

Cisco Firepower Threat Defense Configuration Guide for Firepower Device Manager, Version 7.1

First Published: 2021-12-01

Last Modified: 2026-07-16

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)

© 2021–2026 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1 **Getting Started** 1

Is This Guide for You? 1

New Features in Firewall Device Manager/Firewall Threat Defense Version 7.1.0 1

Default Configuration 7

Default Configuration Prior to Initial Setup 8

Configuration After Initial Setup 10

Logging Into the System 14

Your User Role Controls What You Can See and Do 14

Logging Into the Firewall Device Manager 14

Logging Into the Command Line Interface (CLI) 15

Changing Your Password 16

Setting User Profile Preferences 17

Viewing Pages in Languages Other Than English 18

Setting Up the System 18

Connect the Interfaces 19

(Optional) Change Management Network Settings at the CLI 19

Complete the Initial Configuration Using the Setup Wizard 21

What to Do if You Do Not Obtain an IP Address for the Outside Interface 23

Configuration Basics 24

Configuring the Device 24

Configuring Security Policies 25

Searching for Rules or Objects 26

Deploying Your Changes 27

Configuration Changes that Restart Inspection Engines 28

Configuration Changes that Force a Full Deployment 29

Viewing Interface and Management Status 30

Viewing System Task Status 31
Using the CLI Console to Monitor and Test the Configuration 31
Using Firewall Device Manager and the REST API Together 33
Communication Port and Internet Access Requirements 33
Communication Ports Used by the Device 33
Internet Resources Accessed by the Device 35

CHAPTER 2 **Best Practices: Use Cases for Firewall Threat Defense** 39
How to Configure the Device in Firewall Device Manager 39
How to Gain Insight Into Your Network Traffic 44
How to Block Threats 51
How to Block Malware 56
How to Implement an Acceptable Use Policy (URL Filtering) 59
How to Control Application Usage 64
How to Add a Subnet 67
How to Passively Monitor the Traffic on a Network 72
More Examples 77

CHAPTER 3 **Licensing the System** 81
Smart Licensing for the Firewall System 81
Cisco Smart Software Manager 81
Periodic Communication with the License Authority 82
Smart License Types 82
Firewall Threat Defense Virtual Licensing 83
Firewall Threat Defense Virtual Performance Tier Licensing Guidelines and Limitations 84
Impact of Export Control Setting on Encryption Features 85
Impact of Expired or Disabled Optional Licenses 85
Managing Smart Licenses 86
Registering the Device 87
Changing the Firewall Threat Defense Virtual Performance Tier 88
Enabling or Disabling Optional Licenses 89
Synchronizing with the Cisco Smart Software Manager 90
Unregistering the Device 90
Applying Permanent Licenses 91

Universal Permanent vs. Specific License Reservation 91
Verify That Your Smart Account Can Provide a Universal License 92
Switch to PLR Mode and Apply a Universal License 92
Cancel PLR Registration 93
Unregister the Device in PLR Mode 94

PART I **System Monitoring** 97

CHAPTER 4 **Monitoring the Device** 99
Enable Logging to Obtain Traffic Statistics 99
Event Types 99
Configurable Connection Logging 100
Automatic Connection Logging 101
Tips for Connection Logging 101
Sending Events to an External Syslog Server 102
Evaluating Events Using Cisco Cloud-Based Services 102
Monitoring Traffic and System Dashboards 102
Monitoring Additional Statistics Using the Command Line 105
Viewing Events 105
Configuring Custom Views 106
Filtering Events 107
Event Field Descriptions 108

CHAPTER 5 **Alarms for the Cisco ISA 3000** 119
About alarms 119
Alarm input interfaces 120
Alarm output interface 120
Syslog Alarms 121
SNMP Trap Alarms 121
Defaults for alarms 121
Configuring Alarms for the ISA 3000 122
Configure Alarm Input Contacts 122
Configure Power Supply Alarms 124
Configure Temperature Alarms 126

Monitoring alarms 127
Monitoring alarm status 127
Monitoring syslog messages for alarms 128
Turn off the external alarm 128

PART II **Reusable Objects** 129

CHAPTER 6 **Objects** 131
Object Types 131
Guidelines and Limitations for Objects 134
Managing Objects 134
Configuring Network Objects and Groups 135
Configuring Port Objects and Groups 136
Configuring Security Zones 138
Configuring Application Filter Objects 139
Configuring URL Objects and Groups 141
Configuring Geolocation Objects 142
Configuring Syslog Servers 143
Configuring Security Group Tag (SGT) Groups 145

CHAPTER 7 **Certificates** 147
About Certificates 147
Public Key Cryptography 148
Certificate Types Used by Feature 148
Example: Generating an Internal Certificate using OpenSSL 149
Configuring Certificates 150
Uploading Internal and Internal CA Certificates 151
Generating Self-Signed Internal and Internal CA Certificates 153
Uploading Trusted CA Certificates 154
Configuring Trusted CA Certificate Groups 156

CHAPTER 8 **Identity Sources** 157
About Identity Sources 157
Active Directory (AD) Identity Realms 159

Supported Directory Servers 159
Limitations on Number of Users 159
Determining the Directory Base DN 160
Configuring AD Identity Realms 161
Configuring an AD Realm Sequence 163
Troubleshooting Directory Server Connections 164
RADIUS Servers and Groups 165
Configure RADIUS Servers 165
Configure RADIUS Server Groups 167
Troubleshoot RADIUS Servers and Groups 168
Identity Services Engine (ISE) 169
Guidelines and Limitations for ISE 169
Configure Identity Services Engine 170
Troubleshoot the ISE/ISE-PIC Identity Source 172
SAML Servers 172
Configure SAML Servers 172
Local Users 175
Configure Local Users 175

PART III **The Basics** 177

CHAPTER 9 **Logical Devices on the Firepower 4100/9300** 179
About Interfaces 179
Chassis Management Interface 179
Interface Types 180
FXOS Interfaces vs. Application Interfaces 181
Requirements and Prerequisites for Firepower 9300 Hardware and Software Combinations 181
Guidelines and Limitations for Logical Devices 182
Guidelines and Limitations for Interfaces 182
General Guidelines and Limitations 182
Configure Interfaces 182
Enable or Disable an Interface 182
Configure a Physical Interface 183
Add an EtherChannel (Port Channel) 183

Configure a Logical Device 184
Add a Standalone Firewall Threat Defense for the Firewall Device Manager 184
Add a High Availability Pair 185
Change an Interface on a Firewall Threat Defense Logical Device 186
Connect to the Console of the Application 189
History for Logical Devices 190

CHAPTER 10 High Availability (Failover) 191
About High Availability (Failover) 191
Active/standby failover 191
Primary/secondary roles and active/standby status 192
Active unit determination at startup 192
Failover events 192
Failover and Stateful Failover Links 193
Failover Link 194
Stateful Failover Link 194
Interfaces for the Failover and State Links 194
Connecting the Failover and Stateful Failover Interfaces 195
Avoiding Interrupted Failover and Data Links 195
How Stateful Failover Affects User Connections 196
Supported features 197
Unsupported features 198
Configuration Changes and Actions Allowed on a Standby Unit 199
System Requirements for High Availability 199
Hardware Requirements for HA 199
Software Requirements for HA 200
License Requirements for HA 200
Guidelines for High Availability 201
Configuring High Availability 203
Prepare the Two Units for High Availability 203
Configure the Primary Unit for High Availability 205
Configure the Secondary Unit for High Availability 208
Configure Failover Criteria for Health Monitoring 209
Configure Peer Unit Health Monitoring Failover Criteria 209

Configure Interface Health Monitoring Failover Criteria 210
How the System Tests Interface Health 212
Configure Standby IP and MAC Addresses 213
Verify the High Availability Configuration 214
Managing High Availability 215
Suspending or Resuming High Availability 216
Breaking High Availability 217
Switching the Active and Standby Peers (Forcing Failover) 219
Preserving Undeployed Configuration Changes After Failover 219
Changing Licenses and Registration in High Availability Mode 220
Editing the HA IPsec Encryption Key or HA Configuration 220
Marking a Failed Unit Healthy 220
Upgrading High Availability Firewall Threat Defense 221
Troubleshooting High Availability Threat Defense Upgrades 222
Replacing a Unit in a High Availability Pair 225
Monitoring High Availability 225
Monitoring General Failover Status and History 225
Monitoring Status for HA-Monitored Interfaces 227
Monitoring HA-Related Syslog Messages 227
Remotely Executing CLI Commands on the Peer Unit 228
Troubleshooting High Availability (Failover) 228
Troubleshooting Failed State for a Unit 230
Troubleshooting HA App Sync Failures 231

CHAPTER 11 **Interfaces** 235
About Firewall Threat Defense Interfaces 235
Interface Modes 236
Management Interface 236
Recommendations for Configuring a Separate Management Network 237
Security Zones 237
IPv6 Addressing 238
Auto-MDI/MDIX Feature 238
Guidelines and Limitations for Interfaces 238
Limitations for Interface Configuration 238

Maximum number of VLAN subinterfaces by device model 239
Configure a Physical Interface 239
Configure Bridge Groups 244
Configure EtherChannels 248
About EtherChannels 249
Channel Group Interfaces 249
Connecting to an EtherChannel on Another Device 249
Link Aggregation Control Protocol 251
Load Balancing 251
EtherChannel MAC Address 251
Guidelines for EtherChannels 252
Add an EtherChannel 254
Configure VLAN Interfaces and Switch Ports 259
Understanding switch ports and interfaces 259
Prerequisites for switch ports 260
Guidelines for switch ports 260
Configure a VLAN Interface 261
Configure Switch Ports as Access Ports 265
Configure Switch Ports as Trunk Ports 267
Configure Power Over Ethernet 269
Configure VLAN Subinterfaces and 802.1Q Trunking 270
Configure Passive Interfaces 275
Why Use Passive Interfaces? 276
Limitations for Passive Interfaces 276
Configure the Switch for a Hardware Firewall Threat Defense Passive Interface 277
Configure the VLAN for a Firewall Threat Defense Virtual Passive Interface 278
Configure a Physical Interface in Passive Mode 278
Configure Advanced Interface Options 279
About MAC Addresses 279
About the MTU 280
Path MTU Discovery 280
MTU and Fragmentation 280
MTU and Jumbo Frames 280
Configure Advanced Options 281

Scan for Interface Changes, and Migrate an Interface 283
About Interface Scanning and Migrating 283
Guidelines and Limitations for Interface Scanning and Migrating 284
Scan and Migrate Interfaces 285
Manage the Network Module for the Secure Firewall 3100 288
Add a Network Module 288
Hot Swap the Network Module 290
Replace the Network Module with a Different Type 292
Remove the Network Module 294
Configure Hardware Bypass for Power Failure (ISA 3000) 296
Monitoring Interfaces 298
Examples for Interfaces 299

PART IV **Routing** 301

CHAPTER 12 **Routing Basics and Static Routes** 303
Best Practices for Routing 303
Routing Overview 303
Supported Routing Protocols 304
Route Types 305
The Routing Table and Route Selection 305
How the Routing Table Is Populated 305
How Forwarding Decisions Are Made 307
Routing Table for Management Traffic 308
Equal-Cost Multi-Path (ECMP) Routing 309
Static Routes 309
About Static and Default Routes 309
Default Route 310
Static Routes 310
Backup Static Routes and Static Route Tracking 310
Guidelines for Static Routing 311
Configuring Static Routes 312
Configure SLA Monitor Objects 313
Configuring ECMP Traffic Zones 315

Monitoring Routing 317

CHAPTER 13 **Virtual Routers** 319

Virtual routers and virtual routing and forwarding (VRF) 319

Virtual-router-aware policies 320

Routing Between Virtual Routers 320

Maximum Number of Virtual Routers By Device Model 321

Guidelines for Virtual Routers 322

Managing Virtual Routers 324

Create a Virtual Router or Edit Interface Assignments 325

Configure Static Routes and Routing Processes in a Virtual Router 326

Delete a Virtual Router 327

Examples for Virtual Routers 327

How to Route to a Distant Server through Multiple Virtual Routers 327

How to Provide Internet Access to Multiple Virtual Routers with Overlapping Address Spaces 332

Monitoring Virtual Routers 343

CHAPTER 14 **Route Maps and Other Objects for Route Tuning** 345

Configure Route Maps 345

Route Map Permit and Deny Clauses 345

Route Map Match and Set Statements 346

Configure a Route Map 346

Configure Access Lists 350

Configure Extended Access Lists 351

Configure Standard Access Lists 352

Configure AS Path Access Lists 354

Configure Community Lists 355

Configure Policy Lists 357

Configure Prefix Lists 358

CHAPTER 15 **Open Shortest Path First (OSPF)** 363

Configure the OSPFv2 Process and Areas 363

Customizing OSPF Process and Area Characteristics 365

Configure Advanced Settings for an OSPF Process 365

Configure OSPF Area Properties 369
Configure Static OSPF Neighbors 373
Configure OSPF Summary Addresses 373
Configure OSPF Filter Rules 374
Configure OSPF Redistribution 376
Configure OSPFv2 Interface Settings and OSPF Authentication 377
Configure OSPFv2 Lost Neighbor Detection and Fast Hello Packets (OSPF Interface Settings) 380
Monitoring OSPF 381

CHAPTER 16 **Enhanced Interior Gateway Routing Protocol (EIGRP)** 383
Best Practices for EIGRP 383
About EIGRP 384
DUAL Finite State Machine 384
EIGRP Metric Weights 384
EIGRP Cost Metrics 385
Guidelines for EIGRP 385
Configure the Core EIGRP Process 386
Configure the EIGRP Process for Full Routing 386
Configure the EIGRP Process for Stub Routing 387
Customize the EIGRP Process 389
Configure EIGRP Advanced Settings 389
Configure the Networks for EIGRP to Advertise 391
Configure EIGRP Passive Routing Interfaces 392
Configure Static EIGRP Neighbors 393
Control EIGRP Candidate Default Route Dissemination 394
Configure EIGRP Filter Rules 395
Configure EIGRP Route Redistribution 397
Monitoring EIGRP 399

CHAPTER 17 **Border Gateway Protocol (BGP)** 401
BGP 401
Routing table changes 401
BGP usage scenarios 402
BGP path selection 403

BGP multipath 403
Configure BGP 405
Configure BGP Global Settings 405
Configure the BGP Process 408
Configure BGP General Settings 409
Configure BGP Advanced Settings 410
Configure the Networks for BGP to Advertise 412
Configure BGP Route Injection 413
Configure BGP Aggregate Address Settings 414
Configure BGP Filter Settings for IPv4 416
Configure BGP Neighbors 417
Configure BGP Route Redistribution from Other Routing Protocols 423
Monitoring BGP 425

PART V **Security Policies** 427

CHAPTER 18 **SSL Decryption** 429
About SSL Decryption 429
Why Implement SSL Decryption? 429
Actions You Can Apply to Encrypted Traffic 430
Decrypt Re-Sign 430
Decrypt Known Key 431
Do Not Decrypt 431
Block 432
Automatically Generated SSL Decryption Rules 432
Handling Undecryptable Traffic 432
License Requirements for SSL Decryption 432
Guidelines for SSL Decryption 433
How to Implement and Maintain the SSL Decryption Policy 433
Configuring SSL Decryption Policies 435
Enable the SSL Decryption Policy 436
Configure the Default SSL Decryption Action 437
Configure SSL Decryption Rules 438
Source/Destination Criteria for SSL Decryption Rules 440

Application Criteria for SSL Decryption Rules 441
URL Criteria for SSL Decryption Rules 442
User Criteria for SSL Decryption Rules 443
Advanced Criteria for SSL Decryption Rules 444
Configure Certificates for Known Key and Re-Sign Decryption 445
Downloading the CA Certificate for Decrypt Re-Sign Rules 446
Example: Blocking Older SSL/TLS Versions from the Network 447
Monitoring and Troubleshooting SSL Decryption 449
Monitoring SSL Decryption 449
Handling Web Sites Where Decrypt Re-sign Works for a Browser but not an App (SSL or Certificate Authority Pinning) 450

CHAPTER 19 **Identity Policies** 453
Identity Policy Overview 453
Establishing User Identity Through Passive Authentication 454
Establishing User Identity through Active Authentication 454
Dealing with Unknown Users 454
How to Implement the Identity Policy 455
Best Practice for Active Authentication 456
Configuring Identity Policies 457
Configure Identity Policy Settings 457
Configure the Identity Policy Default Action 459
Configure Identity Rules 459
Enabling Transparent User Authentication 463
Requirements for Transparent Authentication 464
Configuring Internet Explorer for Transparent Authentication 464
Configuring Firefox for Transparent Authentication 465
Monitoring Identity Policies 466
Examples for Identity Policies 466

CHAPTER 20 **Security Intelligence** 467
About Security Intelligence 467
Making Exceptions to the Block Lists 468
Security Intelligence Feed Categories 468

License Requirements for Security Intelligence 469
Configuring Security Intelligence 469
Monitoring Security Intelligence 470
Examples for Security Intelligence 471

CHAPTER 21 **Access Control** 473
Best Practices for Access Control 473
Access Control Overview 476
Access Control Rules and the Default Action 476
Application Filtering 476
Application Control for Encrypted and Decrypted Traffic 477
Filtering on Common Industrial Protocol (CIP) and Modbus Applications (ISA 3000) 477
Best Practices for Application Filtering 477
URL Filtering 478
Filtering URLs by Category and Reputation 478
Looking Up the Category and Reputation for a URL 479
Manual URL Filtering 479
Filtering HTTPS Traffic 480
Comparing URL and Application Filtering 481
Best Practices for Effective URL Filtering 482
What the User Sees When You Block Web Sites 482
DNS Request Filtering 483
Guidelines for DNS Request Filtering 483
Filtering DNS Requests Based on URL Category and Reputation 484
Intrusion, File, and Malware Inspection 485
Best Practices for Access Control Rule Order 485
NAT and Access Rules 486
How Other Security Policies Impact Access Control 486
License Requirements for Access Control 486
Guidelines and Limitations for Access Control Policies 487
Configuring the Access Control Policy 489
Configuring the Default Action 489
Configuring Access Control Policy Settings 490
Configuring Access Control Rules 491

Source/Destination Criteria 492
Application Criteria 494
URL Criteria 496
User Criteria 497
Intrusion Policy Settings 498
File Policy Settings 498
Logging Settings 499
Monitoring Access Control Policies 500
Monitoring Access Control Statistics in the Dashboards 500
Examining Rule Hit Counts 501
Monitoring Syslog Messages for Access Control 502
Monitoring Access Control Policies in the CLI 502
Examples for Access Control 502
How to Control Network Access Using TrustSec Security Group Tags 503
About Security Group Tags (SGT) 503
Configure security groups and SXP publishing in ISE 504
Configure Access Control Based on Security Group Tag (SGT) 506

CHAPTER 22 **Intrusion Policies** 509
About Intrusion and Network Analysis Policies 509
System-Defined Network Analysis and Intrusion Policies 510
Inspection Mode: Prevention vs. Detection 510
Intrusion and Preprocessor Rules 511
Intrusion Rule Attributes 511
Default Intrusion Variable Set 512
Generator Identifiers 513
Network Analysis Policies 514
License Requirements for Intrusion Policies 515
Applying Intrusion Policies in Access Control Rules 515
Switching Between Snort 2 and Snort 3 516
Configuring Syslog for Intrusion Events 517
Configuring the Network Analysis Policy (Snort 3) 517
Configuring Inspector and Binder Overrides 519
Downloading Overrides and the Schema 521

Uploading Overrides 521
Managing Intrusion Policies (Snort 3) 522
Configuring a Custom Intrusion Policy (Snort 3) 523
Viewing or Editing Intrusion Policy Properties (Snort 3) 524
Adding or Removing Rule Groups in an Intrusion Policy (Snort 3) 526
Changing Intrusion Rule Actions (Snort 3) 528
Managing Custom Intrusion Rules and Rule Groups 530
Uploading Custom Intrusion Rules 531
Configuring Individual Custom Intrusion Rules 533
Managing Intrusion Policies (Snort 2) 534
Configure the Inspection Mode for an Intrusion Policy (Snort 2) 535
Changing Intrusion Rule Actions (Snort 2) 535
Monitoring Intrusion Policies 536
Examples for Intrusion Policies 537

CHAPTER 23 **Network Address Translation (NAT)** 539
Why use NAT? 539
NAT basics 540
NAT terminology 540
NAT types 540
NAT in routed mode 541
Auto NAT and Manual NAT 541
Auto NAT 542
Manual NAT 542
Comparing Auto NAT and Manual NAT 542
NAT rule order 543
NAT interfaces 545
Configuring Routing for NAT 545
Addresses on the same network as the mapped interface 545
Addresses on a unique network 546
The same address as the real address (identity NAT) 546
Guidelines for NAT 546
Interface Guidelines 546
IPv6 NAT guidelines 547

IPv6 NAT best practices 547
NAT Support for Inspected Protocols 547
FQDN destination guidelines 549
Additional guidelines for NAT 549
Configure NAT 552
Dynamic NAT 552
About dynamic NAT 552
Dynamic NAT disadvantages and advantages 553
Configure Dynamic Auto NAT 554
Configure Dynamic Manual NAT 555
Dynamic PAT 557
About dynamic PAT 557
Dynamic PAT disadvantages and advantages 558
Configure Dynamic Auto PAT 558
Configure Dynamic Manual PAT 559
Static NAT 562
About static NAT 562
Configure Static Auto NAT 566
Configure Static Manual NAT 567
Identity NAT 570
Configure Identity Auto NAT 570
Configure Identity Manual NAT 572
NAT Rule Properties 574
Packet Translation Properties for Auto NAT 574
Packet Translation Properties for Manual NAT 576
Advanced NAT Properties 577
Translating IPv6 networks 578
NAT64/46: translating IPv6 addresses to IPv4 578
NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet 579
NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet and DNS Translation 581
NAT66: translating IPv6 addresses to different IPv6 addresses 586
NAT66 Example, Static Translation between Networks 586
NAT66 Example, Simple IPv6 Interface PAT 589
Monitoring NAT 592

