether the logging of connection events that are associated with logged intrusion events is enabled or disabled.

Access

Basic

Syntax

```
show log-ips-connection
```

Example

```
> show log-ips-connection
```

managers

Displays the configuration and communication status of the Firepower Management Center. Registration key and NAT ID are only displayed if registration is pending.

If a device is configured as a secondary device in a stacked configuration, information about both the managing FMC and the primary device is displayed.

Access

Basic

Syntax

```
show managers
```

Example

```
> show managers
```

memory

Displays the total memory, the memory in use, and the available memory for the device.

Access

Basic

Syntax

```
show memory
```

Example

```
> show memory
```

model

Displays model information for the device.

Access

Basic

Syntax

```
show model
```

Example

```
> show model
```

mpls-depth

Displays the number of MPLS layers configured on the management interface, from 0 to 6. This command is not available on NGIPSv and ASA FirePOWER.

Access

Basic

Syntax

```
show mpls-depth
```

Example

```
> show mpls-depth
```

NAT Commands

The `show nat` commands display NAT data and configuration information for the management interface. This command is not available on NGIPSv and ASA FirePOWER devices.

Access

Basic

active-dynamic

Displays NAT flows translated according to dynamic rules. These entries are displayed when a flow matches a rule, and persist until the rule has timed out. Therefore, the list can be inaccurate. Timeouts are protocol dependent: ICMP is 5 seconds, UDP is 120 seconds, TCP is 3600 seconds, and all other protocols are 60 seconds.

Syntax

```
show nat active-dynamic
```

Example

```
> show nat active-dynamic
```

active-static

Displays NAT flows translated according to static rules. These entries are displayed as soon as you deploy the rule to the device, and the list does not indicate active flows that match a static NAT rule.

Syntax

```
show nat active-static
```

Example

```
> show nat active-static
```

allocators

Displays information for all NAT allocators, the pool of translated addresses used by dynamic rules.

Syntax

```
show nat allocators
```

Example

```
> show nat allocators
```

config

Displays the current NAT policy configuration for the management interface.

Syntax

```
show nat config
```

Example

```
> show nat config
```

dynamic-rules

Displays dynamic NAT rules that use the specified allocator ID.

Syntax

```
show nat dynamic-rules allocator_id
```

where `allocator_id` is a valid allocator ID number.

Example

```
> show nat dynamic-rules 9
```

flows

Displays the number of flows for rules that use the specified allocator ID.

Syntax

```
show nat flows allocator-id
```

where *allocator_id* is a valid allocator ID number.

Example

```
> show nat flows 81
```

static-rules

Displays all static NAT rules.

Syntax

```
show nat static-rules
```

Example

```
> show nat static-rules
```

netstat

Displays the active network connections for an ASA FirePOWER module.

Access

Basic

Syntax

```
show netstat
```

Example

```
> show netstat
```

network

Displays the IPv4 and IPv6 configuration of the management interface, its MAC address, and HTTP proxy address, port, and username if configured.

Access

Basic

Syntax

```
show network
```

Example

```
> show network
```

network-modules

Displays all installed modules and information about them, including serial numbers. This command is not available on NGIPSv and ASA FirePOWER.

Access

Basic

Syntax

```
show network-modules
```

Example

```
> show network-modules
```

network-static-routes

Displays all configured network static routes and information about them, including interface, destination address, network mask, and gateway address.

Access

Basic

Syntax

```
show network-static-routes
```

Example

```
> show network-static-routes
```

ntp

Displays the ntp configuration.

Access

Basic

Syntax

```
show ntp
```

Example

```
> show ntp
```

perfstats

Displays performance statistics for the device.

Access

Basic

Syntax

```
show perfstats
```

Example

```
> show perfstats
```

portstats

Displays port statistics for all installed ports on the device. This command is not available on NGIPSv and ASA FirePOWER.

Access

Basic

Syntax

```
show portstats [copper | fiber | internal | external | all]
```

where copper specifies for all copper ports, fiber specifies for all fiber ports, internal specifies for all internal ports, external specifies for all external (copper and fiber) ports, and all specifies for all ports (external and internal).

Example

```
> show portstats fiber
```

power-supply-status

Displays the current state of hardware power supplies. This command is not available on NGIPSv and ASA FirePOWER.

Note If an 8000 Series managed device experiences a power failure, it may take up to 15 minutes for the show power-supply-status CLI command to reflect the correct status.

Access

Basic

Syntax

```
show power-supply-status
```

Example

```
> show power-supply-status
```

process-tree

Displays processes currently running on the device, sorted in tree format by type.

Access

Basic

Syntax

```
show process-tree
```

Example

```
> show process-tree
```

processes

Displays processes currently running on the device, sorted by descending CPU usage.

Access

Basic

Syntax

```
show processes sort-flag filter
```

where *sort-flag* can be `-m` to sort by memory (descending order), `-u` to sort by username rather than the process name, or `verbose` to display the full name and path of the command. The *filter* parameter specifies the search term in the command or username by which results are filtered. The header row is still displayed.

Example

```
> show processes -u user1
```

route

Displays the routing information for an ASA FirePOWER module.

Access

Basic

Syntax

```
show route
```

Example

```
> show route
```

routing-table

If no parameters are specified, displays routing information for all virtual routers. If parameters are specified, displays routing information for the specified router and, as applicable, its specified routing protocol type. All parameters are optional. This command is not available on NGIPSv and ASA FirePOWER.

Access

Basic

Syntax

```
show routing-table name [ ospf | rip | static ]
```

where *name* is the name of the specific router for which you want information, and `ospf`, `rip`, and `static` specify the routing protocol type.

Example

```
> show routing-table Vrouter1 static
```

serial-number

Displays the chassis serial number. This command is not available on NGIPSv.

Access

Basic

Syntax

```
show serial-number
```

Example

```
> show serial-number
```

ssl-policy-config

Displays the currently deployed SSL policy configuration, including policy description, default logging settings, all enabled SSL rules and rule configurations, trusted CA certificates, and undecryptable traffic actions.