Examples for NAT 593
Providing Access to an Inside Web Server (Static Auto NAT) 593
Single Address for FTP, HTTP, and SMTP (Static Auto NAT-with-Port-Translation) 596
Different Translation Depending on the Destination (Dynamic Manual PAT) 602
Different Translation Depending on the Destination Address and Port (Dynamic Manual PAT) 608
Rewriting DNS queries and responses using NAT 613
DNS 64 Reply Modification 614
DNS Reply Modification, DNS Server on Outside 620
DNS Reply Modification, DNS Server on Host Network 623

PART VI **Virtual Private Networks (VPN)** 627

CHAPTER 24 **Site-to-Site VPN** 629
VPN Basics 629
Internet Key Exchange (IKE) 630
VPN encryption and performance 630
Decide encryption algorithms for VPN policies 631
Decide which hash algorithms to use 632
Decide which Diffie-Hellman modulus group to use 632
Deciding Which Authentication Method to Use 633
VPN Topologies 634
Establishing Site-to-Site VPN Connections with Dynamically-Addressed Peers 634
Virtual Tunnel Interfaces and Route-Based VPN 634
Overview Process for Configuring Route-Based VPNs 635
Guidelines for Virtual Tunnel Interfaces and Route-Based VPN 635
Managing Site-to-Site VPNs 636
Configuring a Site-to-Site VPN Connection 637
Configuring a Virtual Tunnel Interface 641
Allowing Traffic Through the Site-to-Site VPN 642
Configuring the Global IKE Policy 642
Configuring IKEv1 Policies 643
Configuring IKEv2 Policies 645
Configuring IPsec Proposals 647
Configuring IPsec Proposals for IKEv1 647

Configuring IPsec Proposals for IKEv2 **648**
Verifying Site-to-Site VPN Connections **649**
Monitoring Site-to-Site VPN **652**
Examples for Site-to-Site VPN **652**
Exempting Site-to-Site VPN Traffic from NAT **652**
How to Provide Internet Access on the Outside Interface for External Site-to-Site VPN Users (Hair Pinning) **658**
How to Secure Traffic from Networks in Multiple Virtual Routers over a Site-to-Site VPN **665**

CHAPTER 25 **Remote Access VPN** **671**
Remote Access VPN Overview **671**
Maximum Concurrent VPN Sessions By Device Model **671**
Downloading the AnyConnect Client Software **672**
How Users Can Install the AnyConnect Client Software **673**
Controlling User Permissions and Attributes Using RADIUS and Group Policies **673**
Attributes Sent to the RADIUS Server **674**
Attributes Received from the RADIUS Server **674**
Two-Factor Authentication **675**
RSA Two-Factor Authentication **676**
Duo Two-Factor Authentication Using RADIUS **676**
Duo Two-Factor Authentication using LDAP **677**
Licensing Requirements for Remote Access VPN **677**
Guidelines and Limitations for Remote Access VPN **678**
Configuring Remote Access VPN **678**
Configure and Upload Client Profiles **680**
Allow Traffic Through the Remote Access VPN **682**
Verify the Remote Access VPN Configuration **683**
Managing the Remote Access VPN Configuration **684**
Configure an RA VPN Connection Profile **685**
Configure AAA for a Connection Profile **688**
Configure Certificate Authentication for a Connection Profile **691**
Configure Client Addressing for RA VPN **692**
Configure Group Policies for RA VPN **692**
General Attributes **693**

Session Settings Attributes 694
Address Assignment Attributes 694
Split Tunneling Attributes 695
AnyConnect Client Attributes 696
Traffic Filters Attributes 697
Windows Browser Proxy Attributes 698
Monitoring Remote Access VPN 698
Troubleshooting Remote Access VPNs 698
Troubleshooting SSL Connection Problems 699
Troubleshooting AnyConnect Client Download and Installation Problems 699
Troubleshooting AnyConnect Client Connection Problems 699
Troubleshooting RA VPN Traffic Flow Problems 700
Examples for Remote Access VPN 701
How to Implement RADIUS Change of Authorization 701
System Flow for Change of Authorization 701
Configure Change of Authorization on the Firewall Threat Defense Device 703
Configure Change of Authorization in ISE 706
How to Configure Two-Factor Authentication using Duo LDAP 709
System Flow for Duo LDAP Secondary Authentication 709
Configure Duo LDAP Secondary Authentication 710
How to Provide Internet Access on the Outside Interface for Remote Access VPN Users (Hair Pinning) 716
How to Use a Directory Server on an Outside Network with Remote Access VPN 720
How to Control RA VPN Access By Group 733
How to Allow RA VPN Access to Internal Networks in Different Virtual Routers 737
How to Customize the AnyConnect Client Icon and Logo 740

PART VII **System Administration** 745

CHAPTER 26 **System Settings** 747

Configuring Management Access 747
Configuring the Management Access List 748
Configuring the HTTPS Port for Management Access on Data Interfaces 749
Configuring the Firewall Threat Defense Web Server Certificate 750

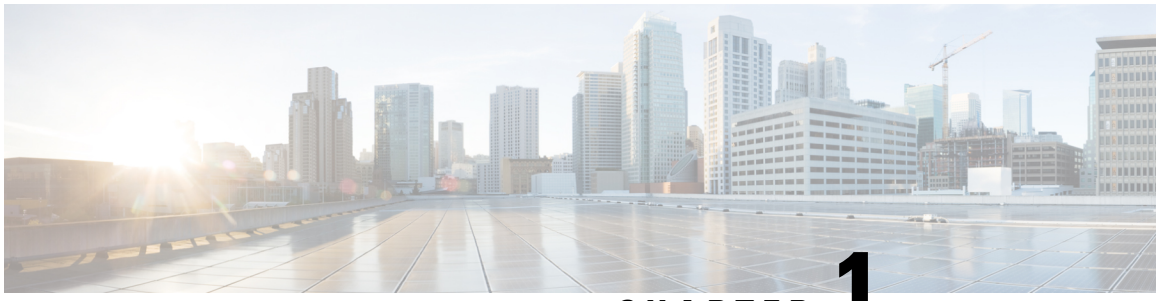
Configuring System Logging Settings 751
Syslog message severity levels 751
Configure Logging to a Remote Syslog Server 752
Configure Logging to the Internal Buffer 753
Configure Logging to the Console 754
Configure Event List Filters 754
Configuring DHCP 755
Configuring the DHCP Server 756
Configuring DHCP Relay 758
Configuring Dynamic DNS 759
Configuring DNS 761
Configuring DNS Groups 762
Configuring DNS for Data and Management Traffic 763
Troubleshooting General DNS Problems 764
Configuring the Management Interface 766
Configuring the Device Hostname 767
Configuring Network Time Protocol (NTP) 768
Configuring Precision Time Protocol (ISA 3000) 769
Configuring HTTP Proxy for Management Connections 771
Configuring Cloud Services 772
Enabling or Disabling Security Cloud Control (Legacy Device Manager Mode) 773
Connecting to the Cisco Success Network 774
Sending Events to the Cisco Cloud 775
Unregistering from Cloud Services 776
Enabling or Disabling Web Analytics 776
Configuring URL Filtering Preferences 777
Switch from the Firewall Device Manager to the Firewall Management Center 778
Switch from the Firewall Management Center to the Firewall Device Manager 783
Configuring TLS/SSL Cipher Settings 784
Configure TLS/SSL Cipher Objects 785

CHAPTER 27 **System Management** 787
Installing Software Updates 787
Updating System Databases and Feeds 787

Overview of System Database and Feed Updates 787
Updating System Databases 788
Updating Cisco Security Intelligence Feeds 790
Upgrading Firewall Threat Defense 791
Running an Upgrade Readiness Check 793
Monitoring Firewall Threat Defense Upgrades 794
Canceling or Retrying Firewall Threat Defense Upgrades 794
Reverting Firewall Threat Defense 794
Troubleshooting Threat Defense Upgrades 795
Reimaging the Device 796
Backing Up and Restoring the System 796
Backing Up the System Immediately 797
Backing Up the System at a Scheduled Time 798
Setting Up a Recurring Backup Schedule 799
Restoring a Backup 800
Managing Backup Files 801
Auditing and Change Management 802
Audit Events 802
Viewing and Analyzing the Audit Log 803
Filtering the Audit Log 804
Examining Deployment and Entity Change History 806
Discarding All Pending Changes 807
Exporting the Device Configuration 808
Managing Firewall Device Manager and Firewall Threat Defense User Access 808
Configuring External Authorization (AAA) for the Firewall Device Manager (HTTPS) Users 809
Configuring External Authorization (AAA) for the Firewall Threat Defense CLI (SSH) Users 810
Managing the Firewall Device Manager User Sessions 812
Enabling the Firewall Device Manager Access on a Standby HA Unit for External Users 812
Creating Local User Accounts for the Firewall Threat Defense CLI 813
Rebooting or Shutting Down the System 815
Troubleshooting the System 816
Pinging Addresses to Test Connectivity 816
Tracing Routes to Hosts 818
Making the Device Appear on Traceroutes 819

Troubleshooting NTP 821
Troubleshooting DNS for the Management Interface 822
Analyzing CPU and Memory Usage 825
Viewing Logs 825
Creating a Troubleshooting File 827
Hardware Management Tasks 827
Replacing an ISA 3000 Device 827
Hot swap an SSD 827

APPENDIX A **Advanced Configuration** 831
About Smart CLI and FlexConfig 831
Recommended Usage for Smart CLI and FlexConfig 832
CLI Commands in Smart CLI and FlexConfig Objects 832
How Software Upgrades Affect the FlexConfig Policy 833
Determine the ASA Software Version and Current CLI Configuration 833
Prohibited CLI Commands 834
Smart CLI Templates 839
Guidelines and Limitations for Smart CLI and FlexConfig 840
Configuring Smart CLI Objects 840
Configuring the FlexConfig Policy 841
Configuring FlexConfig Objects 843
Creating Variables in a FlexConfig Object 845
Referring to FlexConfig Variables and Retrieving Values 846
Variable References: {{variable}} or {{{variable}}} 846
Sections {{#key}} {{/key}} and Inverse Sections {{^key}} {{/key}} 849
Referring to Smart CLI Objects in a FlexConfig Object 851
Configuring Secret Key Objects 852
Troubleshooting the FlexConfig Policy 853
Examples for FlexConfig 854
How to Enable and Disable Global Default Inspections 854
How to Undo Your FlexConfig Changes 860
How to Enable Inspections for Unique Traffic Classes 861



CHAPTER 1

Getting Started

The following topics explain how to get started configuring the Firewall Threat Defense device using Firewall Device Manager.

- Is This Guide for You?, on page 1
- New Features in Firewall Device Manager/Firewall Threat Defense Version 7.1.0, on page 1
- Default Configuration, on page 7
- Logging Into the System, on page 14
- Setting Up the System, on page 18
- Configuration Basics, on page 24
- Communication Port and Internet Access Requirements, on page 33

Is This Guide for You?

This guide explains how to configure Firewall Threat Defense using the Firewall Device Manager web-based configuration interface included on the Firewall Threat Defense devices.

The Firewall Device Manager lets you configure the basic features of the software that are most commonly used for small or mid-size networks. It is especially designed for networks that include a single device or just a few, where you do not want to use a high-powered multiple-device manager to control a large network containing many Firewall Threat Defense devices.

If you are managing large numbers of devices, or if you want to use the more complex features and configurations that Firewall Threat Defense allows, use the Firewall Management Center to configure your devices instead of the integrated Firewall Device Manager.

New Features in Firewall Device Manager/Firewall Threat Defense Version 7.1.0

Released: December 1, 2021

The following table lists the new features available in Firewall Threat Defense 7.1.0 when configured using Firewall Device Manager.

Feature	Description
Platform Features	
Secure Firewall 3100.	We introduced the Secure Firewall 3110, 3120, 3130, and 3140. You can hot swap a network module of the same type while the firewall is powered up without having to reboot; making other module changes requires a reboot. Secure Firewall 3100 25 Gbps interfaces support Forward Error Correction as well as speed detection based on the SFP installed. The SSDs are self-encrypting drives (SEDs), and if you have 2 SSDs, they form a software RAID. Note that the Version 7.1 device manager does not include online help for these devices. See the documentation posted on Cisco.com. New/Modified screens: **Device** > **Interfaces** New/Modified Firewall Threat Defense commands: **configure network speed, configure raid, show raid, show ssd**

Feature	Description
FTDv for AWS instances.	FTDv for AWS adds support for these instances: • c5a.xlarge, c5a.2xlarge, c5a.4xlarge • c5ad.xlarge, c5ad.2xlarge, c5ad.4xlarge • c5d.xlarge, c5d.2xlarge, c5d.4xlarge • c5n.xlarge, c5n.2xlarge, c5n.4xlarge • i3en.xlarge, i3en.2xlarge, i3en.3xlarge • inf1.xlarge, inf1.2xlarge • m5.xlarge, m5.2xlarge, m5.4xlarge • m5a.xlarge, m5a.2xlarge, m5a.4xlarge • m5ad.xlarge, m5ad.2xlarge, m5ad.4xlarge • m5d.xlarge, m5d.2xlarge, m5d.4xlarge • m5dn.xlarge, m5dn.2xlarge, m5dn.4xlarge • m5n.xlarge, m5n.2xlarge, m5n.4xlarge • m5zn.xlarge, m5zn.2xlarge, m5zn.3xlarge • r5.xlarge, r5.2xlarge, r5.4xlarge • r5a.xlarge, r5a.2xlarge, r5a.4xlarge • r5ad.xlarge, r5ad.2xlarge, r5ad.4xlarge • r5b.xlarge, r5b.2xlarge, r5b.4xlarge • r5d.xlarge, r5d.2xlarge, r5d.4xlarge • r5dn.xlarge, r5dn.2xlarge, r5dn.4xlarge • r5n.xlarge, r5n.2xlarge, r5n.4xlarge • z1d.xlarge, z1d.2xlarge, z1d.3xlarge
FTDv for Azure instances.	FTDv for Azure adds support for these instances: • Standard_D8s_v3 • Standard_D16s_v3 • Standard_F8s_v2 • Standard_F16s_v2
Support ends for the ASA 5508-X and 5516-X. The last supported release is 7.0.	You cannot install Firewall Threat Defense 7.1 on an ASA 5508-X or 5516-X. The last supported release for these models is Firewall Threat Defense 7.0.

Feature	Description
Firewall and IPS Features	
Network Analysis Policy (NAP) configuration for Snort 3.	You can use Firewall Device Manager to configure the Network Analysis Policy (NAP) when running Snort 3. Network analysis policies control traffic preprocessing inspection. Inspectors prepare traffic to be further inspected by normalizing traffic and identifying protocol anomalies. You can select which NAP is used for all traffic, and customize the settings to work best with the traffic in your network. You cannot configure the NAP when running Snort 2. We added the Network Analysis Policy to the **Policies** > **Intrusion** settings dialog box, with an embedded JSON editor to allow direct changes, and other features to let you upload overrides, or download the ones you create.
Manual NAT support for fully-qualified domain name (FQDN) objects as the translated destination.	You can use an FQDN network object, such as one specifying www.example.com, as the translated destination address in manual NAT rules. The system configures the rule based on the IP address returned from the DNS server.
Improved active authentication for identity rules.	You can configure active authentication for identity policy rules to redirect the user's authentication to a fully-qualified domain name (FQDN) rather than the IP address of the interface through which the user's connection enters the device. The FQDN must resolve to the IP address of one of the interfaces on the device. By using an FQDN, you can assign a certificate for active authentication that the client will recognize, thus avoiding the untrusted certificate warning users get when being redirected to an IP address. The certificate can specify the FQDN, a wildcard FQDN, or multiple FQDNs in the Subject Alternate Names (SAN) in the certificate. We added the **Redirect to Host Name** option in the identity policy settings.
VPN Features	
Backup remote peers for site-to-site VPN.	You can configure a site-to-site VPN connection to include remote backup peers. If the primary remote peer is unavailable, the system will try to re-establish the VPN connection using one of the backup peers. You can configure separate pre-shared keys or certificates for each backup peer. Backup peers are supported for policy-based connections only, and are not available for route-based (virtual tunnel interface) connections. We updated the site-to-site VPN wizard to include backup peer configuration.

Feature	Description
Password management for remote access VPN (MSCHAPv2).	You can enable password management for remote access VPN. This allows AnyConnect to prompt the user to change an expired password. Without password management, users must change expired passwords directly with the AAA server, and AnyConnect does not prompt the user to change passwords. For LDAP servers, you can also set a warning period to notify users of upcoming password expiration. We added the Enable Password Management option to the authentication settings for remote access VPN connection profiles.
AnyConnect VPN SAML external browser.	When you use SAML as the primary authentication method for a remote access VPN connection profile, you can elect to have the AnyConnect client use the client's local browser instead of the AnyConnect embedded browser to perform the web authentication. This option enables single sign-on (SSO) between your VPN authentication and other corporate logins. Also choose this option if you want to support web authentication methods, such as biometric authentication, that cannot be performed in the embedded browser. We updated the remote access VPN connection profile wizard to allow you to configure the **SAML Login Experience**.
Administrative and Troubleshooting Features	
Dynamic Domain Name System (DDNS) support for updating fully-qualified domain name (FQDN) to IP address mappings for system interfaces.	**Upgrade impact. Redo FlexConfigs after upgrade.** You can configure DDNS for the interfaces on the system to send dynamic updates to DNS servers. This helps ensure that FQDNs defined for the interfaces resolve to the correct address, making it easier for users to access the system using a hostname rather than an IP address. This is especially useful for interfaces that get their addresses using DHCP, but it is also useful for statically-addressed interfaces. After upgrade, if you had used FlexConfig to configure DDNS, you must redo your configuration using Firewall Device Manager or the Firewall Threat Defense API, and remove the DDNS FlexConfig object from the FlexConfig policy, before you can deploy changes again. If you configure DDNS using Firewall Device Manager, then switch to Firewall Management Center management, the DDNS configuration is retained so that Firewall Management Center can find the system using the DNS name. In Firewall Device Manager, we added the **System Settings** > **DDNS Service** page. In the Firewall Threat Defense API, we added the DDNSService and DDNSInterfaceSettings resources.
The **dig** command replaces the **nslookup** command in the device CLI.	To look up the IP address of a fully-qualified domain name (FQDN) in the device CLI, use the **dig** command. The **nslookup** command has been removed.

Feature	Description
DHCP relay configuration using Firewall Device Manager.	You can use Firewall Device Manager to configure DHCP relay. Using DHCP relay on an interface, you can direct DHCP requests to a DHCP server that is accessible through the other interface. You can configure DHCP relay on physical interfaces, subinterfaces, EtherChannels, and VLAN interfaces. You cannot configure DHCP relay if you configure a DHCP server on any interface. We added the **System Settings** > **DHCP** > **DHCP Relay** page, and moved DHCP Server under the new DHCP heading.
Key type and size for self-signed certificates in Firewall Device Manager.	You can specify the key type and size when generating new self-signed internal and internal CA certificates in Firewall Device Manager. Key types include RSA, ECDSA, and EDDSA. The allowed sizes differ by key type. We now warn you if you upload a certificate whose key size is smaller than the minimum recommended length. There is also a weak key pre-defined search filter to help you find weak certificates, which you should replace if possible.
Usage validation restrictions for trusted CA certificates.	You can specify whether a trusted CA certificate can be used to validate certain types of connections. You can allow, or prevent, validation for SSL server (used by dynamic DNS), SSL client (used by remote access VPN), IPsec client (used by site-to-site VPN), or other features that are not managed by the Snort inspection engine, such as LDAPS. The primary purpose of these options is to let you prevent VPN connections from getting established because they can be validated against a particular certificate. We added **Validation Usage** as a property for trusted CA certificates.
Generating the **admin** password in Firewall Device Manager.	During initial system configuration in Firewall Device Manager, or when you change the **admin** password through Firewall Device Manager, you can now click a button to generate a random 16 character password.
Startup time and tmatch compilation status.	The **show version** command now includes information on how long it took to start (boot) up the system. Note that the larger the configuration, the longer it takes to boot up the system. The new **show asp rule-engine** command shows status on tmatch compilation. Tmatch compilation is used for an access list that is used as an access group, the NAT table, and some other items. It is an internal process that can consume CPU resources and impact performance while in progress, if you have very large ACLs and NAT tables. Compilation time depends on the size of the access list, NAT table, and so forth.
Enhancements to **show access-list element-count** output.	The output of the **show access-list element-count** command has been enhanced. When used with object-group search enabled, the output includes details about the number of object groups in the element count. In addition, the **show tech-support** output now includes the output from **show access-list element-count** and **show asp rule-engine**.