Access

Basic

Syntax

```
show ssl-policy-config
```

Example

```
> show ssl-policy-config
```

stacking

Shows the stacking configuration and position on managed devices; on devices configured as primary, also lists data for all secondary devices. For stacks in a high-availability pair, this command also indicates that the stack is a member of a high-availability pair. The user must use the web interface to enable or (in most cases) disable stacking; if stacking is not enabled, the command will return `Stacking not currently configured`. This command is not available on NGIPSv and ASA FirePOWER.

Access

Basic

Syntax

```
show stacking
```

Example

```
> show stacking
```

summary

Displays a summary of the most commonly used information (version, type, UUID, and so on) about the device. For more detailed information, see the following `show` commands: `version`, `interfaces`, `device-settings`, and `access-control-config`.

Access

Basic

Syntax

```
show summary
```

Example

```
> show summary
```

syslog

Displays the system log in reverse chronological order. You can optionally specify a filter to display specific records based on content and the number of records to display per page view (the default is 25).

Access

Basic

Syntax

```
show syslog ["filter" records_per_page]
```

where *filter* specifies a Grep-compatible search filter and *records_per_page* specifies the number of records to display with each page view. See Syntax for System Log Filters, on page 2416 for more information on search filters.

Example

```
> show syslog "ssh" 20
```

The system displays the 20 most recent syslog records containing the string "ssh". To display the next 20 records, press Enter; to stop the display enter q.

time

Displays the current date and time in UTC and in the local time zone configured for the current user.

Access

Basic

Syntax

```
show time
```

Example

```
> show time
```

traffic-statistics

If no parameters are specified, displays details about bytes transmitted and received from all ports. If a port is specified, displays that information only for the specified port. You cannot specify a port for ASA FirePOWER modules; the system displays only the data plane interfaces.

Note In some situations the output of this command may show packet drops when, in point of fact, the device is not dropping traffic. Drop counters increase when malformed packets are received. A malformed packet may be missing certain information in the header or it may have failed a cyclical-redundancy check (CRC). Typically, common root causes of malformed packets are data link layer issues such as bad cables or a bad interface. The dropped packets are not logged. However, if the source is a reliable transport protocol such as TCP, the packets will be retransmitted.

Access

Basic

Syntax

```
show traffic-statistics port
```

where *port* is the specific port for which you want information.

Example

```
> show traffic-statistics s1p1
```

user

Applicable to NGIPSv only. Displays detailed configuration information for the specified user(s). The following values are displayed:

- Login — the login name

- UID — the numeric user ID
- Auth (`Local` or `Remote`) — how the user is authenticated
- Access (`Basic` or `Config`) — the user's privilege level
- Enabled (`Enabled` or `Disabled`) — whether the user is active
- Reset (`Yes` or `No`) — whether the user must change password at next login
- Exp (`Never` or a number) — the number of days until the user's password must be changed
- Warn (`N/A` or a number) — the number of days a user is given to change their password before it expires
- Str (`Yes` or `No`) — whether the user's password must meet strength checking criteria
- Lock (`Yes` or `No`) — whether the user's account has been locked due to too many login failures
- Max (`N/A` or a number) — the maximum number of failed logins before the user's account is locked

Access

Configuration

Syntax

```
show user username username username ...
```

where *username* specifies the name of the user and the usernames are space-separated.

Example

```
> show user jdoe
```

users

Applicable to NGIPSv and ASA FirePOWER only. Displays detailed configuration information for all local users. The following values are displayed:

- Login — the login name
- UID — the numeric user ID
- Auth (`Local` or `Remote`) — how the user is authenticated
- Access (`Basic` or `Config`) — the user's privilege level
- Enabled (`Enabled` or `Disabled`) — whether the user is active
- Reset (`Yes` or `No`) — whether the user must change password at next login
- Exp (`Never` or a number) — the number of days until the user's password must be changed
- Warn (`N/A` or a number) — the number of days a user is given to change their password before it expires
- Str (`Yes` or `No`) — whether the user's password must meet strength checking criteria

- Lock (`Yes` or `No`) — whether the user's account is locked due to too many login failures
- Max (`N/A` or a number) — the maximum number of failed logins before the user's account is locked

Access

Configuration

Syntax

```
show users
```

Example

```
> show users
```

version

Displays the product version and build. If the `detail` parameter is specified, displays the versions of additional components.



Note The `detail` parameter is not available on ASA with FirePOWER Services.

Access

Basic

Syntax

```
show version [detail]
```

Example

```
> show version
```

virtual-routers

If no parameters are specified, displays a list of all currently configured virtual routers with DHCP relay, OSPF, and RIP information. If parameters are specified, displays information for the specified router, limited by the specified route type. All parameters are optional. This command is not available on NGIPSv and ASA FirePOWER.

Access

Basic

Syntax

```
show virtual-routers [ dhcprelay | ospf | rip ] name
```

where `dhcprelay`, `ospf`, and `rip` specify for route types, and *name* is the name of the specific router for which you want information. If you specify `ospf`, you can then further specify `neighbors`, `topology`, or `lsadb` between the route type and (if present) the router name.

Example

```
> show virtual-routers ospf VRouter2
```

virtual-switches

If no parameters are specified, displays a list of all currently configured virtual switches. If parameters are specified, displays information for the specified switch. This command is not available on NGIPSv and ASA FirePOWER.

Access

Basic

Syntax

```
show virtual-switches name
```

Example

```
> show virtual-switches Vswitch1
```

vmware-tools

Indicates whether VMware Tools are currently enabled on a virtual device. This command is available only on NGIPSv.

VMware Tools is a suite of utilities intended to enhance the performance of the virtual machine. These utilities allow you to make full use of the convenient features of VMware products. The system supports the following plugins on all virtual appliances:

- guestInfo
- powerOps
- timeSync
- vmbackup

For more information about VMware Tools and the supported plugins, see the VMware website (http://www.vmware.com).

Access

Basic

Syntax

```
show vmware-tools
```

Example

```
> show vmware-tools
```

VPN Commands

The `show VPN` commands display VPN status and configuration information for VPN connections. This command is not available on NGIPSv and ASA FirePOWER devices.

Access

Basic

config

Displays the configuration of all VPN connections.