Feature	Description
Use Firewall Device Manager to configure the Firewall Threat Defense for management by a Firewall Management Center.	When you perform initial setup using Firewall Device Manager, all interface configuration completed in Firewall Device Manager is retained when you switch to Firewall Management Center for management, in addition to the Management and Firewall Management Center access settings. Note that other default configuration settings, such as the access control policy or security zones, are not retained. When you use the Firewall Threat Defense CLI, only the Management and Firewall Management Center access settings are retained (for example, the default inside interface configuration is not retained). After you switch to Firewall Management Center, you can no longer use Firewall Device Manager to manage the Firewall Threat Defense. New/Modified screens: **System Settings** > **Management Center**
Automatically update CA bundles.	**Upgrade impact. The system connects to Cisco for something new.** The local CA bundle contains certificates to access several Cisco services. The system now automatically queries Cisco for new CA certificates at a daily system-defined time. Previously, you had to upgrade the software to update CA certificates. You can use the CLI to disable this feature. New resources: https://cisco.com/security/pki/ New/modified CLI commands: **configure cert-update auto-update**, **configure cert-update run-now**, **configure cert-update test**, **show cert-update** Version restrictions: Requires Version 7.0.5, 7.1.0.3, or 7.2.4+. See: *Cisco Secure Firewall Threat Defense Command Reference*
FTD REST API version 6.2 (v6).	The Firewall Threat Defense REST API for software version 7.1 is version 6.2. You can use v6 in the API URLs, or preferentially, use /latest/ to signify you are using the most recent API version that is supported on the device. Note that the URL version path element for 6.2 is the same as 6.0/1: v6. Please re-evaluate all existing calls, as changes might have been mode to the resource models you are using. To open the API Explorer, where you can view the resources, log into Firewall Threat Defense, then click the more options button (⋮) and choose **API Explorer**.

Default Configuration

The default configuration of your device depends on whether you have completed initial setup.

Default Configuration Prior to Initial Setup

Before you initially configure the Firewall Threat Defense device using the local manager (Firewall Device Manager), the device includes the following default configuration.

For many models, this configuration assumes that you open the device manager through the inside interface, typically by plugging your computer directly into the interface, and use the DHCP server defined on the inside interface to supply your computer with an IP address. Alternatively, you can plug your computer into the Management interface and use DHCP to obtain an address. However, some models have different default configurations and management requirements. See the table below for details.

Note You can pre-configure many of these settings using the CLI setup ((Optional) Change Management Network Settings at the CLI, on page 19) before you perform setup using the wizard.

Default Configuration Settings

Setting	Default	Can be changed during initial configuration?
Password for admin user.	Admin123 Firepower 4100/9300: Set the password when you deploy the logical device. AWS: The default is the AWS Instance ID, unless you define a default password with user data (**Advanced Details** > **User Data**) during the initial deployment.	Yes. You must change the default password.
Management IP address.	Obtained through DHCP. Firewall Threat Defense Virtual: 192.168.45.45 Firepower 4100/9300: Set the management IP address when you deploy the logical device.	No. For Firepower 4100/9300: Yes.
Management gateway.	The data interfaces on the device. Typically the outside interface becomes the route to the Internet. This gateway works for from-the-device traffic only. If the device receives a default gateway from the DHCP server, then that gateway is used. Firepower 4100/9300: Set the gateway IP address when you deploy the logical device. ISA 3000: 192.168.45.1. Firewall Threat Defense Virtual: 192.168.45.1	No. For Firepower 4100/9300: Yes.

Setting	Default	Can be changed during initial configuration?
DNS servers for the management interface.	The OpenDNS public DNS servers, IPv4: 208.67.220.220 and 208.67.222.222; IPv6: 2620:119:35::35. DNS servers obtained from DHCP are never used. Firepower 4100/9300: Set the DNS servers when you deploy the logical device.	Yes
Inside interface IP address.	192.168.95.1/24 Firepower 4100/9300: Data interfaces are not pre-configured. ISA 3000: BVI1 IP address is not preconfigured. BVI1 includes all inside and outside interfaces. Firewall Threat Defense Virtual: 192.168.45.1/24	No.
DHCP server for inside clients.	Running on the inside interface with the address pool 192.168.95.5 - 192.168.95.254. Firepower 4100/9300: No DHCP server enabled. ISA 3000: No DHCP server enabled. Firewall Threat Defense Virtual: The address pool on the inside interface is 192.168.45.46 - 192.168.45.254.	No.
DHCP auto-configuration for inside clients. (Auto-configuration supplies clients with addresses for WINS and DNS servers.)	Enabled on outside interface.	Yes, but indirectly. If you configure a static IPv4 address for the outside interface, DHCP server auto-configuration is disabled.
Outside interface IP address.	IPv4: Obtained through DHCP from Internet Service Provider (ISP) or upstream router. IPv6: Autoconfiguration. Firepower 4100/9300: Data interfaces are not pre-configured. ISA 3000: BVI1 IP address is not preconfigured. BVI1 includes all inside and outside interfaces.	Yes.

Default Interfaces by Device Model

You cannot select different inside and outside interfaces during initial configuration. To change the interface assignments after configuration, edit the interface and DHCP settings. You must remove an interface from the bridge group before you can configure it as a non-switched interface.

Device Model	Outside Interface	Inside Interface
Firepower 1010	Ethernet1/1	VLAN1, which includes all other switch ports except the outside interface, which is a physical firewall interface.
Firepower 1120, 1140, 1150	Ethernet1/1	Ethernet1/2
Firepower 2100	Ethernet1/1	Ethernet1/2
Secure Firewall 3100	Ethernet1/1	Ethernet1/2
Firepower 4100	Data interfaces are not pre-configured.	Data interfaces are not pre-configured.
Firepower 9300	Data interfaces are not pre-configured.	Data interfaces are not pre-configured.
Firewall Threat Defense Virtual	GigabitEthernet0/0	GigabitEthernet0/1
ISA 3000	GigabitEthernet1/1 and GigabitEthernet1/3 GigabitEthernet1/1 (outside1) and 1/2 (inside1), and GigabitEthernet1/3 (outside2) and 1/4 (inside2) (non-fiber models only) are configured as Hardware Bypass pairs. All inside and outside interfaces are part of BVI1.	GigabitEthernet1/2 and GigabitEthernet1/4

Configuration After Initial Setup

After you complete the setup wizard, the device configuration will include the following settings. The table shows whether a particular setting is something you explicitly chose or whether it was defined for you based on your other selections. Validate any "implied" configurations and edit them if they do not serve your needs.

Note The Firepower 4100/9300 and ISA 3000 do not support the setup wizard. For the Firepower 4100/9300 , all initial configuration is set when you deploy the logical device from the chassis. For the ISA 3000, a special default configuration is applied before shipping.

Setting	Configuration	Explicit, implied, or default configuration
Password for admin user.	Whatever you entered.	Explicit.

Setting	Configuration	Explicit, implied, or default configuration
Management IP address.	Obtained through DHCP. Firewall Threat Defense Virtual: 192.168.45.45 Firepower 4100/9300 : The management IP address you set when you deployed the logical device.	Default.
Management gateway.	The data interfaces on the device. Typically the outside interface becomes the route to the Internet. The management gateway works for from-the-device traffic only. If the device receives a default gateway from the DHCP server, then that gateway is used. Firepower 4100/9300 : The gateway IP address you set when you deployed the logical device. ISA 3000: 192.168.45.1 Firewall Threat Defense Virtual: 192.168.45.1	Default.
DNS servers for the management interface.	The OpenDNS public DNS servers, IPv4: 208.67.220.220, 208.67.222.222; IPv6: 2620:119:35::35, or whatever you entered. DNS servers obtained from DHCP are never used. Firepower 4100/9300 : The DNS servers you set when you deployed the logical device.	Explicit.
Management hostname.	**firepower** or whatever you entered. Firepower 4100/9300 : The hostname you set when you deployed the logical device.	Explicit.
Management access through data interfaces.	A data interface management access list rule allows HTTPS access through the inside interface. SSH connections are not allowed. Both IPv4 and IPv6 connections are allowed. Firepower 4100/9300 : No data interfaces have default management access rules. ISA 3000: No data interfaces have default management access rules. Firewall Threat Defense Virtual: No data interfaces have default management access rules.	Implied.
System time.	The time zone and NTP servers you selected. Firepower 4100/9300 : System time is inherited from the chassis. ISA 3000: Cisco NTP servers: 0.sourcefire.pool.ntp.org, 1.sourcefire.pool.ntp.org, 2.sourcefire.pool.ntp.org.	Explicit.
Smart license.	Either registered with a base license, or the evaluation period activated, whichever you selected. Subscription licenses are not enabled. Go to the smart licensing page to enable them.	Explicit.

Setting	Configuration	Explicit, implied, or default configuration
Inside interface IP address.	192.168.95.1/24 Firepower 4100/9300 : Data interfaces are not pre-configured. ISA 3000: None. You must set the BVI1 IP address manually. Firewall Threat Defense Virtual: 192.168.45.1/24	Default.
DHCP server for inside clients.	Running on the inside interface with the address pool 192.168.95.5 - 192.168.95.254. Firepower 4100/9300 : No DHCP server enabled. ISA 3000: No DHCP server enabled. Firewall Threat Defense Virtual: The address pool on the inside interface is 192.168.45.46 - 192.168.45.254.	Default.
DHCP auto-configuration for inside clients. (Auto-configuration supplies clients with addresses for WINS and DNS servers.)	Enabled on outside interface if you use DHCP to obtain the outside interface IPv4 address. If you use static addressing, DHCP auto-configuration is disabled.	Explicit, but indirectly.
Data interface configuration.	• 1010—The outside interface, Ethernet1/1, is a physical firewall interface. All other interfaces are switch ports that are enabled and part of VLAN1, the inside interface. You can plug end points or switches into these ports and obtain addresses from the DHCP server for the inside interface. • Firepower 4100/9300 —All data interfaces are disabled. • ISA 3000—All data interfaces are enabled and part of the same bridge group, BVI1. GigabitEthernet1/1 and 1/3 are outside interfaces, and GigabitEthernet1/2 and 1/4 are inside interfaces. GigabitEthernet1/1 (outside1) and 1/2 (inside1), and GigabitEthernet1/3 (outside2) and 1/4 (inside2) (non-fiber models only) are configured as Hardware Bypass pairs. • All other models—The outside and inside interfaces are the only ones configured and enabled. All other data interfaces are disabled.	Default.
Outside physical interface and IP address.	The default outside port based on the device model. See Default Configuration Prior to Initial Setup, on page 8. The IP address is obtained by DHCP and IPv6 autoconfiguration, or it is a static address as entered (IPv4, IPv6, or both). Firepower 4100/9300 : Data interfaces are not pre-configured. ISA 3000: None. You must set the BVI1 IP address manually.	Interface is Default. Addressing is Explicit.

Setting	Configuration	Explicit, implied, or default configuration
Static routes.	If you configure a static IPv4 or IPv6 address for the outside interface, a static default route is configured for IPv4/IPv6 as appropriate, pointing to the gateway you defined for that address type. If you select DHCP, the default route is obtained from the DHCP server. Network objects are also created for the gateway and the "any" address, that is, 0.0.0.0/0 for IPv4, ::/0 for IPv6.	Implied.
Security zones.	**inside_zone**, containing the inside interfaces. For the Firepower 4100/9300 , you need to add interfaces manually to this security zone. **outside_zone**, containing the outside interfaces. For the Firepower 4100/9300 , you need to add interfaces manually to this zone. (You can edit these zones to add other interfaces, or create your own zones.)	Implied.
Access control policy.	A rule trusting all traffic from the inside_zone to the outside_zone. This allows without inspection all traffic from users inside your network to get outside, and all return traffic for those connections. The default action for any other traffic is to block it. This prevents any traffic initiated from outside to enter your network. Firepower 4100/9300 : There are no pre-configured access rules. ISA 3000: A rule trusting all traffic from the inside_zone to the outside_zone, and a rule trusting all traffic from the outside_zone to the inside_zone. Traffic is not blocked. The device also has rules trusting all traffic between the interfaces in the inside_zone and in the outside_zone. This allows without inspection all traffic between users on the inside, and between users on the outside.	Implied.
NAT	An interface dynamic PAT rule translates the source address for any IPv4 traffic destined to the outside interface to a unique port on the outside interface's IP address. There are additional hidden PAT rules to enable HTTPS access through the inside interfaces, and routing through the data interfaces for the management address. These do not appear in the NAT table, but you will see them if you use the **show nat** command in the CLI. Firepower 4100/9300 : NAT is not pre-configured. ISA 3000: NAT is not pre-configured.	Implied.

Logging Into the System

There are two interfaces to the Firewall Threat Defense device:

Firewall Device Manager Web Interface

The Firewall Device Manager runs in your web browser. You use this interface to configure, manage, and monitor the system.

Command Line Interface (CLI, Console)

Use the CLI for troubleshooting. You can also use it for initial setup instead of the Firewall Device Manager.

The following topics explain how to log into these interfaces and manage your user account.

Your User Role Controls What You Can See and Do

Your username is assigned a role, and your role determines what you can do or what you can see in the Firewall Device Manager. The locally-defined **admin** user has all privileges, but if you log in using a different account, you might have fewer privileges.

The upper-right corner of the Firewall Device Manager window shows your username and privilege level.



The privileges are:

- **Administrator**—You can see and use all features.
- **Read-Write User**—You can do everything a read-only user can do, but also edit and deploy the configuration. The only restrictions are for system-critical actions, which include installing upgrades, creating and restoring backups, viewing the audit log, and ending the sessions of other Firewall Device Manager users.
- **Read-Only User**—You can view dashboards and the configuration, but you cannot make any changes. If you try to make a change, the error message explains that this is due to lack of permission.

These privileges are not related to those available for CLI users.

Logging Into the Firewall Device Manager

Use the Firewall Device Manager to configure, manage, and monitor the system. The features that you can configure through the browser are not configurable through the command-line interface (CLI); you must use the web interface to implement your security policies.

Use a current version of the following browsers: Firefox, Chrome, Safari, Edge.

Note If you type in the wrong password and fail to log in on 3 consecutive attempts, your account is locked for 5 minutes. You must wait before trying to log in again.

Before you begin

Initially, you can log into the Firewall Device Manager using the **admin** username only. However, you can then configure authorization for additional users defined in an external AAA server, as described in Managing Firewall Device Manager and Firewall Threat Defense User Access, on page 808.

There can be up to 5 active logins at one time. This includes users logged into the device manager and active API sessions, which are represented by non-expired API tokens. If you exceed this limit, the oldest session, either the device manager login or API token, is expired to allow the new session. These limits do not apply to SSH sessions.

Procedure

Step 1 Using a browser, open the home page of the system, for example, https://ftd.example.com.

You can use any of the following addresses. You can use the IPv4 or IPv6 address or the DNS name, if you have configured one.

- The management address. By default (on most platforms), the Management interface is a DHCP client, so the IP address depends on your DHCP server.
- The address of a data interface that you have opened for HTTPS access. By default (on most platforms), the "inside" interface allows HTTPS access, so you can connect to the default inside address 192.168.95.1. See Default Configuration Prior to Initial Setup, on page 8 for details about your model's inside IP address.

  If you changed the HTTPS data port, you must include the custom port in the URL. For example, if you changed the port to 4443: https://ftd.example.com:4443

Tip
If your browser is not configured to recognize the server certificate, you will see a warning about an untrusted certificate. Accept the certificate as an exception, or in your trusted root certificate store.

Step 2 If the login screen includes text, such as a warning or disclaimer, read the information, then select the checkbox to acknowledge your agreement.

Step 3 Enter your username and password defined for the device, then click **Login**.

You can use the **admin** username, which is a pre-defined user. The default admin password is Admin123. On AWS, the default admin password is the AWS Instance ID, unless you define a default password with user data (**Advanced Details** > **User Data**) during the initial deployment.

Your session will expire after 30 minutes of inactivity, and you will be prompted to log in again. You can log out by selecting **Log Out** from the user icon drop-down menu in the upper right of the page.



Logging Into the Command Line Interface (CLI)

Use the command-line interface (CLI) to set up the system and do basic system troubleshooting. You cannot configure policies through a CLI session.

To log into the CLI, do one of the following:

- 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.



Note On the Firepower and Secure Firewall device models, the CLI on the Console port is the Secure Firewall eXtensible Operating System (FXOS). For the some device models, you can get to the Firewall Threat Defense CLI using the **connect ftd** command. For the Firepower 4100/9300, see Connect to the Console of the Application, on page 189. Use the FXOS CLI for chassis-level troubleshooting only. Use the Firewall Threat Defense CLI for basic configuration, monitoring, and normal system troubleshooting. See the FXOS documentation for information on FXOS commands.

- For the Firewall Threat Defense Virtual, open the virtual console.
- Use an SSH client to make a connection to the management IP address. You can also connect to the address on a data interface if you open the interface for SSH connections (see Configuring the Management Access List, on page 748). SSH access to data interfaces is disabled by default. Log in using the **admin** username or another CLI user account. The default admin password is Admin123. On AWS, the default admin password for the Firewall Threat Defense Virtual is the AWS Instance ID, unless you define a default password with user data (**Advanced Details** > **User Data**) during the initial deployment.

Tips

- If the administrator added text to the login process, such as a warning or disclaimer, you need to acknowledge that you read and agree to the statement.
- After logging in, for information on the commands available in the CLI, enter **help** or **?**. For usage information, see Cisco Firepower Threat Defense Command Reference at http://www.cisco.com/c/en/us/td/docs/security/firepower/command_ref/b_Command_Reference_for_Firepower_Threat_Defense.html.
- You can create local user accounts that can log into the CLI using the **configure user add** command. However, these users can log into the CLI only. They cannot log into the Firewall Device Manager web interface.
- You can create user accounts for SSH access in an external server. For information about configuring external authentication for SSH access, see Configuring External Authorization (AAA) for the Firewall Threat Defense CLI (SSH) Users, on page 810.

Changing Your Password

You should periodically change your password. The following procedure explains how to change the password while logged into Firewall Device Manager.



Note If you are logged into the CLI, you can change your password using the **configure password** command. You can change the password for a different CLI user with the **configure user password** *username* command.

Before you begin

This procedure applies to local users only. If your user account is defined on an external AAA server, you must change your password with that server.

Procedure

Step 1 Select **Profile** from the user icon drop-down list in the upper right of the menu.



Step 2 Click the **Password** tab.

Step 3 Enter your current password.

Step 4 Enter your new password and then confirm it.

You can click **Generate** to have a random 16 character password generated for you. Click the Show Password () button to see the passwords unmasked. Then, click the **Copy To Clipboard** link so you can paste the password in the confirmation field.

The page includes the minimum requirements for a password. You cannot change these minimum requirements. Passwords must:

- Be 8 – 128 characters
- Have at least one lower and one upper case letter
- Have at least one digit
- Have at least one special character
- Not contain repeated letters

Step 5 Click **Change**.

Setting User Profile Preferences

You can set preferences for the user interface and change your password.

Procedure

Step 1 Select **Profile** from the user icon drop-down list in the upper right of the menu.



Step 2 On the **Profile** tab, configure the following and click **Save**.

- **Time Zone for Scheduling Tasks**—Select the time zone you want to use for scheduling tasks such as backups and updates. The browser time zone is used for dashboards and events, if you set a different zone.
- **Color Theme**—Select the color theme you want to use in the user interface.

Step 3 On the **Password** tab, you can enter a new password and click **Change**.

Viewing Pages in Languages Other Than English

You can view the GUI and online help in the following languages.

- Chinese
- English (the default)
- Japanese
- Korean

To use these languages, you must select that language in your browser settings. There is no language setting in the product itself.

If your browser does not support a particular language, the product will not appear in that language. For example, the French version appears only if you configure the browser to use Canadian French. Selecting another type of French leaves the product in English.

Setting Up the System

You must complete an initial configuration to make the system function correctly in your network. Successful deployment includes attaching cables correctly and configuring the addresses needed to insert the device into your network and connect it to the Internet or other upstream router. The following procedure explains the process.

Before you begin

Before you start the initial setup, the device includes some default settings. For details, see Default Configuration Prior to Initial Setup, on page 8.

Procedure

Step 1 Connect the Interfaces, on page 19

Step 2 Complete the Initial Configuration Using the Setup Wizard, on page 21

For details about the resulting configuration, see Configuration After Initial Setup, on page 10.

Connect the Interfaces

The default configuration assumes that certain interfaces are used for the inside and outside networks. Initial configuration will be easier to complete if you connect network cables to the interfaces based on these expectations.

The default configuration for most models is designed to let you attach your management computer to the inside interface. Alternatively, you can also directly attach your workstation to the Management port. The interfaces are on different networks, so do not try to connect any of the inside interfaces and the Management port to the same network.

Do not connect any of the inside interfaces to a network that has an active DHCP server. This will conflict with the DHCP server already running on the inside interface . If you want to use a different DHCP server for the network, disable the unwanted DHCP server after initial setup.

See the getting started guide for your model for cabling diagrams.

(Optional) Change Management Network Settings at the CLI

If you cannot use the default management IP address, then you can connect to the console port and perform initial setup at the CLI, including setting the Management IP address, gateway, and other basic networking settings. You can only configure the Management interface settings; you cannot configure inside or outside interfaces, which you can later configure in the GUI.

Note You do not need to use this procedure for the Firepower 4100/9300, because you set the IP address manually when you deployed.

Note You cannot repeat the CLI setup script 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 *Cisco Secure Firewall Threat Defense Command Reference*.

Procedure

Step 1 Connect to the Firewall Threat Defense console port. See Logging Into the Command Line Interface (CLI), on page 15 for more information.

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

The default admin password is Admin123. On AWS, the default admin password for the Firewall Threat Defense Virtual is the AWS Instance ID, unless you define a default password with user data (**Advanced Details** > **User Data**) during the initial deployment.

Step 3 The first time you log into the Firewall Threat Defense, 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.