Syntax

```
show vpn config
```

Example

```
> show vpn config
```

config by virtual router

Displays the configuration of all VPN connections for a virtual router.

Syntax

```
show vpn config virtual router
```

Example

```
> show vpn config VRouter1
```

status

Displays the status of all VPN connections.

Syntax

```
show vpn status
```

Example

```
> show vpn status
```

status by virtual router

Displays the status of all VPN connections for a virtual router.

Syntax

```
show vpn status virtual router
```

Example

```
> show vpn status VRouter1
```

counters

Displays the counters for all VPN connections.

Syntax

```
show vpn counters
```

Example

```
> show vpn counters
```

counters by virtual router

Displays the counters of all VPN connections for a virtual router.

Syntax

```
show vpn counters virtual router
```

Example

```
> show vpn counters VRouter1
```

Classic Device CLI Configuration Commands

The configuration commands enable the user to configure and manage the system. These commands affect system operation; therefore, with the exception of Basic-level `configure password`, only users with configuration CLI access can issue these commands.

audit_cert Commands

The `configure audit_cert` commands configure the device's audit log client certificate for secure audit log streaming.

Access

Configuration

delete

Deletes the current client certificate for secure audit log streaming.

Syntax

```
configure audit_cert delete
```

Example

```
> configure audit_cert delete
```

import

Imports a client certificate for secure audit log streaming. After the user enters the command, the CLI prompts the user to provide either a client certificate and private key, or a certificate chain.

Syntax

```
configure audit_cert import
```

Example

```
> configure audit_cert import
***************Import Audit Client Certificate**************
```

```
1 Import Client Certificate and Private Key
2 Import Certificate Chain
0 Exit

************************************************************
Enter choice: 1
Enter your audit client certificate (PEM format) here:
-----BEGIN CERTIFICATE-----
MIIEoTCCA4mgAwIBAgICAR4wDQYJK0ZIhvcNaQALBWAugYICzAJBqNVBATYAiVT
   ...certificate details ...
Tx*FAhnXeUZ78hFepg1yHQMYWTkD7hCqmSN3UkAb1l0IoBcxTA==
-----END CERTIFICATE-----

Enter your private key (PEM format) here:
-----BEGIN RSA PRIVATE KEY-----
miiieOWobabkc3qwaOgVx0Tt61eY83Mrqa+bek_qPetcHRAw6ea4p0TlMVVsE7qr
   ...private key details ...
nRI6QNkoumLUT9EvjF6bFoT3M6eDI7+NdDIhjVeOP*E4+hxEX50jM
-----END RSA PRIVATE KEY-----

Client certificate import succeed, exiting...
```

bypass

On 7000 or 8000 Series devices, places an inline pair in fail-open (hardware bypass) or fail-close mode. You can use this command only when the inline set **Bypass Mode** option is set to **Bypass**.

Note that rebooting a device takes an inline set out of fail-open mode.

Access

Configuration

Syntax

```
configure bypass {open | close} {interface}
```

where `interface` is the name of either hardware port in the inline pair.

Example

```
> configure bypass open s1p1
```

high-availability

Disables or configures bypass for high availability on the device. This command is not available on NGIPSv, ASA FirePOWER, or on devices configured as secondary stack members.

Access

Configuration

Syntax

```
configure high-availability {disable | bypass}
```

Example

```
> configure high-availability disable
```

gui

Enables or disables the device web interface, including the streamlined upgrade web interface that appears during major updates to the system. This command is not available on NGIPSv and ASA FirePOWER.

Access

Configuration

Syntax

```
configure gui [enable | disable]
```

Example

```
> configure gui disable
```

lcd

Enables or disables the LCD display on the front of the device. This command is not available on NGIPSv and ASA FirePOWER.

Access

Configuration

Syntax

```
configure lcd {enable | disable}
```

Example

```
> configure lcd disable
```

log-ips-connections

Enables or disables logging of connection events that are associated with logged intrusion events.

Access

Configuration

Syntax

```
configure log-ips-connections {enable | disable}
```

Example

```
> configure log-ips-connections disable
```

manager Commands

The `configure manager` commands configure the device's connection to its managing Firepower Management Center.

Access

Configuration

add

Configures the device to accept a connection from a managing Firepower Management Center. This command works only if the device is not actively managed.

A unique alphanumeric registration key is always required to register a device to a Firepower Management Center. In most cases, you must provide the hostname or the IP address along with the registration key. However, if the device and the Firepower Management Center are separated by a NAT device, you must enter a unique NAT ID, along with the registration key, and specify `DONTRESOLVE` instead of the hostname.

Syntax

```
configure manager add {hostname | IPv4_address | IPv6_address | DONTRESOLVE} regkey [nat_id]
```

where `{`hostname `|` IPv4_address `|` IPv6_address `|` `DONTRESOLVE}` specifies the DNS host name or IP address (IPv4 or IPv6) of the Firepower Management Center that manages this device. If the Firepower Management Center is not directly addressable, use `DONTRESOLVE`. If you use`DONTRESOLVE`, nat_id is required. regkey is the unique alphanumeric registration key required to register a device to the Firepower Management Center. nat_id is an optional alphanumeric string used during the registration process between the Firepower Management Center and the device. It is required if the hostname is set to `DONTRESOLVE`.

Example

```
> configure manager add DONTRESOLVE abc123 efg456
```

delete

Removes the Firepower Management Center's connection information from the device. This command only works if the device is not actively managed.

Syntax

```
configure manager delete
```

Example

```
> configure manager delete
```

mpls-depth

Configures the number of MPLS layers on the management interface. This command is not available on NGIPSv and ASA FirePOWER.

Access

Configuration

Syntax

```
configure mpls-depth depth
```

where *depth* is a number between 0 and 6.

Example

```
> configure mpls-depth 3
```

network Commands

The `configure network` commands configure the device's management interface.

Access

Configuration

dns searchdomains

Replaces the current list of DNS search domains with the list specified in the command.

Syntax

```
configure network dns searchdomains {searchlist}
```

where `searchlist` is a comma-separated list of domains.