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**—If you set a manual IP address, enter either **data-interfaces** or the IP address of the gateway router. The **data-interfaces** setting sends outbound management traffic over the backplane to exit a data interface. This setting is useful if you do not have a separate Management network that can access the internet. Traffic originating on the Management interface includes license registration and database updates that require internet access. If you use **data-interfaces**, you can still use the Firewall Device Manager (or SSH) on the Management interface if you are directly-connected to the Management network, but for remote management for specific networks or hosts, you should add a static route using the **configure network static-routes** command. Note that the Firewall Device Manager management on data interfaces is not affected by this setting. If you use DHCP, the system uses the gateway provided by DHCP and uses the **data-interfaces** as a fallback method if DHCP doesn't provide a gateway.
- **If your networking information has changed, you will need to reconnect**—If you are connected with SSH to the default IP address 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.
- **Manage the device locally?**—Enter **yes** to use the Firewall Device Manager. A **no** answer means you intend to use the Firewall Management Center to manage the device.

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.
For HTTP Proxy configuration, run 'configure network http-proxy'

Manage the device locally? (yes/no) [yes]: yes

>
```

Step 4 Log into the Firewall Device Manager on the new Management IP address.

Complete the Initial Configuration Using the Setup Wizard

When you initially log into the Firewall Device Manager, you are taken through the device setup wizard to complete the initial system configuration.

If you plan to use the device in a high availability configuration, please read Prepare the Two Units for High Availability, on page 203.



Note The Firepower 4100/9300 and ISA 3000 do not support the setup wizard, so this procedure does not apply to these models. For the Firepower 4100/9300, all initial configuration is set when you deploy the logical device from the chassis. For the ISA 3000, a special default configuration is applied before shipping.

Before you begin

Ensure that you connect a data interface to your gateway device, for example, a cable modem or router. For edge deployments, this would be your Internet-facing gateway. For data center deployments, this would be a back-bone router. Use the default "outside" interface for your model (see Connect the Interfaces, on page 19 and Default Configuration Prior to Initial Setup, on page 8).

Then, connect your management computer to the "inside" interface for your hardware model. Alternatively, you can connect to the Management interface. For the Firewall Threat Defense Virtual, simply ensure that you have connectivity to the management IP address.

(Except for the Firewall Threat Defense Virtual, which requires connectivity to the internet from the management IP address.) The Management interface does not need to be connected to a network. By default, the system obtains system licensing and database and other updates through the data interfaces, typically the outside interface, that connect to the internet. If you instead want to use a separate management network, you can connect the Management interface to a network and configure a separate management gateway after you complete initial setup.

To change the Management interface network settings if you cannot access the default IP address, see (Optional) Change Management Network Settings at the CLI, on page 19.

Procedure

Step 1 Log into the Firewall Device Manager.

a) Assuming you did not go through initial configuration in the CLI, open the Firewall Device Manager at **https://***ip-address*, where the address is one of the following.

- If you are connected to the inside interface: **https://192.168.95.1**.
- (Firewall Threat Defense Virtual) If you are connected to the Management interface: **https://192.168.45.45**.
- (All other models) If you are connected to the Management interface: **https://***dhcp_client_ip*

b) Log in with the username **admin**. The default admin password is Admin123. On AWS, the default admin password for the Firewall Threat Defense Virtual is the AWS Instance ID, unless you define a default password with user data (**Advanced Details** > **User Data**) during the initial deployment..

Step 2 If this is the first time logging into the system, and you did not use the CLI setup wizard, you are prompted to read and accept the End User License Agreement and change the admin password.

You must complete these steps to continue.

Step 3 Configure the following options for the outside and management interfaces and click **Next**.

Caution
Your settings are deployed to the device when you click **Next**. The interface will be named "outside" and it will be added to the "outside_zone" security zone. Ensure that your settings are correct.

Outside Interface

- **Configure IPv4**—The IPv4 address for the outside interface. You can use DHCP or manually enter a static IP address, subnet mask, and gateway. You can also select **Off** to not configure an IPv4 address. Do not configure an IP address on the same subnet as the default inside address (see Default Configuration Prior to Initial Setup, on page 8), either statically or through DHCP. PPPoE is not supported on WAN interfaces, including interfaces that you use as VPN tunnel sources. Use DHCP or a static IPv4 address for WAN interfaces..
- **Configure IPv6**—The IPv6 address for the outside interface. You can use DHCP or manually enter a static IP address, prefix, and gateway. You can also select **Off** to not configure an IPv6 address.

Management Interface

- **DNS Servers**—The DNS server for the system's management address. Enter one or more addresses of DNS servers for name resolution. The default is the OpenDNS public DNS servers, or the DNS servers you obtain from the DHCP server. If you edit the fields and want to return to the default, click **Use OpenDNS** to reload the appropriate IP addresses into the fields. Your ISP might require that you use specific DNS servers. If after completing the wizard, you find that DNS resolution is not working, see Troubleshooting DNS for the Management Interface, on page 822.
- **Firewall Hostname**—The hostname for the system's management address.

Step 4 Configure the system time settings and click **Next**.

- **Time Zone**—Select the time zone for the system.
- **NTP Time Server**—Select whether to use the default NTP servers or to manually enter the addresses of your NTP servers. You can add multiple servers to provide backups.

Step 5 Configure the smart licenses for the system.

You must have a smart license account to obtain and apply the licenses that the system requires. Initially, you can use the 90-day evaluation license and set up smart licensing later.

To register the device now, select the option to register the device, click the link to log into your Smart Software Manager account, generate a new token, and copy the token into the edit box. You must also select your services region, and decide whether to send usage data to the Cisco Success Network. The on-screen text explains these settings in more detail.

If you do not want to register the device yet, select the evaluation mode option. The evaluation period last up to 90 days. To later register the device and obtain smart licenses, click **Device**, then click the link in the **Smart Licenses** group.

Step 6 Click **Finish**.

What to do next

- If you want to use features covered by optional licenses, such as category-based URL filtering, intrusion inspection, or malware prevention, enable the required licenses. See Enabling or Disabling Optional Licenses, on page 89.
- Connect the other data interfaces to distinct networks and configure the interfaces. For information on configuring interfaces, see How to Add a Subnet, on page 67 and Interfaces, on page 235.
- If you are managing the device through the inside interface, and you want to open CLI sessions through the inside interface, open the inside interface to SSH connections. See Configuring the Management Access List, on page 748.
- Go through the use cases to learn how to use the product. See Best Practices: Use Cases for Firewall Threat Defense, on page 39.

What to Do if You Do Not Obtain an IP Address for the Outside Interface

The default device configuration includes a static IPv4 address for the inside interface. You cannot change this address through the initial device setup wizard, although you can change it afterwards.

The default inside IP address might conflict with other networks attached to the device. This is especially true if you use DHCP on the outside interface to obtain an address from your Internet Service Provider (ISP). Some ISPs use the same subnet as the inside network as the address pool. Because you cannot have two data interfaces with addresses on the same subnet, conflicting addresses from the ISP cannot be configured on the outside interface.

If there is a conflict between the inside static IP address and the DHCP-provided address on the outside interface, the connection diagram should show the outside interface as administratively UP, but with no IPv4 address.

The setup wizard will complete successfully in this case, and all the default NAT, access, and other policies and settings will be configured. Simply follow the procedure below to eliminate the conflict.

Before you begin

Verify that you have a healthy connection to the ISP. Although a subnet conflict will prevent you from getting an address on the outside interface, you will also fail to get one if you simply do not have a link to the ISP.

Procedure

Step 1 Click **Device**, then click the link in the **Interfaces** summary.

Step 2 Mouse over the **Actions** column for the inside interface and click the edit icon ( ).

Step 3 On the **IPv4 Address** tab, enter a static address on a unique subnet, for example, 192.168.2.1/24 or 192.168.46.1/24. Note that the default management address is 192.168.45.45/24, so do not use that subnet.

You also have the option to use DHCP to obtain an address if you have a DHCP server already running on the inside network. However, you must first click **Delete** in the **DHCP SERVER IS DEFINED FOR THIS INTERFACE** group to remove the DHCP server from the interface.

Step 4 In the **DHCP SERVER IS DEFINED FOR THIS INTERFACE** area, click **Edit** and change the DHCP pool to a range on the new subnet, for example, 192.168.2.5-192.168.2.254.

Step 5 Click **OK** to save the interface changes.

Step 6 Click the **Deploy** button in the menu to deploy your changes.



Step 7 Click **Deploy Now**.

After deployment completes, the connection graphic should show that the outside interface now has an IP address. Use a client on the inside network to verify you have connectivity to the Internet or other upstream network.

Configuration Basics

The following topics explain the basic methods for configuring the device.

Configuring the Device

When you initially log into Firewall Device Manager, you are guided through a setup wizard to help you configure basic settings. Once you complete the wizard, use the following method to configure other features and to manage the device configuration.

If you have trouble distinguishing items visually, select a different color scheme in the user profile. Select **Profile** from the user icon drop-down menu in the upper right of the page.



Procedure

Step 1 Click **Device** to get to the **Device Summary**.

The dashboard shows a visual status for the device, including enabled interfaces and whether key settings are configured (colored green) or still need to be configured. For more information, see Viewing Interface and Management Status, on page 30.

Above the status image is a summary of the device model, software version, VDB (System and Vulnerability Database) version, and the last time intrusion rules were updated. This area also shows high availability status, including links to configure the feature; see High Availability (Failover), on page 191. It also shows cloud registration status, where you see the account to which the device is registered if you are using cloud management; see Configuring Cloud Services, on page 772.

Below the image are groups for the various features you can configure, with summaries of the configurations in each group, and actions you can take to manage the system configuration.

Step 2 Click the links in each group to configure the settings or perform the actions.

Following is a summary of the groups:

- **Interface**—You should have at least two data interfaces configured in addition to the management interface. See Interfaces, on page 235.
- **Routing**—The routing configuration. You must define a default route. Other routes might be necessary depending on your configuration. See Routing, on page 301.
- **Updates**—Geolocation, intrusion rule, and vulnerability database updates, and system software upgrades. Set up a regular update schedule to ensure that you have the latest database updates if you use those features. You can also go to this page if you need to download an update before the regularly schedule update occurs. See Updating System Databases and Feeds, on page 787.
- **System Settings**—This group includes a variety of settings. Some are basic settings that you would configure when you initially set up the device and then rarely change. See System Settings, on page 747.
- **Smart License**—Shows the current state of the system licenses. You must install the appropriate licenses to use the system. Some features require additional licenses. See Licensing the System, on page 81.
- **Backup and Restore**—Back up the system configuration or restore a previous backup. See Backing Up and Restoring the System, on page 796.
- **Troubleshoot**—Generate a troubleshooting file at the request of the Cisco Technical Assistance Center. See Creating a Troubleshooting File, on page 827.
- **Site-to-Site VPN**—The site-to-site virtual private network (VPN) connections between this device and remote devices. See Managing Site-to-Site VPNs, on page 636.
- **Remote Access VPN**—The remote access virtual private network (VPN) configuration that allows outside clients to connect to your inside network. See Configuring Remote Access VPN, on page 678.
- **Advanced Configuration**—Use FlexConfig and Smart CLI to configure features that you otherwise cannot configure using Firewall Device Manager. See Advanced Configuration, on page 831.
- **Device Administration**—View the audit log or export a copy of the configuration. See Auditing and Change Management, on page 802.

Step 3 Click the **Deploy** button in the menu to deploy your changes.



Changes are not active on the device until you deploy them. See Deploying Your Changes, on page 27.

What to do next

Click **Policies** in the main menu and configure the security policy for the system. You can also click **Objects** to configure the objects needed in those policies.

Configuring Security Policies

Use the security policies to implement your organization's acceptable use policy and to protect your network from intrusions and other threats.

Procedure

Step 1 Click **Policies**.

The Security Policies page shows the general flow of a connection through the system, and the order in which security policies are applied.

Step 2 Click the name of a policy and configure it.

You might not need to configure each policy type, although you must always have an access control policy. Following is a summary of the policies:

- **SSL Decryption**—If you want to inspect encrypted connections (such as HTTPS) for intrusions, malware, and so forth, you must decrypt the connections. Use the SSL decryption policy to determine which connections need to be decrypted. The system re-encrypts the connection after inspecting it. See Configuring SSL Decryption Policies, on page 435.
- **Identity**—If you want to correlate network activity to individual users, or control network access based on user or user group membership, use the identity policy to determine the user associated with a given source IP address. See Configuring Identity Policies, on page 457.
- **Security Intelligence**—Use the Security Intelligence policy to quickly drop connections from or to selected IP addresses or URLs. By blocking known bad sites, you do not need to account for them in your access control policy. Cisco provides regularly updated feeds of known bad addresses and URLs so that the Security Intelligence block lists update dynamically. Using feeds, you do not need to edit the policy to add or remove items in the block lists. See Configuring Security Intelligence, on page 469.
- **NAT** (Network Address Translation)—Use the NAT policy to convert internal IP addresses to externally routeable addresses. See Configure NAT, on page 552.
- **Access Control**—Use the access control policy to determine which connections are allowed on the network. You can filter by security zone, IP address, protocol, port, application, URL, user or user group. You also apply intrusion and file (malware) policies using access control rules. Use this policy to implement URL filtering. See Configuring the Access Control Policy, on page 489.
- **Intrusion**—Use the intrusion policies to inspect for known threats. Although you apply intrusion policies using access control rules, you can edit the intrusion policies to selectively enable or disable specific intrusion rules. See Intrusion Policies, on page 509.

Step 3 Click the **Deploy** button in the menu to deploy your changes.



Changes are not active on the device until you deploy them. See Deploying Your Changes, on page 27.

Searching for Rules or Objects

You can use full-text search on lists of policy rules or objects to help you find the item you want to edit. This is especially helpful when dealing with policies that have hundreds of rules, or long object lists.

The method for using search on rules and objects is the same for any type of policy (except the intrusion policy) or object: in the **Search** field, enter a string to find, and press Enter.

This string can exist in any part of the rule or object, and it can be a partial string. You can use the asterisk * as a wildcard that matches zero or more characters. Do not include the following characters, they are not supported as part of the search string: ?~!{}<>:%. The following characters are ignored: ;#&.

The string can appear within an object in the group. For example, you can enter an IP address and find the network objects or groups that specify that address.

When done, click the **x** on the right side of the search box to clear the filter.

Deploying Your Changes

When you update a policy or setting, the change is not immediately applied to the device. There is a two step process for making configuration changes:

1. Make your changes.
2. Deploy your changes.

This process gives you the opportunity to make a group of related changes without forcing you to run a device in a "partially configured" manner. In most cases, the deployment includes just your changes. However, if necessary, the system will reapply the entire configuration, which might be disruptive to your network. In addition, some changes require inspection engines to restart, with traffic dropping during the restart. Thus, consider deploying changes when potential disruptions will have the least impact.