Example

```
> configure network dns searchdomains foo.bar.com,bar.com
```

dns servers

Replaces the current list of DNS servers with the list specified in the command.

Syntax

```
configure nework dns servers {dnslist}
```

where dnslist is a comma-separated list of DNS servers.

Example

```
> configure network dns servers 10.123.1.10,10.124.1.10
```

hostname

Sets the hostname for the device.

Syntax

```
configure network hostname {name}
```

where name is the new hostname.

Example

```
> configure network hostname sfrocks
```

http-proxy

On 7000 & 8000 Series and NGIPSv devices, configures an HTTP proxy. After issuing the command, the CLI prompts the user for the HTTP proxy address and port, whether proxy authentication is required, and if it is required, the proxy username, proxy password, and confirmation of the proxy password.

Use this command on NGIPSv to configure an HTTP proxy server so the virtual device can submit files to the AMP cloud for dynamic analysis.

Syntax

The proxy password can use only alphanumeric characters.

```
configure network http-proxy
```

Example

```
> configure network http-proxy
Manual proxy configuration
Enter HTTP Proxy address:
Enter HTTP Proxy Port:
Use Proxy Authentication? (y/n) [n]:
Enter Proxy Username:
Enter Proxy Password:
Confirm Proxy Password:
```

http-proxy-disable

On 7000 Series, 8000 Series, or NGIPSv devices, deletes any HTTP proxy configuration.

Syntax

```
configure network http-proxy-disable
```

Example

```
> configure network http-proxy-disable
Are you sure that you wish to delete the current
http-proxy configuration? (y/n):
```

ipv4 delete

Disables the IPv4 configuration of the device's management interface.

Syntax

```
configure network ipv4 delete [management_interface]
```

where *management_interface* is the management interface ID. If you do not specify an interface, this command configures the default management interface. This parameter is needed only if you use the **configure management-interface** commands to enable more than one management interface. Multiple management interfaces are supported on 8000 series devices and the ASA 5585-X with FirePOWER services only. Do not specify this parameter for other platforms. The management interface IDs are **eth0** for the default management interface and **eth1** for the optional event interface.

Example

```
> configure network ipv4 delete eth1
```

ipv4 dhcp

Sets the IPv4 configuration of the device's management interface to DHCP. The management interface communicates with the DHCP server to obtain its configuration information.

Syntax

```
configure network ipv4 dhcp [management_interface]
```

where *management_interface* is the management interface ID. DHCP is supported only on the default management interface, so you do not need to use this argument.

Example

```
> configure network ipv4 dhcp
```

ipv4 manual

Manually configures the IPv4 configuration of the device's management interface.

Syntax

```
configure network ipv4 manual ipaddr netmask [gw] [management_interface]
```

where *ipaddr* is the IP address, *netmask* is the subnet mask, and *gw* is the IPv4 address of the default gateway. The *management_interface* is the management interface ID. If you do not specify an interface, this command configures the default management interface. This parameter is needed only if you use the **configure management-interface** commands to enable more than one management interface. Multiple management interfaces are supported on 8000 series devices and the ASA 5585-X with FirePOWER services only. Do not specify this parameter for other platforms. The management interface IDs are **eth0** for the default management interface and **eth1** for the optional event interface.

Example

```
> configure network ipv4 manual 10.123.1.10 255.255.0.0 10.123.1.1
```

ipv6 delete

Disables the IPv6 configuration of the device's management interface.

Syntax

```
configure network ipv6 delete [management_interface]
```

where *management_interface* is the management interface ID. If you do not specify an interface, this command configures the default management interface. This parameter is needed only if you use the **configure management-interface** commands to enable more than one management interface. Multiple management interfaces are supported on 8000 series devices and the ASA 5585-X with FirePOWER services only. Do not specify this parameter for other platforms. The management interface IDs are **eth0** for the default management interface and **eth1** for the optional event interface.

Example

```
> configure network ipv6 delete
```

ipv6 dhcp

Sets the IPv6 configuration of the device's management interface to DHCP. The management interface communicates with the DHCP server to obtain its configuration information.

Syntax

```
configure network ipv6 dhcp [management_interface]
```

where *management_interface* is the management interface ID. DHCP is supported only on the default management interface, so you do not need to use this argument.

Example

```
> configure network ipv6 dhcp
```

ipv6 manual

Manually configures the IPv6 configuration of the device's management interface.

Syntax

```
configure network ipv6 manual ip6addr/ip6prefix [ip6gw] [management_interface]
```

where *ip6addr*/*ip6prefix* is the IP address and prefix length and *ip6gw* is the IPv6 address of the default gateway. The *management_interface* is the management interface ID. If you do not specify an interface, this command configures the default management interface. This parameter is needed only if you use the **configure management-interface** commands to enable more than one management interface. Multiple management interfaces are supported on 8000 series devices and the ASA 5585-X with FirePOWER services only. Do not specify this parameter for other platforms. The management interface IDs are **eth0** for the default management interface and **eth1** for the optional event interface.

Example

```
> configure network ipv6 manual 2001:DB8:3ffe:1900:4545:3:200:f8ff:fe21:67cf 64
```

ipv6 router

Sets the IPv6 configuration of the device's management interface to Router. The management interface communicates with the IPv6 router to obtain its configuration information.

Syntax

```
configure network ipv6 router [management_interface]
```

where *management_interface* is the management interface ID. If you do not specify an interface, this command configures the default management interface. This parameter is needed only if you use the **configure management-interface** commands to enable more than one management interface. Multiple management interfaces are supported on 8000 series devices and the ASA 5585-X with FirePOWER services only. Do not specify this parameter for other platforms. The management interface IDs are **eth0** for the default management interface and **eth1** for the optional event interface.

Example

```
> configure network ipv6 router
```

management-interface disable

Disables a management interface. Multiple management interfaces are supported on 8000 series devices and the ASA 5585-X with FirePOWER services only.

Syntax

```
configure network management-interface disable ethn
```

where *n* is the number of the management interface you want to configure. **eth0** is the default management interface and **eth1** is the optional event interface. Cisco recommends that you leave the eth0 default management interface enabled, with both management and event channels enabled. See Management Interfaces, on page 881 for detailed information about using a separate event interface on the Firepower Management Center and on the managed device.

Example

```
> configure network management-interface disable eth1
```

management-interface disable-event-channel

Disables the event traffic channel on the specified management interface. Multiple management interfaces are supported on 8000 series devices and the ASA 5585-X with FirePOWER services only.

Syntax

```
configure network management-interface disable-event-channel ethn
```

where *n* is the number of the management interface you want to configure. **eth0** is the default management interface and **eth1** is the optional event interface. Cisco recommends that you leave the eth0 default management interface enabled, with both management and event channels enabled. See Management Interfaces, on page 881 for detailed information about using a separate event interface on the Firepower Management Center and on the managed device.

Example

```
> configure network management-interface disable-event-channel eth1
```

management-interface disable-management-channel

Disables the management traffic channel on the specified management interface. Multiple management interfaces are supported on 8000 series devices and the ASA 5585-X with FirePOWER services only.

Syntax

```
configure network management-interface disable-management-channel ethn
```

where *n* is the number of the management interface you want to configure. **eth0** is the default management interface and **eth1** is the optional event interface. Cisco recommends that you leave the eth0 default management interface enabled, with both management and event channels enabled. See Management Interfaces, on page 881 for detailed information about using a separate event interface on the Firepower Management Center and on the managed device.

Example

```
> configure network management-interface disable-management-channel eth1
```

management-interface enable

Enables the specified management interface. Multiple management interfaces are supported on 8000 series devices and the ASA 5585-X with FirePOWER services only.

Syntax

```
configure network management-interface enable ethn
```

where *n* is the number of the management interface you want to enable. **eth0** is the default management interface and **eth1** is the optional event interface.

For device management, the Firepower Management Center management interface carries two separate traffic channels: the management traffic channel carries all internal traffic (such as inter-device traffic specific to the management of the device), and the event traffic channel carries all event traffic (such as web events). You can optionally configure a separate event-only interface on the Management Center to handle event traffic (see the Firepower Management Center web interface do perform this configuration). You can only configure one event-only interface. Event traffic can use a large amount of bandwidth, so separating event traffic from management traffic can improve the performance of the Management Center.

The default eth0 interface includes both management and event channels by default. You can optionally enable the eth0 interface as an event-only interface. Event traffic is sent between the device event interface and the Firepower Management Center event interface if possible. If the event network goes down, then event traffic reverts to the default management interface. Separate event interfaces are used when possible, but the management interface is always the backup.

When you enable a management interface, both management and event channels are enabled by default. We recommend that you use the default management interface for both management and eventing channels; and then enable a separate event-only interface. The Firepower Management Center event-only interface cannot accept management channel traffic, so you should simply disable the management channel on the device event interface.

Use the **configure network** {**ipv4** | **ipv6**} **manual** commands to configure the address(es) for management interfaces.

Example

```
> configure network management-interface enable eth1
> configure network management-interface disable-management-channel eth1
```

management-interface enable-event-channel

Enables the event traffic channel on the specified management interface. Multiple management interfaces are supported on 8000 series devices and the ASA 5585-X with FirePOWER services only.

Syntax

```
configure network management-interface enable-event-channel ethn
```

where *n* is the number of the management interface you want to configure. **eth0** is the default management interface and **eth1** is the optional event interface. Cisco recommends that you leave the eth0 default management interface enabled, with both management and event channels enabled. See Management Interfaces, on page 881 for detailed information about using a separate event interface on the Firepower Management Center and on the managed device.

Example

```
> configure network management-interface enable-event-channel eth1
```

management-interface enable-management-channel

Enables the management traffic channel on the specified management interface. Multiple management interfaces are supported on 8000 series devices and the ASA 5585-X with FirePOWER services only.

Syntax

```
configure network management-interface enable-management-channel ethn
```

where *n* is the number of the management interface you want to configure. **eth0** is the default management interface and **eth1** is the optional event interface. Cisco recommends that you leave the eth0 default management interface enabled, with both management and event channels enabled. See Management Interfaces, on page 881 for detailed information about using a separate event interface on the Firepower Management Center and on the managed device.

Example

```
> configure network management-interface enable-management-channel eth1
```

management-interface tcpport

Changes the value of the TCP port for management.

Syntax

```
configure network management-interface tcpport port
```

where *port* is the management port value you want to configure.

Example

```
> configure network management-interface tcpport 8500
```

management-port

Sets the value of the device's TCP management port.

Syntax

```
configure network management-port number
```

where *number* is the management port value you want to configure.

Example

```
> configure network management-port 8500
```

static-routes ipv4 add

Adds an IPv4 static route for the specified management interface.

Syntax

```
configure network static-routes ipv4
add interface destination netmask gateway
```

where interface is the management interface, destination is the destination IP address, netmask is the network mask address, and gateway is the gateway address you want to add.

Example

```
> configure network static-routes ipv4
add eth1 10.115.24.0 255.255.255.0 10.115.9.2
```

static-routes ipv4 delete

Deletes an IPv4 static route for the specified management interface.

Syntax

```
configure network static-routes ipv4
delete interface destination netmask gateway
```

where interface is the management interface, destination is the destination IP address, netmask is the network mask address, and gateway is the gateway address you want to delete.

Example

```
> configure network static-routes ipv4
delete eth1 10.115.24.0 255.255.255.0 10.115.9.2
```

static-routes ipv6 add

Adds an IPv6 static route for the specified management interface.

Syntax

```
configure network static-routes ipv6
add interface destination prefix gateway
```

where interface is the management interface, destination is the destination IP address, prefix is the IPv6 prefix length, and gateway is the gateway address you want to add.

Example

```
> configure network static-routes ipv6
add eth1 2001:DB8:3ffe:1900:4545:3:200: f8ff:fe21:67cf 64
```

static-routes ipv6 delete

Deletes an IPv6 static route for the specified management interface.

Syntax

```
configure network static-routes ipv6
delete interface destination prefix gateway
```

where interface is the management interface, destination is the destination IP address, prefix is the IPv6 prefix length, and gateway is the gateway address you want to delete.