Note If the deployment job fails, the system must roll back any partial changes to the previous configuration. Rollback includes clearing the data plane configuration and redeploying the previous version. This will disrupt traffic until the rollback completes.

After you complete the changes you want to make, use the following procedure to deploy them to the device.

Caution The Firewall Threat Defense device drops traffic when the inspection engines are busy because of a software resource issue, or down because a configuration requires the engines to restart during configuration deployment. For detailed information on changes that require a restart, see Configuration Changes that Restart Inspection Engines, on page 28.

Procedure

Step 1 Click the **Deploy Changes** icon in the upper right of the web page.

The icon is highlighted with a dot when there are undeployed changes.



The Pending Changes window shows a comparison of the deployed version of the configuration with the pending changes. These changes are color-coded to indicate removed, added, or edited elements. See the legend in the window for an explanation of the colors.

If the deployment requires that inspection engines be restarted, the page includes a message that provides detail on what changed that requires a restart. If momentary traffic loss at this time would be unacceptable, close the dialog box and wait until a better time to deploy changes.

If the icon is not highlighted, you can still click it to see the date and time of the last successful deployment job. There is also a link to show you the deployment history, which takes you to the audit page filtered to show deployment jobs only.



Step 2 If you are satisfied with the changes, you can click **Deploy Now** to start the job immediately.

The window will show that the deployment is in progress. You can close the window, or wait for deployment to complete. If you close the window while deployment is in progress, the job does not stop. You can see results in the task list or audit log. If you leave the window open, click the **Deployment History** link to view the results.

Optionally, you can do the following:

- **Name the Job**—To name the deployment job, click the drop-down arrow on the **Deploy Now** button and select **Name the Deployment Job**. Enter a name, then click **Deploy**. The name will appear in the audit and deployment history as part of the job, which might make it easier for you to find the job.

  For example, if you name a job "DMZ Interface Configuration," a successful deployment will be named "Deployment Completed: DMZ Interface Configuration." In addition, the name is used as the Event Name in Task Started and Task Completed events related to the deployment job.

- **Discard Changes**—To discard all pending changes, click **More Options** > **Discard All**. You are prompted for confirmation.

- **Copy Changes**—To copy the list of changes to the clipboard, click **More Options** > **Copy to Clipboard**. This option works only if there are fewer than 500 changes.

- **Download Changes**—To download the list of changes as a file, click **More Options** > **Download as Text**. You are prompted to save the file to your workstation. The file is in YAML format. You can view it in a text editor if you do not have an editor that specifically supports YAML format.

Configuration Changes that Restart Inspection Engines

Any of the following configurations or actions restart inspection engines when you deploy configuration changes.



Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations requires inspection engines to restart, which interrupts traffic inspection and drops traffic.

Deployment

Some changes require that inspection engines be restarted, which will result in momentary traffic loss. Following are the changes that require inspection engine restart:

- SSL decryption policy is enabled or disabled.
- The MTU changed on one or more physical interfaces (but not subinterfaces).
- You add or remove a file policy on an access control rule.
- The VDB was updated.
- Creating or breaking the high availability configuration.

In addition, some packets might be dropped during deployment if the Snort process is busy, with the total CPU utilization exceeding 60%. You can check the current CPU utilization for Snort using the **show asp inspect-dp snort** command.

System Database Updates

If you download an update to the Rules database or VDB, you must deploy the update for it to become active. This deployment might restart inspection engines. When you manually download an update, or schedule an update, you can indicate whether the system should automatically deploy changes after the download is complete. If you do not have the system automatically deploy the update, the update is applied the next time you deploy changes, at which time inspection engines might restart.

System Updates

Installing a system update or patch that does not reboot the system and includes a binary change requires inspection engines to restart. Binary changes can include changes to inspection engines, a preprocessor, the vulnerability database (VDB), or a shared object rule. Note also that a patch that does not include a binary change can sometimes require a Snort restart.

Configuration Changes that Force a Full Deployment

In most cases, the deployment includes just your changes. However, if necessary, the system will reapply the entire configuration, which might be disruptive to your network. Following are some changes that force a full deployment.

- The Security Intelligence or Identity policies are initially enabled.
- Both the Security Intelligence and Identity policies are disabled.
- Creating an EtherChannel when you reuse data.
- Deleting an EtherChannel.
- Modifying the member interface associations of an EtherChannel.
- Deleting any interface that is used in the configuration. For example, deleting a subinterface that is part of a security zone used by an access control rule.
- Changing a FlexConfig object that is part of the FlexConfig policy, or deleting an object from the policy, when that object does not include negate lines. Omitting negate lines forces the system to full deploy, because there is no specific way to remove the configuration produced by the FlexConfig object. You can avoid this problem by always including the appropriate negate lines in each FlexConfig object.

Viewing Interface and Management Status

The Device Summary includes a graphical view of your device and select settings for the management address. To open the Device Summary, click **Device**.

Elements on this graphic change color based on the status of the element. Mousing over elements sometimes provides additional information. Use this graphic to monitor the following items.

Note The interface portion of the graphic, including interface status information, is also available on the **Interfaces** page and the **Monitoring** > **System** dashboard.

Interface Status

Mouse over a port to see its IP addresses, and enabled and link statuses. The IP addresses can be statically assigned or obtained using DHCP. Mousing over a Bridge Virtual Interface (BVI) also shows the list of member interfaces.

Interface ports use the following color coding:

- Green—The interface is configured, enabled, and the link is up.
- Gray—The interface is not enabled.
- Orange/Red—The interface is configured and enabled, but the link is down. If the interface is wired, this is an error condition that needs correction. If the interface is not wired, this is the expected status.

Inside, Outside Network Connections

The graphic indicates which port is connected to the outside (or upstream) and inside networks, under the following conditions.

- Inside Network—The port for the inside network is shown for the interface named “inside” only. If there are additional inside networks, they are not shown. If you do not name any interface “inside,” no port is marked as the inside port.
- Outside Network—The port for the outside network is shown for the interface named “outside” only. As with the inside network, this name is required, or no port is marked as the outside port.

Management Setting Status

The graphic shows whether the gateway, DNS servers, NTP servers, and Smart Licensing are configured for the management address, and whether those settings are functioning correctly.

Green indicates that the feature is configured and functioning correctly, gray indicates that it is not configured or not functioning correctly. For example, the DNS box is gray if the servers cannot be reached. Mouse over the elements to see more information.

If you find problems, correct them as follows:

- Management port and gateway—Select **System Settings** > **Management Interface**.
- DNS servers—Select **System Settings** > **DNS Server**.
- NTP servers—Select **System Settings** > **NTP**. Also see Troubleshooting NTP, on page 821.

- Smart License—Click the **View Configuration** link in the Smart License group.

Viewing System Task Status

System tasks include actions that occur without your direct involvement, such as retrieving and applying various database updates. You can view a list of these tasks and their status to verify that these system tasks are completing successfully.

The task list shows consolidated status for system tasks and deployment jobs. The audit log contains more detailed information, and is available under **Device** > **Device Administration** > **Audit Log**. For example, the audit log shows separate events for task start and task end, whereas the task list merges those events into a single entry. In addition, the audit log entry for a deployment includes detailed information about the deployed changes.

Procedure

Step 1 Click the **Task List** button in the main menu.



The task list opens, displaying the status and details of system tasks.

Step 2 Evaluate the task status.

If you find a persistent problem, you might need to fix the device configuration. For example, a persistent failure to obtain database updates could indicate that there is no path to the Internet for the device's management IP address. You might need to contact the Cisco Technical Assistance Center (TAC) for some issues as indicted in the task descriptions.

You can do the following with the task list:

- Click the **Success** or **Failures** buttons to filter the list based on these statuses.
- Click the delete icon () for a task to remove it from the list.
- Click **Remove All Completed Tasks** to empty the list of all tasks that are not in progress.

Using the CLI Console to Monitor and Test the Configuration

Firewall Threat Defense devices include a command line interface (CLI) that you can use for monitoring and troubleshooting. Although you can open an SSH session to get access to all of the system commands, you can also open a CLI Console in the Firewall Device Manager to use read-only commands, such as the various **show** commands and **ping**, **traceroute**, and **packet-tracer**. If you have Administrator privileges, you can also enter the **failover**, **reboot**, and **shutdown** commands.

You can keep the CLI Console open as you move from page to page, configure, and deploy features. For example, after deploying a new static route, you could use **ping** in the CLI Console to verify that the target network is reachable.

The CLI Console uses the base Firewall Threat Defense CLI. You cannot enter the diagnostic CLI, expert mode, or FXOS CLI (on models that use FXOS) using the CLI Console. Use SSH if you need to enter those other CLI modes.

For detailed information on commands, see Cisco Firepower Threat Defense Command Reference, https://www.cisco.com/c/en/us/td/docs/security/firepower/command_ref/b_Command_Reference_for_Firepower_Threat_Defense.html.

Notes:

- Although **ping** is supported in CLI Console, the **ping system** command is not supported.
- The system can process at most 2 concurrent commands. Thus, if another user is issuing commands (for example, using the REST API), you might need to wait for other commands to complete before entering a command. If this is a persistent problem, use an SSH session instead of the CLI Console.
- Commands return information based on the deployed configuration. If you make a configuration change in the Firewall Device Manager, but do not deploy it, you will not see the results of your change in the command output. For example, if you create a new static route but do not deploy it, that route will not appear in **show route** output.

Procedure

Step 1 Click the **CLI Console** button in the upper right of the web page.



Step 2 Type the commands at the prompt and press **Enter**.

Some commands take longer to produce output than others, please be patient. If you get a message that the command execution timed out, please try again. You will also get a time out error if you enter a command that requires interactive responses, such as **show perfstats**. If the problem persists, you might need to use an SSH client instead of the CLI Console.

Following are some tips on how to use the window.

- Press the **Tab** key to automatically complete a command after partially typing it. Also, Tab will list out the parameters available at that point in the command. Tab works down to three levels of keyword. After three levels, you need to use the command reference for more information.
- You can stop command execution by pressing Ctrl+C.
- To move the window, click and hold anywhere in the header, then drag the window to the desired location.
- Click the **Expand** () or **Collapse** () button to make the window bigger or smaller.
- Click the **Undock Into Separate Window** () button to detach the window from the web page into its own browser window. To dock it again, click the **Dock to Main Window** () button.
- Click and drag to highlight text, then press Ctrl+C to copy output to the clipboard.
- Click the **Clear CLI** () button to erase all output.

- Click the **Copy Last Output** () button to copy the output from the last command you entered to the clipboard.

Step 3 When you are finished, simply close the console window. Do not use the **exit** command.

Although the credentials you use to log into the Firewall Device Manager validate your access to the CLI, you are never actually logged into the CLI when using the console.

Using Firewall Device Manager and the REST API Together

When you set up the device in local management mode, you can configure the device using the Firewall Device Manager and the Firewall Threat Defense REST API. In fact, the Firewall Device Manager uses the REST API to configure the device.

However, please understand that the REST API can provide additional features than the ones available through the Firewall Device Manager. Thus, for any given feature, you might be able to configure settings using the REST API that cannot appear when you view the configuration through the Firewall Device Manager.

If you do configure a feature setting that is available in the REST API but not in the Firewall Device Manager, and then make a change to the overall feature (such as remote access VPN) using the Firewall Device Manager, that setting might be undone. Whether an API-only setting is preserved can vary, and in many cases, API changes to settings not available in the Firewall Device Manager are preserved through the Firewall Device Manager edits. For any given feature, you should verify whether your changes are preserved.

In general, you should avoid using both the Firewall Device Manager and the REST API simultaneously for any given feature. Instead, choose one method or the other, feature by feature, for configuring the device.

You can view, and try out, the API methods using API Explorer. Click the more options button (⋮) and choose **API Explorer**.

Communication Port and Internet Access Requirements

The following topics describe the ports that need to be open on the device, and the internet sites to which access is required, to operate all features available on the device. If you are operating in a limited or air-gapped manner, some or all of these ports and internet sites can be blocked to meet your requirements. Otherwise, ensure these ports and sites are open and accessible.

Communication Ports Used by the Device

The following tables list the ports that need to be open for various services. Use access control rules to ensure access to these ports are open on the interfaces that face the services, or that face users who should be able to gain access to the device using the associated protocol. If you do not use a feature, you can leave those ports closed.

Table 1: Inbound Ports

Inbound Port	Protocol/Feature	Details
TCP/22	SSH	Secure remote connections to the appliance command line interface. **Note** If you use the **copy** command, or another command that performs external communication, you need to open that outbound port. For example, TCP/20, 21 if you intend to use FTP.
UDP/161	SNMP	Allow access to MIBs via SNMP polling.
TCP/443	HTTPS	Used for the following services: • Management connections to Firewall Device Manager. • Firewall Threat Defense REST API. • Remote access VPN SSL connections. If you configure a custom port for RA VPN, open that port.
TCP/885	Captive portal	Communicate with a captive portal identity source. This is the default port. If you configure a different port, open your custom port. For more information, see Configure Identity Policy Settings, on page 457.
TCP/8989	Cisco Support Diagnostics	Accepts authorized requests. Also initiates connections on this port.

Table 2: Outbound Ports

Outbound Port	Protocol/Feature	Details
UDP/53 (If used.) TCP/53	DNS	DNS. Normally, UDP/53 is used for DNS. But if you use DNS over TCP, open TCP/53 as well.
UDP/67 UDP/68	DHCP	DHCP.
UDP/123	NTP	Synchronize time.
UDP/162	SNMP	Send SNMP alerts to a remote trap server.
TCP/389 TCP/636	LDAP, LDAPS	Communicate with an LDAP server for external authentication. If you configure custom ports, open those ports. See Configuring AD Identity Realms, on page 161.
TCP/443	HTTPS	Send and receive data from the internet, such as downloading database updates..
UDP/514	Syslog	Send system log messages to a remote syslog server.

Outbound Port	Protocol/Feature	Details
UDP/1812 UDP/1813	RADIUS	Communicate with a RADIUS server for external authentication and accounting. If you configure custom ports, open those ports. See Configure RADIUS Servers, on page 165.
UDP/8514	Secure Network Analytics Manager	Send syslog messages to the cloud.
TCP/8989	Cisco Support Diagnostics	Transmits usage information and statistics. Also accepts connections on this port.

Internet Resources Accessed by the Device

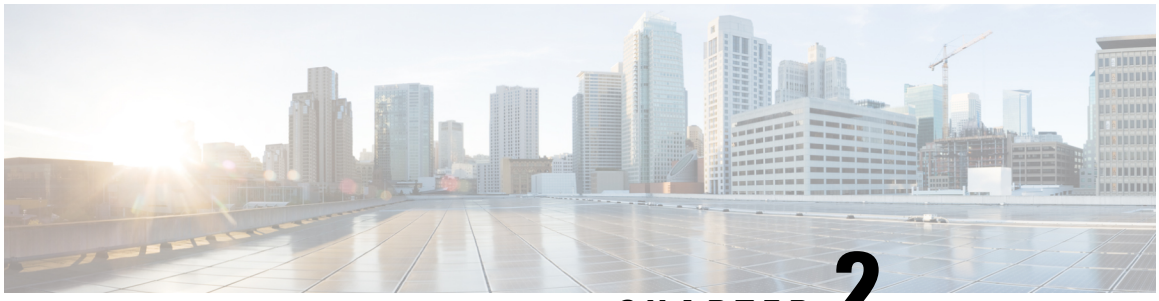
The following features must have access to the associated internet resources to operate correctly. The device uses ports TCP/80 and TCP/443 as needed.

Table 3: Internet Resources Accessed by the Device

Feature	Reason	High Availability	Resource
CA certificate bundles	Queries for new CA certificates at a daily system-defined time. The local CA bundle contains certificates to access several Cisco services. This feature is configured in the CLI using the **configure cert-update auto-update** command. Available in versions 7.0(5), 7.1(0.3), 7.2(4), 7.3 and higher.	Each peer downloads its own certificates.	cisco.com/security/pki
Malware Defense 1	AMP Cloud lookups.	Both peers perform lookups.	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. Query for dynamic analysis results.	Both peers submit or query for dynamic analysis reports.	fmc.api.threatgrid.com fmc.api.threatgrid.eu fmc.api.threatgrid.ca fmc.api.threatgrid.com.au fmc.api.threatgrid.in

Feature	Reason	High Availability	Resource
Security intelligence	Download security intelligence feeds.	Active peer downloads, syncs to standby.	intelligence.sourcefire.com
URL filtering	Download URL category and reputation data. Manually query (look up) URL category and reputation data. Query for uncategorized URLs.	Active peer downloads, syncs to standby.	URLs: • regsvc.sco.cisco.com • est.sco.cisco.com • updates-talos.sco.cisco.com • updates-dyn-talos.sco.cisco.com • updates.ironport.com IPv4 blocks: • 146.112.62.0/24 • 146.112.63.0/24 • 146.112.255.0/24 • 146.112.59.0/24 IPv6 blocks: • 2a04:e4c7:ffff::/48 • 2a04:e4c7:fffe::/48
Cisco Smart Software Manager	Communicate with the Smart Software Manager.	Active peer communicates.	www.cisco.com
Cisco Success Network	Transmit usage information and statistics.	Active peer communicates.	api-sse.cisco.com:8989 dex.sse.itd.cisco.com dex.eu.sse.itd.cisco.com
Cisco Support Diagnostics	Accepts authorized requests and transmits usage information and statistics.	Active peer communicates.	api-sse.cisco.com:8989
General cloud services	—	—	api-sse.cisco.com
Cisco XDR integration	Configure devices to send events to the Security Cloud Control.	Active peer communicates.	Cisco Secure Firewall Threat Defense and Cisco XDR Integration Guide

Feature	Reason	High Availability	Resource
Time synchronization	Synchronize time in your deployment. Not supported with a proxy server.	Both peers communicate with the NTP server.	User configured. Default servers: • 0.sourcefire.pool.ntp.org • 1.sourcefire.pool.ntp.org • 2.sourcefire.pool.ntp.org
Intrusion rules	Download intrusion rules (SRU/LSP).	Active peer downloads, syncs to standby.	talosintelligence.com
Vulnerability database	Download VDB updates.	Active peer downloads, syncs to standby.	support.sourcefire.com
Geolocation database	Download GeoDB updates.	Active peer downloads, syncs to standby.	support.sourcefire.com



CHAPTER 2

Best Practices: Use Cases for Firewall Threat Defense

The following topics explain some common tasks you might want to accomplish with Firewall Threat Defense using the Firewall Device Manager. These use cases assume that you completed the device configuration wizard and that you retained this initial configuration. Even if you modified the initial configuration, you should be able to use these examples to understand how to use the product.

- How to Configure the Device in Firewall Device Manager, on page 39
- How to Gain Insight Into Your Network Traffic, on page 44
- How to Block Threats, on page 51
- How to Block Malware, on page 56
- How to Implement an Acceptable Use Policy (URL Filtering), on page 59
- How to Control Application Usage, on page 64
- How to Add a Subnet, on page 67
- How to Passively Monitor the Traffic on a Network, on page 72
- More Examples, on page 77

How to Configure the Device in Firewall Device Manager

After you complete the setup wizard, you should have a functioning device with a few basic policies in place:

- An outside and an inside interface. No other data interfaces are configured.
- (Firepower 4100/9300) No data interfaces are pre-configured.
- (ISA 3000) A bridge group contains 2 inside interfaces and 2 outside interfaces. You need to manually set the IP address of BVI1 to complete your setup.
- (Except for the Firepower 4100/9300) Security zones for the inside and outside interfaces.
- (Except for the Firepower 4100/9300) An access rule trusting all inside to outside traffic. For the ISA 3000, there are access rules that allow all traffic from inside to outside, and from outside to inside.
- (Except for the Firepower 4100/9300 and ISA 3000) An interface NAT rule that translates all inside to outside traffic to unique ports on the IP address of the outside interface.
- (Except for the Firepower 4100/9300 and ISA 3000) A DHCP server running on the inside interface.

The following steps provide an overview of additional features you might want to configure. Please click the help button (**?**) on a page to get detailed information about each step.

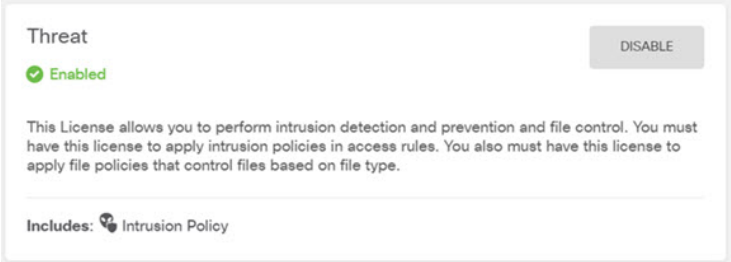
Procedure

Step 1 Choose **Device**, then click **View Configuration** in the **Smart License** group.

Click **Enable** for each of the optional licenses you want to use: Threat, Malware, URL. If you registered the device during setup, you can also enable the RA VPN license desired. Read the explanation of each license if you are unsure of whether you need it.

If you have not registered, you can do so from this page. Click **Register Device** and follow the instructions. Please register before the evaluation license expires.

For example, an enabled ThreatThreat license should look like the following:



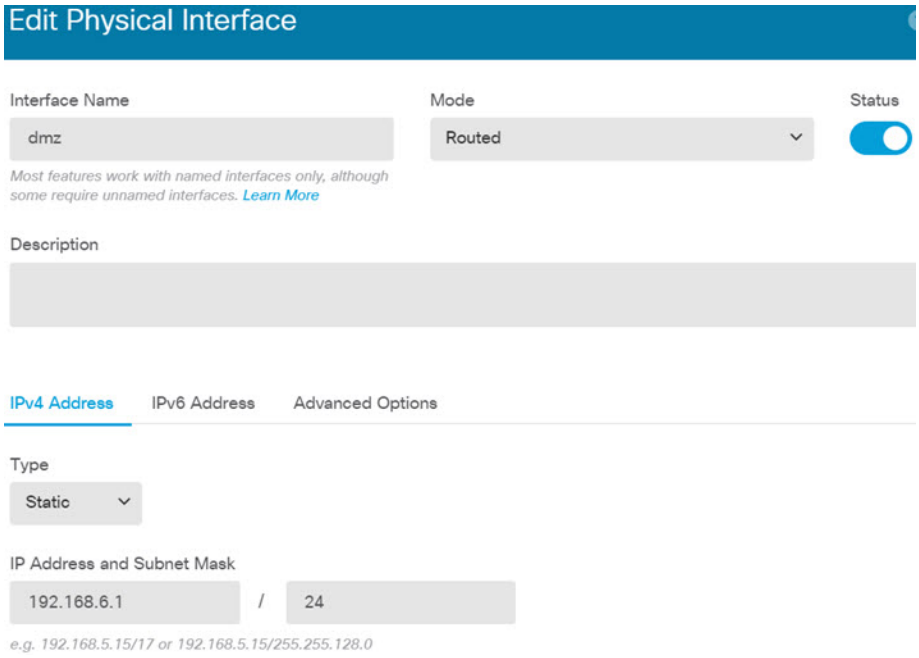
Step 2 If you wired other interfaces, choose **Device**, then click the link in the **Interfaces** summary, and then click the interfaces type to view the list of interfaces.

- For the Firepower 4100/9300, no data interfaces are pre-configured with names, IP addresses, or security zones, so you need to enable and configure any interfaces that you want to use.
- Because the ISA 3000 comes pre-configured with a bridge group containing all data interfaces, there is no need to configure these interfaces. However, you must manually configure an IP address for the BVI. If you want to break apart the bridge group, you can edit it to remove the interfaces you want to treat separately. Then you can configure those interfaces as hosting separate networks.

  For other models, you can create a bridge group for the other interfaces, or configure separate networks, or some combination of both.
- For the Firepower 1010, all interfaces except for Ethernet1/1 (outside) are access mode switch ports assigned to VLAN1 (inside). You can change switch ports to be firewall ports; add new VLAN interfaces and assign switch ports to them; or configure trunk mode switch ports.

Click the edit icon ( ) for each interface to define the IP address and other settings.

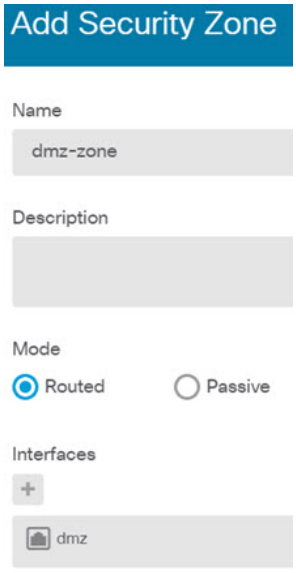
The following example configures an interface to be used as a “demilitarized zone” (DMZ), where you place publically-accessible assets such as your web server. Click **Save** when you are finished.



Step 3 If you configured new interfaces, choose **Objects**, then select **Security Zones** from the table of contents.

Edit or create new zones as appropriate. Each interface must belong to a zone, because you configure policies based on security zones, not interfaces. You cannot put the interfaces in zones when configuring them, so you must always edit the zone objects after creating new interfaces or changing the purpose of existing interfaces.

The following example shows how to create a new dmz-zone for the dmz interface.

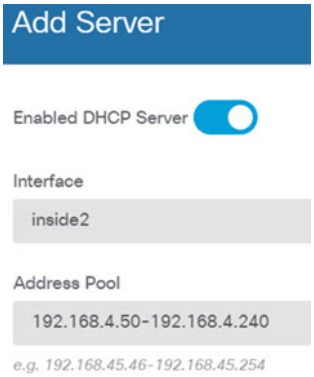


Step 4 If you want internal clients to use DHCP to obtain an IP address from the device, choose **Device**, then **System Settings** > **DHCP Server**. Select the **DHCP Servers** tab.

There is already a DHCP server configured for the inside interface, but you can edit the address pool or even delete it. If you configured other inside interfaces, it is very typical to set up a DHCP server on those interfaces. Click + to configure the server and address pool for each inside interface.

You can also fine-tune the WINS and DNS list supplied to clients on the **Configuration** tab.

The following example shows how to set up a DHCP server on the inside2 interface with the address pool 192.168.4.50-192.168.4.240.



Step 5 Choose **Device**, then click **View Configuration** in the **Routing** group and configure a default route.

The default route normally points to the upstream or ISP router that resides off the outside interface. A default IPv4 route is for any-ipv4 (0.0.0.0/0), whereas a default IPv6 route is for any-ipv6 (::0/0). Create routes for each IP version you use. If you use DHCP to obtain an address for the outside interface, you might already have the default routes that you need.

The routes you define on this page are for the data interfaces only. They do not impact the management interface. Set the management gateway on **System Settings** > **Management Interface**.

The following example shows a default route for IPv4. In this example, isp-gateway is a network object that identifies the IP address of the ISP gateway (you must obtain the address from your ISP). You can create this object by clicking **Create New Network** at the bottom of the **Gateway** drop-down list.



Step 6 Choose **Policies** and configure the security policies for the network.

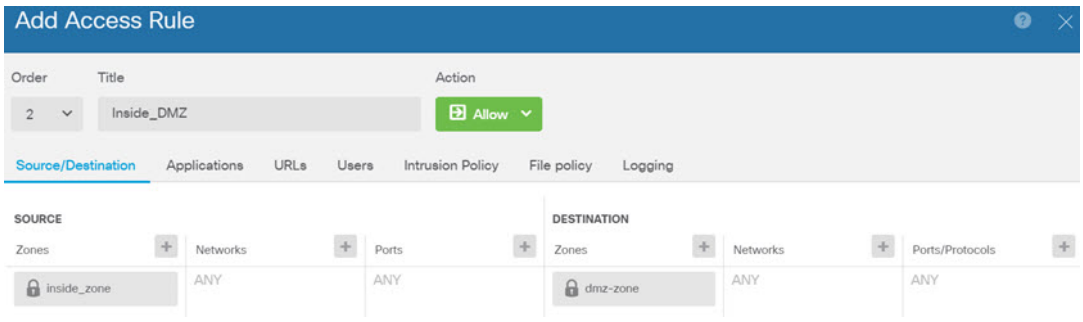
The device setup wizard enables traffic flow between the inside-zone and outside-zone, and interface NAT for all interfaces when going to the outside interface. Even if you configure new interfaces, if you add them to the inside-zone object, the access control rule automatically applies to them.

However, if you have multiple inside interfaces, you need an access control rule to allow traffic flow from inside-zone to inside-zone. If you add other security zones, you need rules to allow traffic to and from those zones. These would be your minimum changes.

In addition, you can configure other policies to provide additional services, and fine-tune NAT and access rules to get the results that your organization requires. You can configure the following policies:

- **SSL Decryption**—If you want to inspect encrypted connections (such as HTTPS) for intrusions, malware, and so forth, you must decrypt the connections. Use the SSL decryption policy to determine which connections need to be decrypted. The system re-encrypts the connection after inspecting it.
- **Identity**—If you want to correlate network activity to individual users, or control network access based on user or user group membership, use the identity policy to determine the user associated with a given source IP address.
- **Security Intelligence**—Use the Security Intelligence policy to quickly drop connections from or to selected IP addresses or URLs. By blocking known bad sites, you do not need to account for them in your access control policy. Cisco provides regularly updated feeds of known bad addresses and URLs so that the Security Intelligence block lists update dynamically. Using feeds, you do not need to edit the policy to add or remove items in the block lists.
- **NAT** (Network Address Translation)—Use the NAT policy to convert internal IP addresses to externally routeable addresses.
- **Access Control**—Use the access control policy to determine which connections are allowed on the network. You can filter by security zone, IP address, protocol, port, application, URL, user or user group. You also apply intrusion and file (malware) policies using access control rules. Use this policy to implement URL filtering.
- **Intrusion**—Use the intrusion policies to inspect for known threats. Although you apply intrusion policies using access control rules, you can edit the intrusion policies to selectively enable or disable specific intrusion rules.

The following example shows how to allow traffic between the inside-zone and dmz-zone in the access control policy. In this example, no options are set on any of the other tabs except for **Logging**, where **At End of Connection** is selected.



Step 7 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.



b) Click the **Deploy Now** button.

You can wait until deployment completes, or click **OK** and check the task list or deployment history later.

How to Gain Insight Into Your Network Traffic

After completing initial device setup, you have an access control policy that allows all inside traffic access to the Internet or other upstream network, and a default action to block all other traffic. Before you create additional access control rules, you might find it beneficial to gain insight into the traffic that is actually occurring on your network.

You can use the monitoring capabilities of the Firewall Device Manager to analyze network traffic. Firewall Device Manager reporting helps you answer the following questions:

- What is my network being used for?
- Who is using the network the most?
- Where are my users going?
- What devices are they using?
- What access control rules (policies) are being hit the most?

The initial access rule can provide some insight into traffic, including policies, destinations, and security zones. But to obtain user information, you need to configure an identity policy that requires users to authenticate (identify) themselves. To obtain information on applications used on the network, you need to make some additional adjustments.

The following procedure explains how to set up the Firewall Threat Defense device to monitor traffic and provides an overview of the end-to-end process of configuring and monitoring policies.



Note This procedure does not provide insight into the web site categories and reputations of sites visited by users, so you cannot see meaningful information in the URL categories dashboard. You must implement category-based URL filtering, and enable the URL license, to obtain category and reputation data. If you just want to obtain this information, you can add a new access control rule that allows access to an acceptable category, such as Finance, and make it the first rule in the access control policy. For details on implementing URL filtering, see How to Implement an Acceptable Use Policy (URL Filtering), on page 59.

Procedure

Step 1 To gain insight into user behavior, you need to configure an identity policy to ensure that the user associated with a connection is identified.

By enabling the identity policy, you can collect information about who is using the network, and what resources they are using. This information is available in the User monitoring dashboard. User information is also available for connection events shown in Event Viewer.

In this example, we will implement active authentication to acquire user identity. With active authentication, the device prompts the user for username and password. Users are authenticated only when they use a web browser for HTTP connections.

If a user fails to authenticate, the user is not prevented from making web connections. This just means that you do not have user identity information for the connections. If you want, you can create an access control rule to drop traffic for Failed Authentication users.

a) Click **Policies** in the main menu, then click **Identity**.

   The identity policy is initially disabled. When using active authentication, the identity policy uses your Active Directory server to authenticate users and associate them with the IP address of the workstation they are using. Subsequently, the system will identify traffic for that IP address as being the user's traffic.

b) Click **Enable Identity Policy**.

c) Click the **Create Identity Rule** button, or the **+** button, to create the rule to require active authentication.

   In this example, we will assume you want to require authentication for everyone.

d) Enter a **Name** for the rule, which can be anything you choose, for example, Require_Authentication.

e) On the **Source/Destination** tab, leave the defaults, which apply to Any criteria.

   You can constrain the policy as you see fit to a more limited set of traffic. However, active authentication will only be attempted for HTTP traffic, so it does not matter that non-HTTP traffic matches the source/destination criteria. For more details about identity policy properties, see Configure Identity Rules, on page 459

f) For **Action**, select **Active Auth**.

   Assuming you have not configured the identity policy settings, the Identity Policy Configuration dialog box will open because there are some undefined settings.

g) Configure the Captive Portal and SSL Decryption settings that are required for active authentication.

   When an identity rule requires active authentication for a user, the user is redirected to the captive portal port and then they are prompted to authenticate. Captive portal requires SSL decryption rules, which the system will generate automatically, but you must select the certificate to use for the SSL decryption rules.

   - **Server Certificate**—Select the internal certificate to present to users during active authentication. You can select the predefined self-signed DefaultInternalCertificate, or you can click **Create New Internal Certificate** and upload a certificate that your browsers already trust.

     Users will have to accept the certificate if you do not upload a certificate that their browsers already trust.

   - **Redirect to Host Name**—Select the network object that defines the fully-qualified host name of the interface that should be used as the captive portal for active authentication requests. Click **Create New Network** if the object does not exist.

     The FQDN must resolve to the IP address of one of the interfaces on the device. By using an FQDN, you can assign a certificate for active authentication that the client will recognize, thus avoiding the untrusted certificate warning users get when being redirected to an IP address. The certificate can specify the FQDN, a wildcard FQDN, or multiple FQDNs in the Subject Alternate Names (SAN) in the certificate.

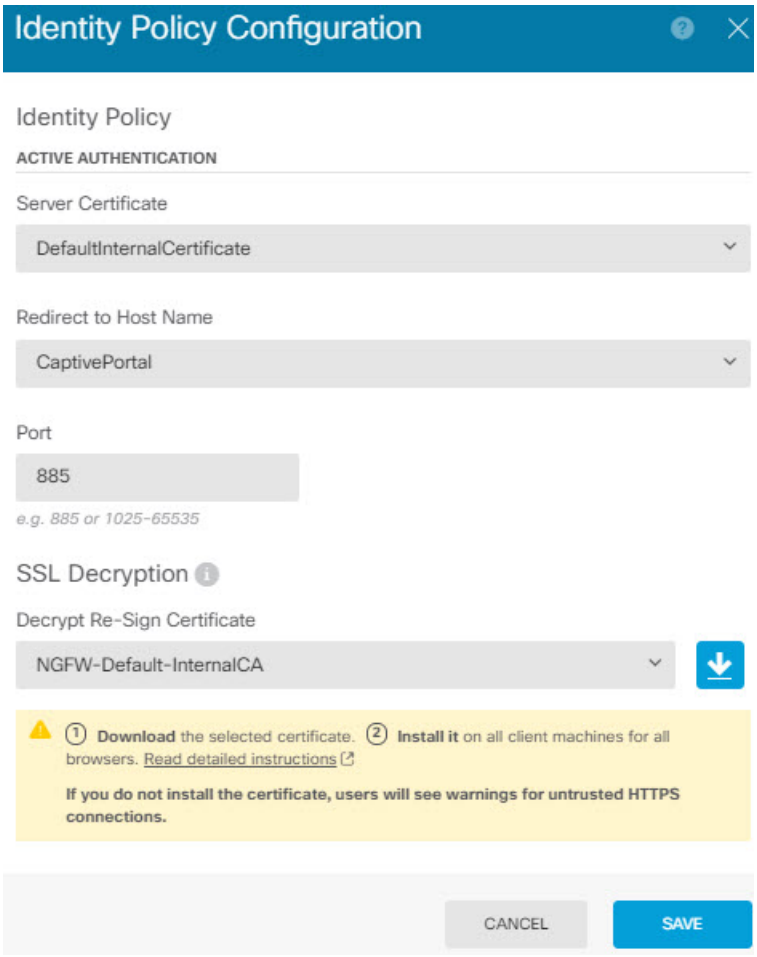
If an identity rule requires active authentication for a user, but you do not specify a redirect FQDN, the user is redirected to the captive portal port on the interface through which they are connected.

- **Port**—The captive portal port. The default is 885 (TCP). If you configure a different port, it must be in the range 1025-65535.
- **Decrypt Re-Sign Certificate**—Select the internal CA certificate to use for rules that implement decryption with re-signed certificates. You can use the pre-defined NGFW-Default-InternalCA certificate (the default), or one that you created or uploaded. If the certificate does not yet exist, click **Create Internal CA** to create it. (You are prompted for the decrypt re-sign certificate only if you have not yet enabled the SSL decryption policy.)

  If you have not already installed the certificate in client browsers, click the download button ( ) to obtain a copy. See the documentation for each browser for information on how to install the certificate. Also see Downloading the CA Certificate for Decrypt Re-Sign Rules, on page 446.

Example:

The Identity Policy Configuration dialog box should now look like the following.



h) Click **Save** to save the active authentication settings.

The Active Authentication tab now appears below the Action setting.

i) On the **Active Authentication** tab, select **HTTP Negotiate**.

This allows the browser and directory server to negotiate the strongest authentication protocol, in order, NTLM, then HTTP basic.

Note
If you do not supply a **Redirect to Host Name** FQDN, the HTTP Basic, HTTP Response Page, and NTLM authentication methods redirect the user to the captive portal using the IP address of the interface. However, for HTTP Negotiate, the user is redirected using the fully-qualified DNS name *firewall-hostname.AD-domain-name*. If you want to use HTTP Negotiate without a **Redirect to Host Name** FQDN, you must also update your DNS server to map this name to the IP addresses of all inside interfaces where you are requiring active authentication. Otherwise, the redirection cannot complete, and users cannot authenticate. We recommend that you always provide a **Redirect to Host Name** FQDN to ensure consistent behavior regardless of authentication method. If you cannot, or do not want to, update the DNS server, select one of the other authentication methods.