Example

```
> configure network static-routes ipv6
delete eth1 2001:DB8:3ffe:1900:4545:3:200:f8ff: fe21:67cf 64
```

password

Allows the current user to change their password. After issuing the command, the CLI prompts the user for their current (or old) password, then prompts the user to enter the new password twice.

Access

Basic

Syntax

```
configure password
```

Example

```
> configure password
Enter current password:
Enter new password:
Confirm new password:
```

stacking disable

On 7000 and 8000 Series devices, removes any stacking configuration present on that device:

- On devices configured as primary, the stack is removed entirely.
- On devices configured as secondary, that device is removed from the stack.

This command is not available on NGIPSv or ASA FirePOWER modules, and you cannot use it to break a device high-availability pair.

Use this command when you cannot establish communication with appliances higher in the stacking hierarchy. If the Firepower Management Center is available for communication, a message appears instructing you to use the Firepower Management Center web interface instead; likewise, if you enter `stacking disable` on a device configured as secondary when the primary device is available, a message appears instructing you to enter the command from the primary device.

Access

Configuration

Syntax

```
configure stacking disable
```

Example

```
> configure stacking disable
```

user Commands

Applicable only to NGIPSv, the `configure user` commands manage the device's local user database.

Access

Configuration

access

Modifies the access level of the specified user. This command takes effect the next time the specified user logs in.

Syntax

```
configure user access username [basic | config]
```

where *username* specifies the name of the user for which you want to modify access, `basic` indicates basic access, and `config` indicates configuration access.

Example

```
> configure user access jdoe basic
```

add

Creates a new user with the specified name and access level. This command prompts for the user's password.

Syntax

```
configure user add username [basic | config]
```

where username specifies the name of the new user, `basic` indicates basic access, and `config` indicates configuration access.

Example

```
> configure user add jdoe basic
Enter new password for user jdoe:
Confirm new password for user jdoe:
```

aging

Forces the expiration of the user's password.

Syntax

```
configure user aging username max_days warn_days
```

where `username` specifies the name of the user, `max_days` indicates the maximum number of days that the password is valid, and`warn_days` indicates the number of days that the user is given to change the password before it expires.

Example

```
> configure user aging jdoe 100 3
```

delete

Deletes the user and the user's home directory.

Syntax

```
configure user delete username
```

where `username` specifies the name of the user.

Example

```
> configure user delete jdoe
```

disable

Disables the user. Disabled users cannot login.

Syntax

```
configure user disable username
```

where `username` specifies the name of the user.

Example

```
> configure user disable jdoe
```

enable

Enables the user.

Syntax

```
configure user enable username
```

where username specifies the name of the user.

Example

```
> configure user enable jdoe
```

forcereset

Forces the user to change their password the next time they login. When the user logs in and changes the password, strength checking is automatically enabled.

Syntax

```
configure user forcereset username
```

where *username* specifies the name of the user.

Example

```
> configure user forcereset jdoe
```

maxfailedlogins

Sets the maximum number of failed logins for the specified user.

Syntax

```
configure user maxfailedlogins username number
```

where *username* specifies the name of the user, and *number* specifies the maximum number of failed logins.

Example

```
> configure user maxfailedlogins jdoe 3
```

minpasswdlen

Sets the minimum number of characters a user password must contain.

Syntax

```
configure user minpasswdlen username number
```

Where *username* specifies the name of the user account, and *number* specifies the minimum number of characters the password for that account must contain (ranging from 1 to 127).

Example

```
> configure user minpasswdlen jdoe 13
```

password

Sets the user's password. This command prompts for the user's password.

Syntax

```
configure user password username
```

where *username* specifies the name of the user.

Example

```
> configure user pasword jdoe
Enter new password for user jdoe:
Confirm new password for user jdoe:
```

strengthcheck

Enables or disables the strength requirement for a user's password. When a user's password expires or if the configure user forcereset command is used, this requirement is automatically enabled the next time the user logs in.

Syntax

```
configure user strengthcheck username {enable | disable}
```

where *username* specifies the name of the user, `enable` sets the requirement for the specified users password, and `disable` removes the requirement for the specified user's password.

Example

```
> configure user strengthcheck jdoe enable
```

unlock

Unlocks a user that has exceeded the maximum number of failed logins.

Syntax

```
configure user unlock username
```

where *username* specifies the name of the user.

Example

```
> configure user unlock jdoe
```

vmware-tools

Enables or disables VMware Tools functionality on NGIPSv. This command is available only on NGIPSv.

VMware Tools is a suite of utilities intended to enhance the performance of the virtual machine. These utilities allow you to make full use of the convenient features of VMware products. The system supports the following plugins on all virtual appliances:

- guestInfo
- powerOps
- timeSync
- vmbackup

For more information about VMware Tools and the supported plugins, see the VMware website (http://www.vmware.com).

Access

Basic

Syntax

```
configure vmware-tools [enable | disable]
```

Example

```
> configure vmware-tools enable
```

Classic Device CLI System Commands

The system commands enable the user to manage system-wide files and access control settings. Only users with configuration CLI access can issue commands in system mode.

access-control Commands

The `system access-control` commands enable the user to manage the access control configuration on the device.

Access

Configuration

archive

Saves the currently deployed access control policy as a text file on `/var/common`.

Syntax

```
system access-control archive
```

Example

```
> system access-control archive
```

clear-rule-counts

Resets the access control rule hit count to 0.

Syntax

```
system access-control clear-rule-counts
```

Example

```
> system access-control clear-rule-counts
```

rollback

Reverts the system to the previously deployed access control configuration. You cannot use this command with devices in stacks or high-availability pairs.

Syntax

```
system access-control rollback
```

Example

```
> system access-control rollback
```

compliance Commands

The `compliance` commands display and configure the device's security certifications compliance mode.

Caution After you enable this setting, you cannot disable it. If you need to do so, contact Support for assistance.

Access

Configuration

enable cc

Configures the device's security certifications compliance to Common Criteria (CC) mode.

Caution After you enable this setting, you cannot disable it. If you need to do so, contact Support for assistance.