j) For **AD Identity Source**, click **Create New Identity Realm**.

If you already created your realm server object, simply select it and skip the steps for configuring the server.

Fill in the following fields, then click **OK**.

- **Name**—A name for the directory realm.
- **Type**—The type of directory server. Active Directory is the only supported type, and you cannot change this field.
- **Directory Username**, **Directory Password**—The distinguished username and password for a user with appropriate rights to the user information you want to retrieve. For Active Directory, the user does not need elevated privileges. You can specify any user in the domain. The username must be fully qualified; for example, Administrator@example.com (not simply Administrator).

  Note
  The system generates ldap-login-dn and ldap-login-password from this information. For example, Administrator@example.com is translated as cn=adminisntrator,cn=users,dc=example,dc=com. Note that cn=users is always part of this translation, so you must configure the user you specify here under the common name “users” folder.

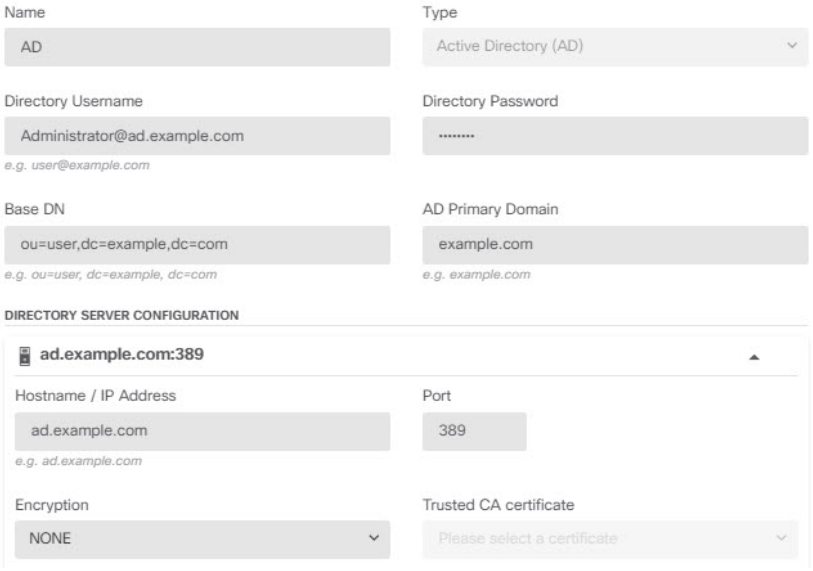
- **Base DN**—The directory tree for searching or querying user and group information, that is, the common parent for users and groups. For example, dc=example,dc=com. For information on finding the base DN, see Determining the Directory Base DN, on page 160.
- **AD Primary Domain**— The fully qualified Active Directory domain name that the device should join. For example, example.com.
- **Hostname/IP Address**—The hostname or IP address of the directory server. If you use an encrypted connection to the server, you must enter the fully-qualified domain name, not the IP address.
- **Port**—The port number used for communications with the server. The default is 389. Use port 636 if you select LDAPS as the encryption method.
- **Encryption**—To use an encrypted connection for downloading user and group information, select the desired method, **STARTTLS** or **LDAPS**. The default is **None**, which means that user and group information is downloaded in clear text.

- **STARTTLS** negotiates the encryption method, and uses the strongest method supported by the directory server. Use port 389. This option is not supported if you use the realm for remote access VPN.
- **LDAPS** requires LDAP over SSL. Use port 636.

- **Trusted CA Certificate**—If you select an encryption method, upload a Certificate Authority (CA) certificate to enable a trusted connection between the system and the directory server. 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 ad.example.com in the certificate, the connection fails.

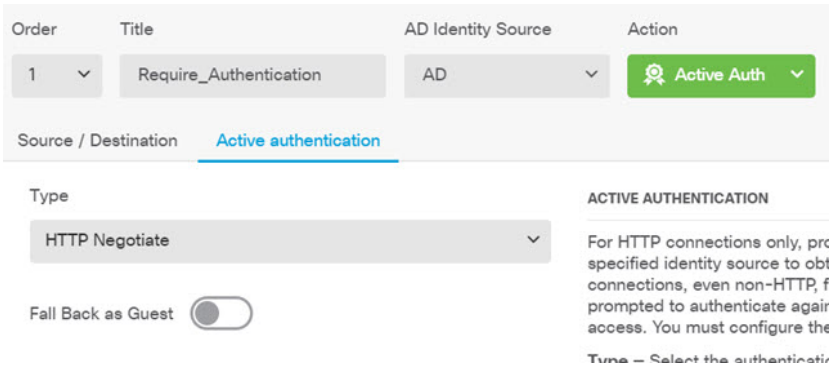
Example:

For example, the following image shows how to create an unencrypted connection for the ad.example.com server. The primary domain is example.com, and the directory username is Administrator@ad.example.com. All user and group information is under the Distinguished Name (DN) ou=user,dc=example,dc=com.



k) For **AD Identity Source**, select the object you just created.

The rule should look similar to the following.



l) Click **OK** to add the rule.

If you look in the upper right of the window, you can see that the **Deploy** icon button now has a dot, which indicates that there are undeployed changes. Making changes in the user interface is not sufficient for getting the changes configured on the device, you must deploy changes. Thus, you can make a set of related changes before you deploy them, so that you do not face the potential problems of having a partially-configured set of changes running on the device. You will deploy changes later in this procedure.

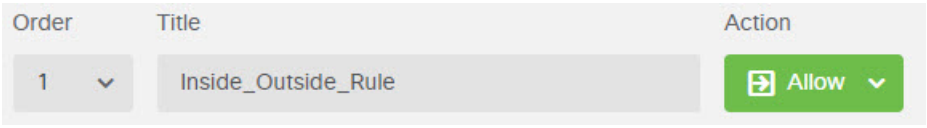


Step 2 Change the action on the Inside_Outside_Rule access control rule to **Allow**.

The Inside_Outside_Rule access rule is created as a trust rule. However, trusted traffic is not inspected, so the system cannot learn about some of the characteristics of trusted traffic, such as application, when the traffic matching criteria does not include application or other conditions besides zone, IP address, and port. If you change the rule to allow rather than trust traffic, the system fully inspects the traffic.

Note
(ISA 3000.) Also consider changing the Outside_Inside_Rule, Inside_Inside_Rule and Outside_Outside_Rule from Trust to Allow.

a) Click **Access Control** on the **Policies** page.
b) Hover over the **Actions** cell on the right side of the Inside_Outside_Rule row to expose the edit and delete icons, and click the edit icon ( ) to open the rule.
c) Select **Allow** for the **Action**.

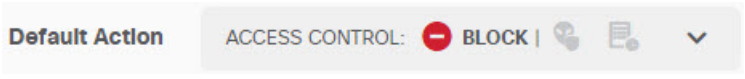


d) Click **OK** to save the change.

Step 3 Enable logging on the access control policy default action.

Dashboards contain information about connections only if the connection matches an access control rule that enables connection logging. The Inside_Outside_Rule enables logging, but the default action has logging disabled. Thus, dashboards show information for the Inside_Outside_Rule only, and do not reflect connections that do not match any rules.

a) Click anywhere in the default action at the bottom of the access control policy page.



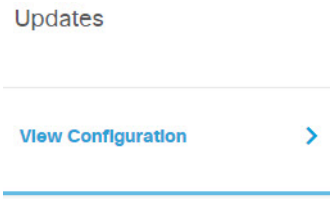
b) Select **Select Log Action** > **At Beginning and End of Connection**.
c) Click **OK**.

Step 4 Set an update schedule for the vulnerability database (VDB).

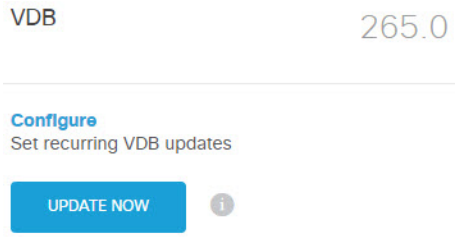
Cisco regularly releases updates to the VDB, which includes the application detectors that can identify the application used in a connection. You should update the VDB on a regular basis. You can either manually download updates, or you can set up a regular schedule. The following procedure shows how to set up a schedule. By default, VDB updates are disabled, so you need to take action to get VDB updates.

a) Click **Device**.

b) Click **View Configuration** in the Updates group.



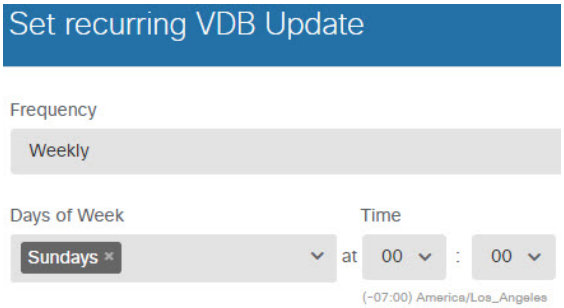
c) Click **Configure** in the VDB group.



d) Define the update schedule.

Choose a time and frequency that will not be disruptive to your network. Also, please understand that the system will do an automatic deployment after downloading the update. This is necessary to activate the new detectors. Thus, any configuration changes that you have made and saved but have not yet deployed will also be deployed.

For example, the following schedule updates the VDB once a week on Sunday at 12:00 AM (using the 24-hour clock notation).



e) Click **Save**.

Step 5 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.



b) Click the **Deploy Now** button.

You can wait until deployment completes, or click **OK** and check the task list or deployment history later.

What to do next

At this point, the monitoring dashboards and events should start showing information about users and applications. You can evaluate this information for undesirable patterns and develop new access rules to constrain unacceptable use.

If you want to start collecting information about intrusions and malware, you need to enable intrusion and file policies on one or more access rule. You also need to enable the licenses for those features.

If you want to start collecting information about URL categories, you must implement URL filtering.

How to Block Threats

You can implement next generation Intrusion Prevention System (IPS) filtering by adding intrusion policies to your access control rules. Intrusion policies analyze network traffic, comparing the traffic contents against known threats. If a connection matches a threat you are monitoring, the system drops the connection, thus preventing the attack.

All other traffic handling occurs before network traffic is examined for intrusions. By associating an intrusion 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.

You can configure intrusion policies on rules that **allow** traffic only. Inspection is not performed on rules set to **trust** or **block** traffic. In addition, you can configure an intrusion policy as part of the default action if the default action is **allow**.

The intrusion policies are designed by the Cisco Talos Intelligence Group (Talos), who set the intrusion and preprocessor rule states and advanced settings. You can create your own custom policies based on the Talos policies if you are using Snort 3 as the inspection engine.

Besides inspecting traffic that you allow for potential intrusions, you can use the Security Intelligence policy to preemptively block all traffic to or from known bad IP addresses, or to known bad URLs.

Procedure

Step 1 If you have not already done so, enable the Threat license.

You must enable the Threat license to use intrusion policies and Security Intelligence. If you are currently using the evaluation license, you are enabling an evaluation version of the license. If you have registered the device, you must purchase the required license and add it to your Smart Software Manager account on Cisco.com.

a) Click **Device**.
b) Click **View Configuration** in the Smart License group.



c) Click **Enable** in the **Threat** group.

The system registers the license with your account, or activates the evaluation license, as appropriate. The group should indicate that the license is enabled, and the button changes to a Disable button.



Step 2 Select an intrusion policy for one or more access rules.

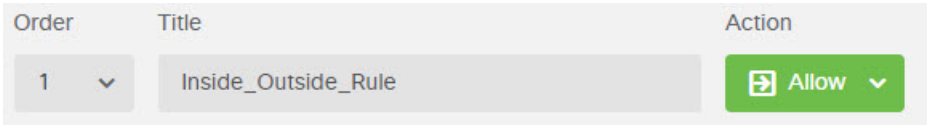
Determine which rules cover traffic that should be scanned for threats. For this example, we will add intrusion inspection to the Inside_Outside_Rule.

a) Click **Policies** in the main menu.

Ensure that the **Access Control** policy is displayed.

b) Hover over the **Actions** cell on the right side of the Inside_Outside_Rule row to expose the edit and delete icons, and click the edit icon ( ) to open the rule.

c) If you have not already done so, select **Allow** for the **Action**.

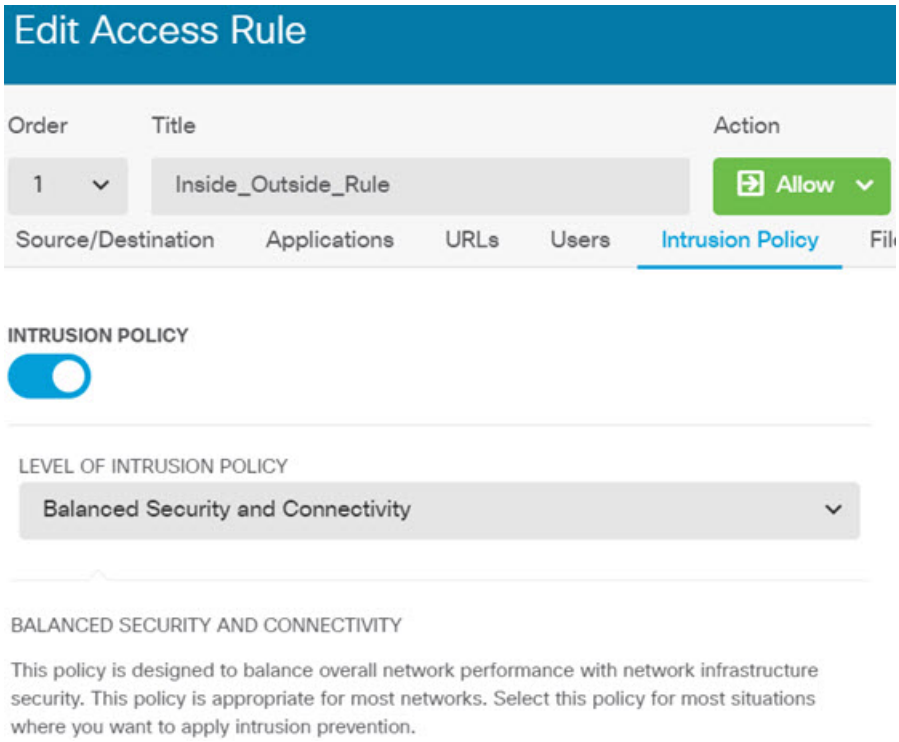


d) Click the **Intrusion Policy** tab.

e) Click the **Intrusion Policy** toggle to enable it, then select the intrusion policy.

The **Balanced Security and Connectivity** policy is appropriate for most networks. It provides a good intrusion defense without being overly aggressive, which has the potential of dropping traffic that you might not want to be dropped. If you determine that too much traffic is getting dropped, you can ease up on intrusion inspection by selecting the **Connectivity over Security** policy.

If you need to be aggressive about security, try the **Security over Connectivity** policy. The **Maximum Detection** policy offers even more emphasis on network infrastructure security with the potential for even greater operational impact.



f) Click **OK** to save the change.

Step 3 (Optional.) Go to **Policies** > **Intrusion**, click the gear icon, and configure a syslog server for the intrusion policy.

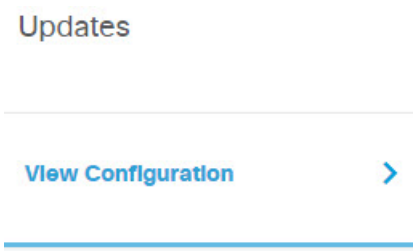
Intrusion events do not use the syslog server configured for the access control rule.

Step 4 Set an update schedule for the intrusion rule database.

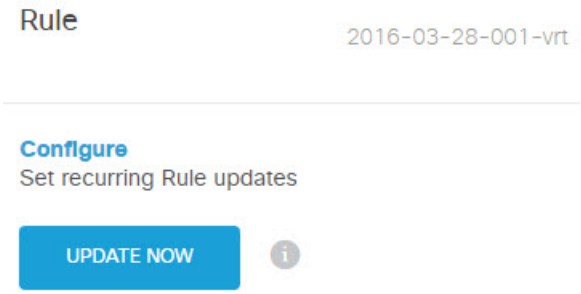
Cisco regularly releases updates to the intrusion rule database, which is used by intrusion policies to determine whether connections should be dropped. You should update the rule database on a regular basis. You can either manually download updates, or you can set up a regular schedule. The following procedure shows how to set up a schedule. By default, database updates are disabled, so you need to take action to get updated rules.

a) Click **Device**.

b) Click **View Configuration** in the Updates group.



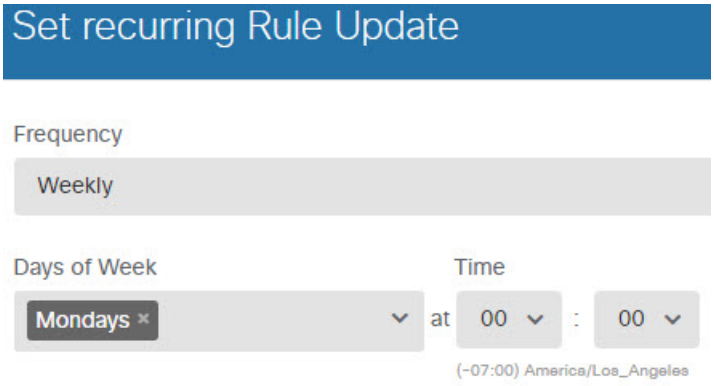
c) Click **Configure** in the Rule group.



d) Define the update schedule.

Choose a time and frequency that will not be disruptive to your network. Also, please understand that the system will do an automatic deployment after downloading the update. This is necessary to activate the new rules. Thus, any configuration changes that you have made and saved but have not yet deployed will also be deployed.

For example, the following schedule updates the rule database once a week on Monday at 12:00 AM (using the 24-hour clock notation).



e) Click **Save**.

Step 5 Configure the Security Intelligence policy to preemptively drop connections with known bad hosts and sites.

By using Security Intelligence to block connections with hosts or sites that are known to be threats, you save your system the time needed to do deep packet inspection to identify threats in each connection. Security Intelligence provides an early block of undesirable traffic, leaving more system time to handle the traffic you really care about.

a) Click **Device**, then click **View Configuration** in the **Updates** group.
b) Click **Update Now** in the Security Intelligence Feeds group.
c) Also, click **Configure** and set a recurring update for the feeds. The default, **Hourly**, is appropriate for most networks but you can decrease the frequency if necessary.
d) Click **Policies**, then click the **Security Intelligence** policy.
e) Click **Enable Security Intelligence** if you have not already enabled the policy.
f) On the **Network** tab, click **+** in the block/drop list, and select all of the feeds on the **Network Feeds** tab. You can click the **i** button next to a feed to read a description of each feed.

If you see a message that there are no feeds yet, please try again later. The feeds download has not yet completed. If this problem persists, ensure that there is a path between the management IP address and the Internet.

g) Click **OK** to add the selected feeds.

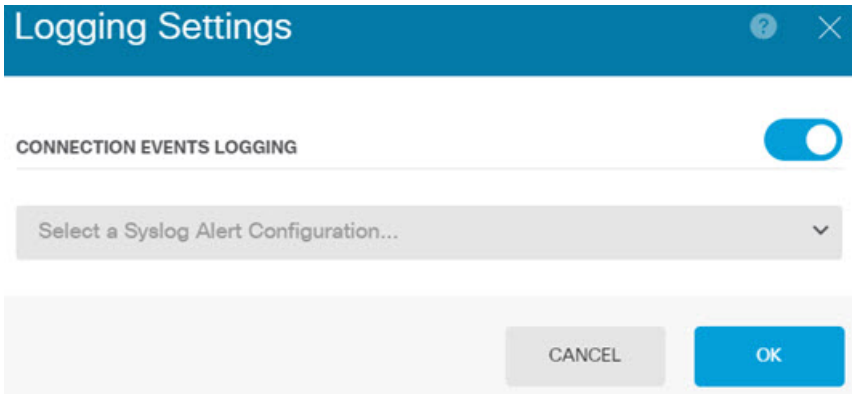
If you know of additional bad IP addresses, you can click + > **Network Objects** and add the objects that contain the addresses. You can click **Create New Network Object** at the bottom of the list to add them now.

h) Click the **URL** tab, then click + > **URL Feeds** in the block/drop list, and select all of the URL feeds. Click **OK** to add them to the list.

Similar to the network list, you can add your own URL objects to the list to block additional sites that are not in the feeds. Click + > **URL Objects**. You can add new objects by clicking **Create New URL Object** at the end of the list.

i) Click the gear icon, and enable **Connection Events Logging**, to enable the policy to generate Security Intelligence events for matched connections. Click **OK** to save your changes.

If you do not enable connection logging, you will have no data to use to evaluate whether the policy is performing to expectations. If you have an external syslog server defined, you can select it now so that the events are also sent to that server.



j) As needed, you can add network or URL objects to the **Do Not Block** list on each tab to create exceptions to the blocked list.

The **Do Not Block** lists are not real "allow" lists. They are exception lists. If an address or URL in the exception list also appears in the blocked list, the connection for the address or URL is allowed to pass on to the access control policy. This way, you can block a feed, but if you later find that a desirable address or site is being blocked, you can use the exception list to override that block without needing to remove the feed entirely. Keep in mind that these connections are subsequently evaluated by access control, and if configured, an intrusion policy. Thus, if any connections do contain threats, they can be identified and blocked during intrusion inspection.

Use the Access and SI Rules dashboard, and the Security Intelligence view in the Event Viewer, to determine what traffic is actually being dropped by the policy, and whether you need to add addresses or URLs to the **Do Not Block** lists.

Step 6 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.



b) Click the **Deploy Now** button.

You can wait until deployment completes, or click **OK** and check the task list or deployment history later.

What to do next

At this point, the monitoring dashboards and events should start showing information about attackers, targets, and threats, if any intrusions are identified. You can evaluate this information to determine if your network needs more security precautions, or if you need to reduce the level of intrusion policy you are using.

For Security Intelligence, you can see policy hits on the Access and SI Rules dashboard. You can also see Security Intelligence events in the Event Viewer. Security Intelligence blocks are not reflected in intrusion threat information, because the traffic is blocked before it can be inspected.

How to Block Malware

Users are continually at risk of obtaining malicious software, or *malware*, from Internet sites or other communication methods, such as e-mail. Even trusted web sites can be hijacked to serve malware to unsuspecting users. Web pages can contain objects coming from different sources. These objects can include images, executables, Javascript, advertisements, and so forth. Compromised web sites often incorporate objects hosted on external sources. Real security means looking at each object individually, not just the initial request.

Use file policies to detect malware using malware defense. You can also use file policies to perform file control, which allows control over all files of a specific type regardless of whether the files contain malware.

Malware defense uses the AMP Cloud to retrieve dispositions for possible malware detected in network traffic. The management interface must have a path to the Internet to reach the AMP Cloud and perform malware lookups. When the device detects an eligible file, it uses the file's SHA-256 hash value to query the AMP Cloud for the file's disposition. The possible disposition can be **clean**, **malware**, or **unknown** (no clear verdict). If the AMP Cloud is unreachable, the disposition is **unknown**.

By associating a 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 any files in the connection.

You can configure file policies on rules that **allow** traffic only. Inspection is not performed on rules set to **trust** or **block** traffic.

Procedure

Step 1 If you have not already done so, enable the Malware and Threat licenses.

You must enable the Malware to use file policies in addition to the Threat license, which is required for intrusion policies. If you are currently using the evaluation license, you are enabling an evaluation version of the licenses. If you have registered the device, you must purchase the required licenses and add them to your Smart Software Manager account on Cisco.com.

a) Click **Device**.
b) Click **View Configuration** in the Smart License group.



c) Click **Enable** in the **Malware** group, and if not already enabled, the **Threat** group.

The system registers the license with your account, or activates the evaluation license, as appropriate. The group should indicate that the license is enabled, and the button changes to a Disable button.



Step 2 Select a file policy for one or more access rules.

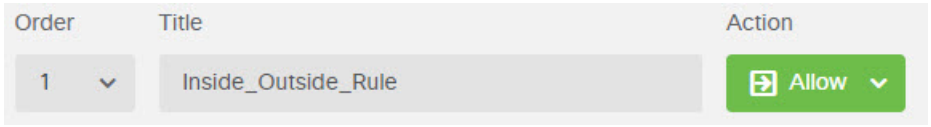
Determine which rules cover traffic that should be scanned for malware. For this example, we will add file inspection to the Inside_Outside_Rule.

a) Click **Policies** in the main menu.

Ensure that the **Access Control** policy is displayed.

b) Hover over the **Actions** cell on the right side of the Inside_Outside_Rule row to expose the edit and delete icons, and click the edit icon ( ) to open the rule.