Syntax

```
system compliance enable cc
```

Example

```
> system compliance enable cc
```

enable ucapl

Configures the device's security certifications compliance to Unified Capabilities Approved Products List (UCAPL) mode.

Caution After you enable this setting, you cannot disable it. If you need to do so, contact Support for assistance.

Syntax

```
system compliance enable ucapl
```

Example

```
> system compliance enable ucapl
```

show

Displays the device's current security certifications compliance mode.

Syntax

```
system compliance show
```

Example

```
> system compliance show
```

disable-http-user-cert

Disables the requirement that the browser present a valid client certificate.

Access

Configuration

Syntax

```
system disable-http-user-cert
```

Example

```
> system disable-http-user-cert
```

file Commands

The `system file` commands enable the user to manage the files in the common directory on the device.

Access

Configuration

copy

Uses FTP to transfer files to a remote location on the host using the login username. The local files must be located in the common directory.

Syntax

```
system file copy hostname username path filenames filenames ...
```

where `hostname` specifies the name or ip address of the target remote host, `username` specifies the name of the user on the remote host, `path` specifies the destination path on the remote host, and `filenames` specifies the local files to transfer; the file names are space-separated.

Example

```
> system file copy sfrocks jdoe /pub *
```

delete

Removes the specified files from the common directory.

Syntax

```
system file delete filenames filenames ...
```

where `filenames` specifies the files to delete; the file names are space-separated.

Example

```
> system file delete *
```

list

If no file names are specified, displays the modification time, size, and file name for all the files in the common directory. If file names are specified, displays the modification time, size, and file name for files that match the specified file names.

Syntax

```
system file list filenames
```

where *filenames* specifies the files to display; the file names are space-separated.

Example

```
> system file list
```

secure-copy

Uses SCP to transfer files to a remote location on the host using the login username. The local files must be located in the `/var/common` directory.

Syntax

```
system file secure-copy hostname username path filenames filenames ...
```

where *hostname* specifies the name or ip address of the target remote host, *username* specifies the name of the user on the remote host, *path* specifies the destination path on the remote host, and *filenames* specifies the local files to transfer; the file names are space-separated.

Example

```
> system file secure-copy 10.123.31.1 jdoe /tmp *
```

generate-troubleshoot

Generates troubleshooting data for analysis by Cisco.

Caution Generating troubleshooting files for lower-memory devices can trigger Automatic Application Bypass (AAB) when AAB is enabled, At a minimum, triggering AAB restarts the Snort process, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 350 for more information. In some such cases, triggering AAB can render the device temporarily inoperable. If inoperability persists, contact Cisco Technical Assistance Center (TAC), who can propose a solution appropriate to your deployment. Susceptible devices include Firepower 7010, 7020, and 7030; ASA 5506-X, 5508-X, 5516-X, 5512-X , 5515-X, and 5525-X; NGIPSv.

Access

Configuration

Syntax

```
system generate-troubleshoot option1 optionN
```

Where options are one or more of the following, space-separated:

- `ALL`: Run all of the following options.
- `SNT`: Snort Performance and Configuration
- `PER`: Hardware Performance and Logs
- `SYS`: System Configuration, Policy, and Logs
- `DES`: Detection Configuration, Policy, and Logs
- `NET`: Interface and Network Related Data
- `VDB`: Discover, Awareness, VDB Data, and Logs
- `UPG`: Upgrade Data and Logs
- `DBO`: All Database Data
- `LOG`: All Log Data
- `NMP`: Network Map Information

Example

```
> system generate-troubleshoot VDB NMP
starting /usr/local/sf/bin/sf_troubleshoot.pl…
Please, be patient. This may take several minutes.
The troubleshoot options codes specified are VDB,NMP.
Getting filenames from [usr/local/sf/etc/db_updates/index]
Getting filenames from [usr/local/sf/etc/db_updates/base-6.2.3]
Troubleshooting information successfully created at
/var/common/results-06-14-2018—222027.tar.gz
```

ldapsearch

Enables the user to perform a query of the specified LDAP server. Note that all parameters are required.

Access

Configuration

Syntax

```
system ldapsearch host port baseDN userDN basefilter
```

where host specifies the LDAP server domain, port specifies the LDAP server port, baseDN specifies the DN (distinguished name) that you want to search under, userDN specifies the DN of the user who binds to the LDAP directory, and basefilter specifies the record or records you want to search for.

Example

```
> system ldapsearch ldap.example.com 389 cn=users,

dc=example,dc=com cn=user1,cn=users,dc=example,dc=com, cn=user2
```

lockdown-sensor

Removes the `expert` command and access to the Linux shell on the device.



Caution This command is irreversible without a hotfix from Support. Use with care.

Access

Configuration

Syntax

```
system lockdown-sensor
```

Example

```
> system
```

nat rollback

Reverts the system to the previously applied NAT configuration. This command is not available on NGIPSv or ASA FirePOWER. You cannot use this command with devices in stacks or high-availability pairs.

Access

Configuration

Syntax

```
system nat rollback
```

Example

```
> system nat rollback
```

reboot

Reboots the device.

Access

Configuration

Syntax

```
system reboot
```

Example

```
> system reboot
```

restart

Restarts the device application.

Access

Configuration

Syntax

```
system restart
```

Example

```
> system restart
```

support Commands

The `system support` commands enable the user to manage special SSL ClientHello processing on the device.

Access

Configuration

ssl-client-hello-display

Displays the current settings for processing the ClientHello message during an SSL handshake. For a description of these settings, see the `ssl-client-hello-enabled` and `ssl-client-hello-tuning` commands.

Access

Basic

Syntax

```
system support ssl-client-hello-display
```

Example

```
> system support ssl-client-hello-display
```

ssl-client-hello-enabled

Controls special processing of the ClientHello message during the SSL handshake.



Caution Use these commands *only* if advised to do so by Cisco TAC.

Access

Configuration

Syntax

```
system support ssl-client-hello-enabled setting {true | false}
```

Possible *setting* values are:

feature
Controls all special handling of ClientHello messages.

curves
Controls stripping of elliptic curves that the Firepower System does not support:

- `true` (enabled)—The system strips any unsupported elliptic curves from the ClientHello message, increasing the likelihood of traffic decryption. You must also enable the `extensions` setting.
- `false` (disabled)—The system retains unsupported elliptic curves in the ClientHello message, decreasing the likelihood of traffic decryption.

ciphers
Controls stripping of cipher suites that the Firepower System does not support:

- `true` (enabled)—The system strips unsupported cipher suites from ClientHello messages, increasing the likelihood of traffic decryption.
- `false` (disabled)—The system retains unsupported cipher suites in ClientHello messages. This decreases the likelihood of traffic decryption and can result in a number of `Unsupported` or `Unknown Cipher` errors in the SSL Flow Error field of associated connection events.

extensions
Controls stripping of TLS extensions that prevent decryption:

- `true` (enabled)—The system identifies TLS extensions that prevent decryption and strips them from the ClientHello message. This value is required if you want to enable **curves**, **session_ticket**, and **alpn**.
- `false` (disabled)—The system retains all TLS extensions in the ClientHello message. This decreases the likelihood of traffic decryptions and can result in `Unknown Session` errors in the SSL Flow Error field of associated connection events.

session_ticket
Controls processing of the SessionTicket extension in ClientHello messages. If the system can match a SessionTicket value in an incoming ClientHello message to cached session data, it can resume the session without the client and server performing the full SSL handshake.

- `true` (enabled)—The system strips unrecognized SessionTicket values from the ClientHello message. This increases the likelihood of traffic decryption for the resumed session. You must also enable the `extensions` setting.
- `false` (disabled)—The system retains all SessionTicket values in the ClientHello message. This decreases the likelihood of traffic decryption and can result in `Uncached Session` errors in the SSL Flow Error field of associated connection events.

session_id
Controls processing of the Session Identifier element in ClientHello messages. If the system can match the Session Identifier in an incoming ClientHello message to cached session data, it can resume the session without the client and server performing the full SSL handshake.

- `true` (enabled)—The system strips unrecognized Session Identifier values from the ClientHello message. This increases the likelihood of traffic decryption for the resumed session.
- `false` (disabled)—The system retains all Session Identifier values in the ClientHello message. This decreases the likelihood of traffic decryption and can result in `Uncached Session` errors in the SSL Flow Error field of associated connection events.

alpn
Controls stripping of ALPN protocol values that cannot be decrypted, specifically, the SPDY and HTTP2 protocols:

- `true` (enabled)—The system prevents the client from establishing SPDY or HTTP2 sessions, increasing the likelihood of traffic decryption and inspection. You must also enable the `extensions` setting.
- `false` (disabled)—The system allows the client to establish SPDY or HTTP2 sessions with the server, decreasing the likelihood of traffic decryption and inspection.

compression
Controls stripping of TLS compression requests from ClientHello messages:

- `true` (enabled)—The system prevents the client from establishing a TLS compressed session with the server.
- `false` (disabled)—The system allows the client to establish a TLS compressed session with the server. This prevents traffic decryption for the session and can result in `Compression Used` errors in the SSL Flow Error field of associated connection events.

Example

```
> system support ssl-client-hello-enabled feature false
```

ssl-client-hello-force-reset

Resets the configurable settings for ClientHello message processing to default values. The system does not require user confirmation before proceeding.



Caution Do **not** use this command unless you are directed to do so by Support.

Access

Configuration

Syntax

```
system support ssl-client-hello-force-reset
```

Example

```
> system support ssl-client-hello-force-reset
```

ssl-client-hello-reset

Resets the configurable settings for ClientHello message processing to default values. The system requires user confirmation before proceeding.

Caution Do **not** use this command unless you are directed to do so by Support.

Access

Configuration

Syntax

```
system support ssl-client-hello-reset
```

Example

```
> system support ssl-client-hello-reset
```

ssl-client-hello-tuning

Allows you to refine how the managed device modifies ClientHello messages during SSL handshakes. This command tunes the default lists of cipher suites, elliptic curves, and extensions that the system allows in ClientHello messages. This command only adds entries to or removes entries from the default lists of allowed values. It does not overwrite the default lists.

Caution Do **not** use this command unless you are directed to do so by Support.

Access

Configuration

Syntax

```
system support ssl-client-hello-tuning setting value
```

The *value* element supports a comma-delimited list of values. Possible values for the *setting* and *value* elements include:

Setting	System Action	*Value*
`ciphers_allow`	Allows the specified cipher suites in ClientHello messages. If you use this command, the system retains the specified cipher suites in any ClientHello messages it modifies.	Obtain individual cipher suite numbers from the IANA website: https://www.iana.org/assignments/tls-parameters/tls-parameters.xhtml#tls-parameters-4 IANA provides values in hexadecimal. Convert them to decimal for use in this command.
`ciphers_remove`	Disallows the specified cipher suites in ClientHello messages. If you use this command, the system strips the specified cipher suites from any ClientHello message it modifies.	

Setting	System Action	*Value*
`curves_allow`	Allows the specified elliptic curves in ClientHello messages. If you use this command, the system retains the specified elliptic curves in any ClientHello message it modifies.	Obtain curve numbers from the IANA website: https://www.iana.org/assignments/tls-parameters/tls-parameters.xhtml#tls-parameters-8
`curves_remove`	Disallows the specified elliptic curves in ClientHello messages. If you use this command, the system strips the specified elliptic curves from any ClientHello message it modifies.	
`extensions_allow`	Allows the specified extensions in ClientHello messages. If you use this command, the system retains the specified extensions in any ClientHello message it modifies.	Obtain extension numbers from the IANA website: https://www.iana.org/assignments/tls-extensiontype-values/tls-extensiontype-values.xhtml
`extensions_remove`	Disallows the specified elliptic curves in ClientHello messages. The system strips the specified extensions from any ClientHello message it modifies. By default, the system disallows extensions 22, 23, and 30032.	

Example

```
> system support ssl-client-hello-tuning ciphers_allow 4,7,16,22
```

shutdown

Shuts down the device. This command is not available on ASA FirePOWER modules.

Access

Configuration

Syntax

```
system shutdown
```

Example

```
> system shutdown
```