c) If you have not already done so, select **Allow** for the **Action**.



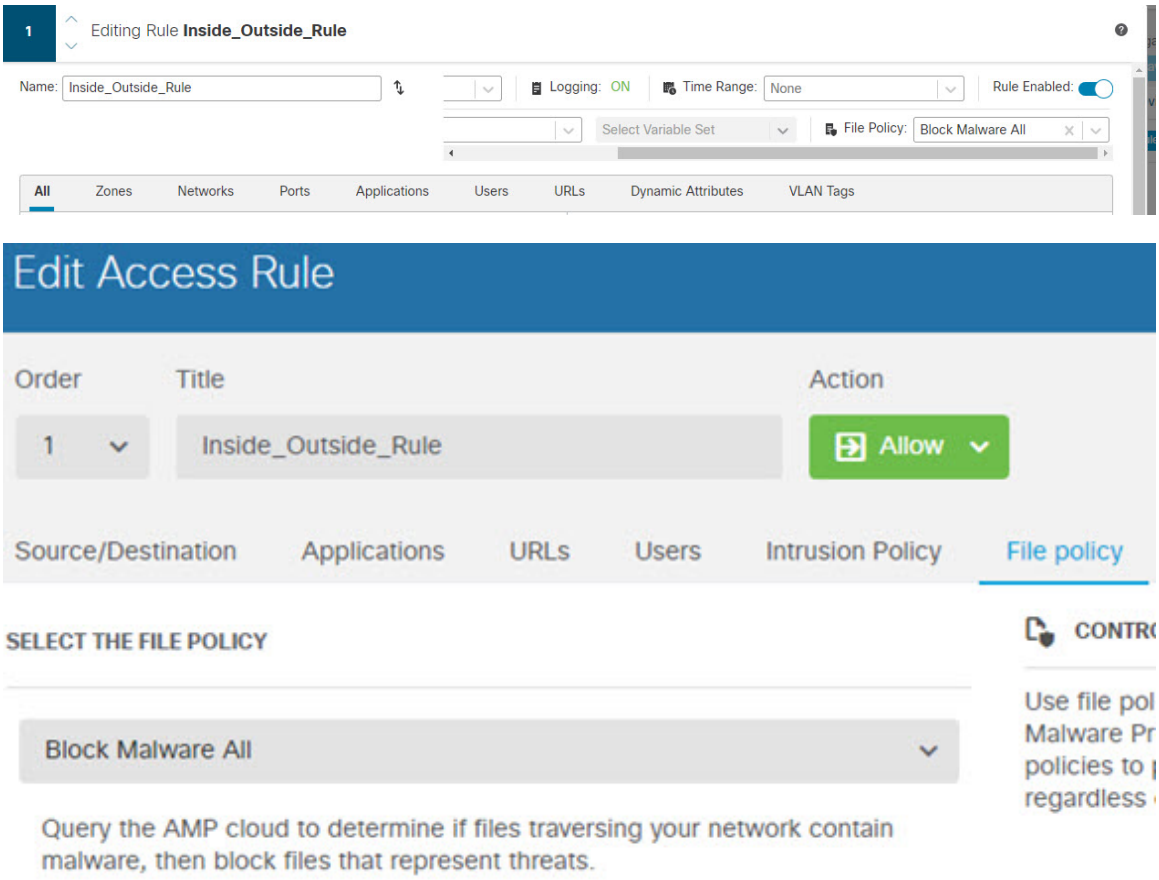
d) Click the **File Policy** tab.

e) Click the file policy you want to use.

Your main choice is between **Block Malware All**, which drops any files that are considered malware, or **Cloud Lookup All**, which queries the AMP Cloud to determine the file's disposition, but does no blocking. If you want to first see how files are being evaluated, use cloud lookup. You can switch to the blocking policy later if you are satisfied with how files are being evaluated.

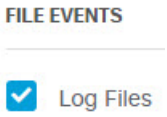
There are other policies available that block malware. These policies are coupled with file control, blocking the upload of Microsoft Office, or Office and PDF, documents. That is, these policies prevent users from sending these file types to other networks in addition to blocking malware. You can select these policies if they fit your needs.

For this example, select **Block Malware All**.



f) Click the **Logging** tab and verify that **Log Files** under File Events is selected.

By default, file logging is enabled whenever you select a file policy. You must enable file logging to get file and malware information in events and dashboards.



g) Click **OK** to save the change.

Step 3 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.



b) Click the **Deploy Now** button.

You can wait until deployment completes, or click **OK** and check the task list or deployment history later.

What to do next

At this point, the monitoring dashboards and events should start showing information about file types and file and malware events, if any files or malware are transmitted. You can evaluate this information to determine if your network needs more security precautions related to file transmissions.

How to Implement an Acceptable Use Policy (URL Filtering)

You might have an acceptable use policy for your network. Acceptable use policies differentiate between network activity that is appropriate in your organization and activity that is considered inappropriate. These policies are typically focused on Internet usage, and are geared towards maintaining productivity, avoiding legal liabilities (for example, maintaining a non-hostile workplace), and in general controlling web traffic.

You can use URL filtering to define an acceptable use policy with access policies. You can filter on broad categories, such as Gambling, so that you do not need to identify every individual web site that should be blocked. For category matches, you can also specify the relative reputation of sites to allow or block. If a user attempts to browse to any URL with that category and reputation combination, the session is blocked.

Using category and reputation data also simplifies policy creation and administration. It grants you assurance that the system will control web traffic as expected. Finally, because Cisco's threat intelligence is continually updated with new URLs, as well as new categories and risks for existing URLs, you can ensure that the system uses up-to-date information to filter requested URLs. Malicious sites that represent security threats such as malware, spam, botnets, and phishing may appear and disappear faster than you can update and deploy new policies.

The following procedure explains how to implement an acceptable use policy using URL filtering. For purposes of this example, we will block sites of any reputation in several categories, risky Social Networking sites, and an unclassified site, badsite.example.com.

Procedure

Step 1 If you have not already done so, enable the **URL** license.

You must enable the URL license to use URL category and reputation information, or to see the information in dashboards and events. If you are currently using the evaluation license, you are enabling an evaluation version of the license. If you have registered the device, you must purchase the required license and add it to your Smart Software Manager account on Cisco.com.

a) Click **Device**.
b) Click **View Configuration** in the Smart License group.



c) Click **Enable** in the **URL License** group.

The system registers the license with your account, or activates the evaluation license, as appropriate. The group should indicate that the license is enabled, and the button changes to a Disable button.



Step 2 Create a URL filtering access control rule.

You might want to first see the categories for sites your users are visiting before making a blocking rule. If that is the case, you can create a rule with the Allow action for an acceptable category, such as Finance. Because all web connections must be inspected to determine if the URL belongs to this category, you would get category information even for non-Finance sites.

But there are probably URL categories that you already know you want to block. A blocking policy also forces inspection, so you get category information on connections to unblocked categories, not just the blocked categories.

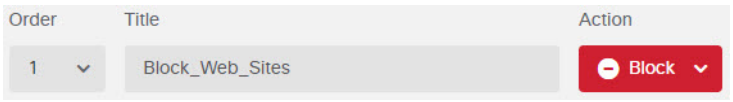
a) Click **Policies** in the main menu.

Ensure that the **Access Control** policy is displayed.

b) Click + to add a new rule.

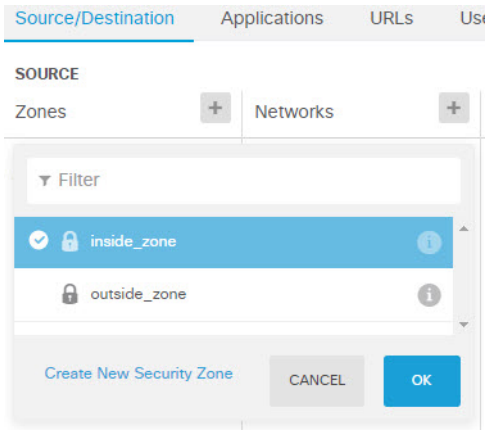
c) Configure the order, title, and action.

- **Order**—The default is to add new rules to the end of the access control policy. However, you must place this rule ahead of (above) any rule that would match the same Source/Destination and other criteria, or the rule will never be matched (a connection matches one rule only, and that is the first rule it matches in the table). For this rule, we will use the same Source/Destination as the Inside_Outside_Rule created during initial device configuration. You might have created other rules as well. To maximize access control efficiency, it is best to have specific rules early, to ensure the quickest decision on whether a connection is allowed or dropped. For the purposes of this example, select **1** as the rule order.
- **Title**—Give the rule a meaningful name, such as Block_Web_Sites.
- **Action**—Select **Block**.

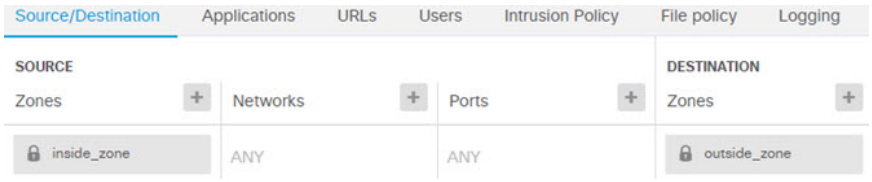


d) On the **Source/Destination** tab, click + for **Source** > **Zones**, select **inside_zone**, then click **OK** in the zones dialog box.

Adding any of the criteria works the same way. Clicking + opens a little dialog box, where you click the items you want to add. You can click multiple items, and clicking a selected item de-selects it; the check marks indicate the selected items. But nothing is added to the policy until you click the **OK** button; simply selecting the items is not sufficient.



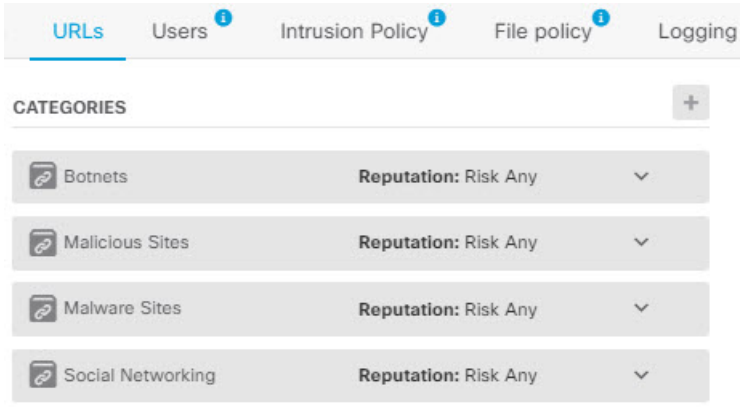
e) Using the same technique, select **outside_zone** for **Destination** > **Zones**.



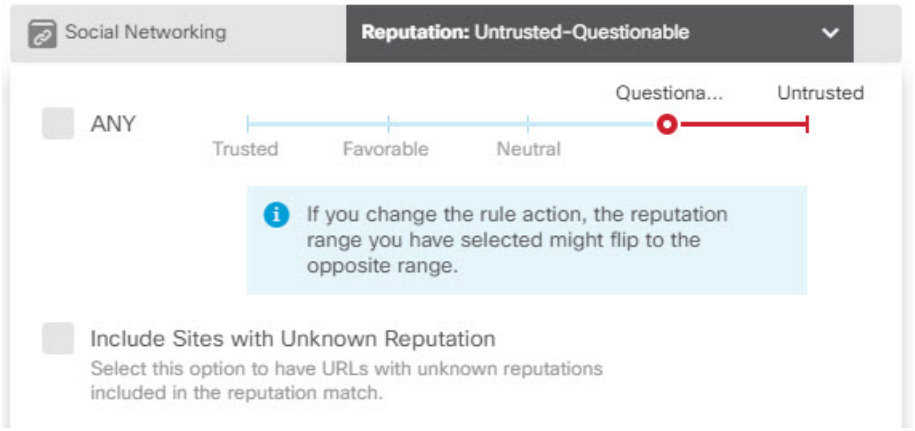
f) Click the **URLs** tab.

g) Click the + for **Categories**, and select the categories you want to fully or partially block.

For purposes of this example, select Botnets, Malicious Sites, Malware Sites, and Social Networking. There are additional categories that you would most likely want to block. If you know of a site that you want to block, but you are unsure of the category, enter the URL in the in the **URL to Check** field and click **Go**. You will be taken to a web site that shows the results of the lookup.



h) To implement reputation-sensitive blocking for the Social Networking category, click **Reputation: Risk Any** for that category, deselect **Any**, then move the slider to **Questionable**. Click away from the slider to close it.



The left of the reputation slider indicates sites that will be allowed, the right side are sites that will be blocked. In this case, only Social Networking sites with reputations in the Questionable and Untrusted ranges will be blocked. Thus, your users should be able to get to commonly-used Social Networking sites, where there are fewer risks.

Select the **Include Sites with Unknown Reputation** option to have URLs with unknown reputation included in the reputation match. New sites typically are unrated, and there can be other reasons a site's reputation is unknown or cannot be determined.

Using reputation, you can selectively block sites within a category you otherwise want to allow.

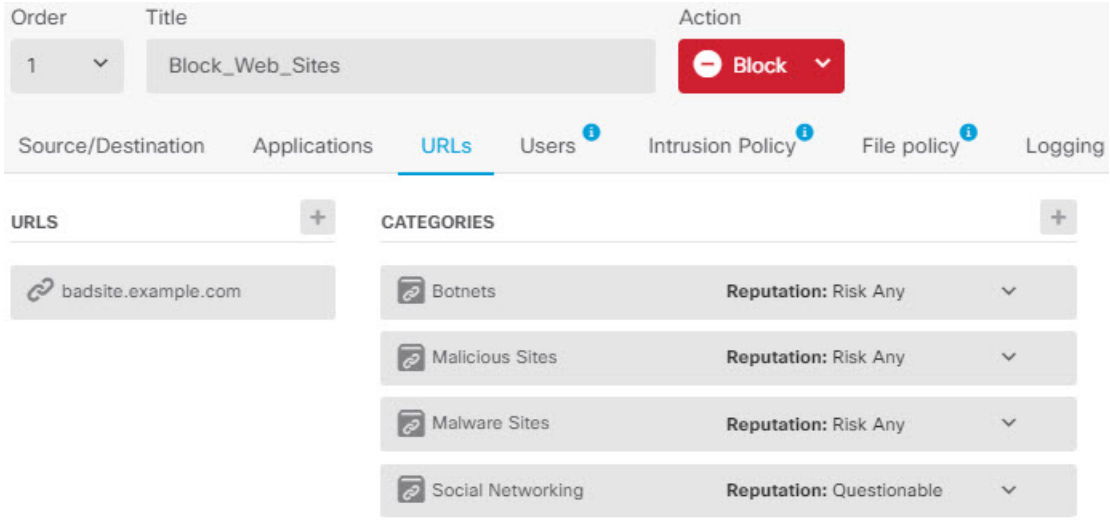
i) Click the + next to the **URLS** list to the left of the categories list.
j) At the bottom of the popup dialog box, click the **Create New URL** link.
k) Enter **badsite.example.com** for both the name and URL, then click **OK** to create the object.

You can name the object the same as the URL or give the object a different name. For the URL, do not include the protocol portion of the URL, just add the server name.



l) Select the new object, then click **OK**.

Adding new objects while editing policies simply adds the object to the list. The new object is not automatically selected.



m) Click the **Logging** tab and select **Select Log Action** > **At Beginning and End of Connection**.

You must enable logging to get category and reputation information into the web category dashboard and connection events.

n) Click **OK** to save the rule.

Step 3 (Optional.) Set preferences for URL filtering.

When you enable the URL license, the system automatically enables updates to the web category database. The system checks for updates every 30 minutes, although the data is typically updated once per day. You can turn off these updates if for some reason you do not want them.

You can also elect to send URLs that are not categorized to Cisco for analysis. Thus, if the installed URL database does not have a categorization for a site, the Cisco Cloud might have one. The cloud returns the category and reputation, and your category-based rules can then be applied correctly to the URL request. Selecting this option is important for lower-end systems, which install a smaller URL database due to memory limitations. You can set a time to live for the lookup results: the default is Never, which means lookup results are never refreshed.

a) Click **Device**.
b) Click **System Settings** > **Traffic Settings** > **URL Filtering Preferences**.
c) Select **Query Cisco CSI for Unknown URLs**.
d) Select a reasonable **URL Time to Live**, such as 24 hours.
e) Click **Save**.

Step 4 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.



b) Click the **Deploy Now** button.

You can wait until deployment completes, or click **OK** and check the task list or deployment history later.

What to do next

At this point, the monitoring dashboards and events should start showing information about URL categories and reputations, and which connections were dropped. You can evaluate this information to determine if your URL filtering is dropping just those sites that are objectionable, or if you need to ease up on the reputation setting for certain categories.

Consider informing users beforehand that you will be blocking access to web sites based on their categorization and reputation.

How to Control Application Usage

The Web has become the ubiquitous platform for application delivery in the enterprise, whether that is browser based application platforms, or rich media applications that use web protocols as the transport in and out of enterprise networks.

Firewall Threat Defense inspects connections to determine the application being used. This makes it possible to write access control rules targeted at applications, rather than just targeting specific TCP/UDP ports. Thus, you can selectively block or allow web-based applications even though they use the same port.

Although you can select specific applications to allow or block, you can also write rules based on type, category, tag, risk, or business relevance. For example, you could 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 session is blocked.

Cisco frequently updates and adds additional application detectors via system and vulnerability database (VDB) updates. Thus, a rule blocking high risk applications can automatically apply to new applications without you having to update the rule manually.

In this use case, we will block any application that belongs to the **anonymizer/proxy** category.

Before you begin

This use case assumes that you completed the use case How to Gain Insight Into Your Network Traffic, on page 44. That use case explains how to collect application usage information, which you can analyze in the Applications dashboard. Understanding what applications are actually being used can help you design effective application-based rules. The use case also explains how to schedule VDB updates, which will not be repeated here. Ensure that you update the VDB regularly so that applications can be correctly identified.

Procedure

Step 1 Create the application-based access control rule.

a) Click **Policies** in the main menu.

Ensure that the **Access Control** policy is displayed.

b) Click + to add a new rule.

c) Configure the order, title, and action.

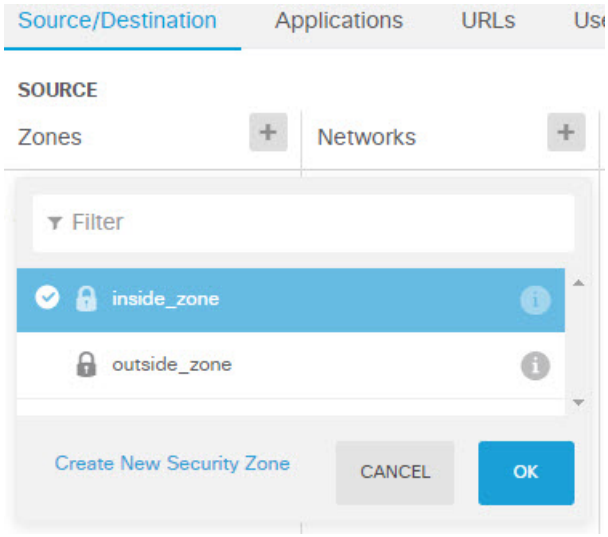
- **Order**—The default is to add new rules to the end of the access control policy. However, you must place this rule ahead of (above) any rule that would match the same Source/Destination and other criteria, or the rule will never be matched (a connection matches one rule only, and that is the first

rule it matches in the table). For this rule, we will use the same Source/Destination as the Inside_Outside_Rule created during initial device configuration. You might have created other rules as well. To maximize access control efficiency, it is best to have specific rules early, to ensure the quickest decision on whether a connection is allowed or dropped. For the purposes of this example, select **1** as the rule order.

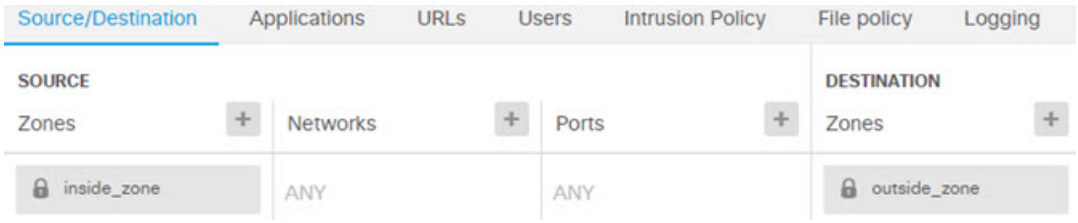
- **Title**—Give the rule a meaningful name, such as Block_Anonymizers.
- **Action**—Select **Block**.



d) On the **Source/Destination** tab, click **+** for **Source** > **Zones**, select **inside_zone**, then click **OK** in the zones dialog box.



e) Using the same technique, select **outside_zone** for **Destination** > **Zones**.



f) Click the **Applications** tab.

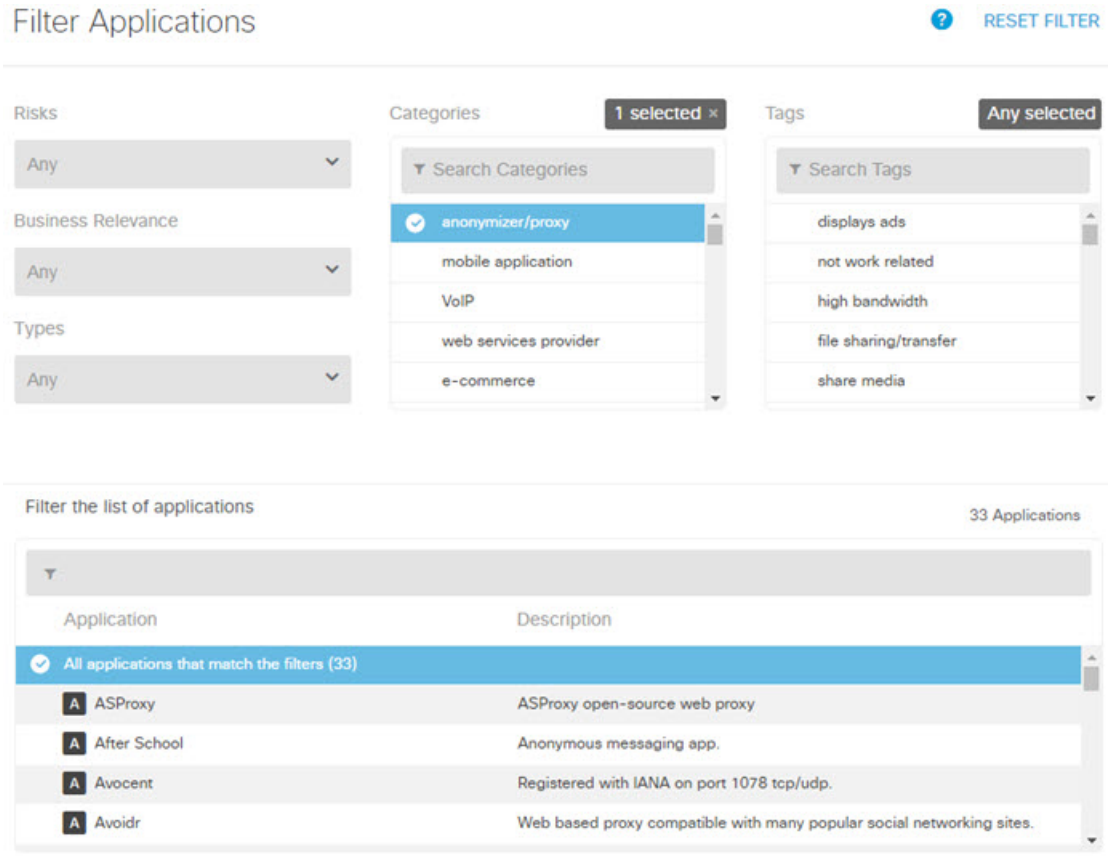
g) Click the **+** for **Applications**, and then click the **Advanced Filter** link at the bottom of the popup dialog box.

Although you can create application filter objects beforehand and select them on the Application Filters list here, you can also specify criteria directly in the access control rule, and optionally save the criteria as a filter object. Unless you are writing a rule for a single application, it is easier to use the Advanced Filter dialog box to find applications and construct appropriate criteria.

As you select criteria, the Applications list at the bottom of the dialog box updates to show exactly which applications match the criteria. The rule you are writing applies to these applications.

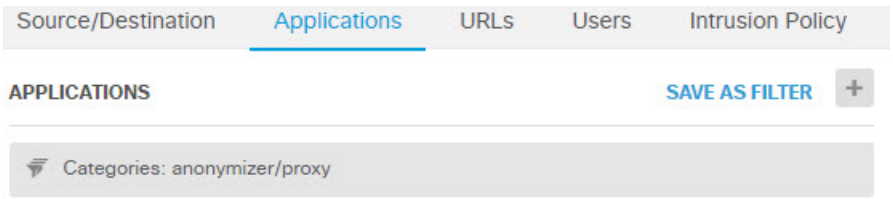
Look at this list carefully. For example, you might be tempted to block all very high risk applications. However, as of this writing, TFPT is classified as very high risk. Most organizations would not want to block this application. Take the time to experiment with various filter criteria to see which applications match your selections. Keep in mind that these lists can change with every VDB update.

For purposes of this example, select anonymizers/proxies from the Categories list.



h) Click **Add** in the Advanced Filters dialog box.

The filter is added and shown on the Applications tab.



i) Click the **Logging** tab and select **Select Log Action** > **At Beginning and End of Connection**.

You must enable logging to get information about any connections blocked by this rule.

j) Click **OK** to save the rule.

Step 2 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.



b) Click the **Deploy Now** button.

You can wait until deployment completes, or click **OK** and check the task list or deployment history later.

Step 3 Click **Monitoring** and evaluate the results.

You might now see dropped connections on the Applications widget on the **Network Overview** dashboard. Use the **All/Denied/Allowed** drop-down options to focus just on dropped applications.

You can also find information about the applications on the **Web Applications** dashboard. The **Applications** dashboards show protocol-related results. If someone tries to use these applications, you should be able to correlate the application with the user attempting the connection, assuming that you enable identity policies and require authentication.

How to Add a Subnet

If you have an available interface on your device, you can wire it to a switch (or another router) to provide services to another subnet.

There are many potential reasons you would add a subnet. For this use case, we will address the following typical scenario.

- The subnet is an inside network using the private network 192.168.2.0/24.
- The interface for the network has the static address 192.168.2.1. In this example, the physical interface is devoted to the network. Another option is to use an already-wired interface and create a subinterface for the new network.
- The device will provide addresses to workstations on the network using DHCP, using 192.168.2.2-192.168.2.254 as the address pool.
- Network access to other inside networks, and to the outside network, will be allowed. Traffic going to the outside network will use NAT to obtain a public address.



Note This example assumes the unused interface is not part of a bridge group. If it is currently a bridge group member, you must first remove it from the bridge group before following this procedure.

Before you begin

Physically connect the network cable to the interface and to the switch for the new subnet.

Procedure

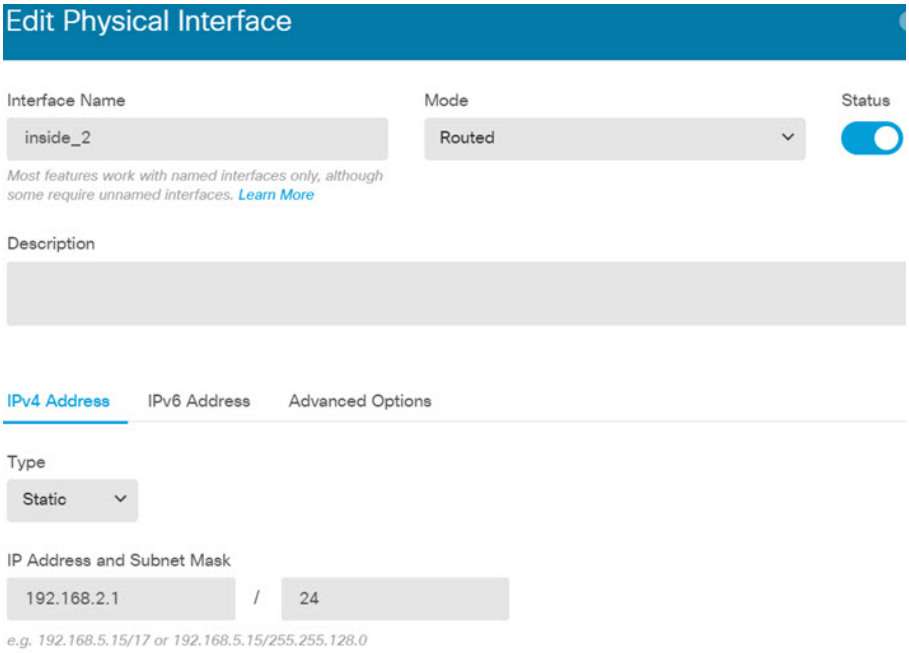
Step 1 Configure the interface.

a) Click **Device**, click the link in the **Interfaces** summary, and then click the interfaces type to view the list of interfaces.

b) Hover over the **Actions** cell on the right side of the row for the interface you wired, and click the edit icon ( ).

c) Configure the basic interface properties.

- **Name**—A unique name for the interface. For this example, **inside_2**.
- **Mode**—Select **Routed**.
- **Status**—Click the status toggle to enable the interface.
- **IPv4 Address** tab—Select **Static** for **Type**, then enter **192.168.2.1/24**.



d) Click **Save**.

The interface list shows the updated interface status and the configured IP address.



Step 2 Configure the DHCP server for the interface.

a) Click **Device**.

b) Click **System Settings** > **DHCP Server**.

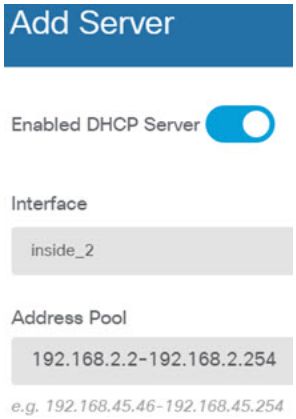
c) Click the **DHCP Servers** tab.

The table lists any existing DHCP servers. If you are using the default configuration, the list includes one for the inside interface.

d) Click + above the table.

e) Configure the server properties.

- **Enable DHCP Server**—Click this toggle to enable the server.
- **Interface**—Select the interface on which you are providing DHCP services. In this example, select inside_2.
- **Address Pool**—The addresses the server can supply to devices on the network. Enter 192.168.2.2-192.168.2.254. Make sure you do not include the network address (.0), the interface address (.1), or the broadcast address (.255). Also, if you need static addresses for any devices on the network, exclude those addresses from the pool. The pool must be a single continuous series of addresses, so choose static addresses from the beginning or ending of the range.



f) Click **Add**.



Step 3 Add the interface to the inside security zone.

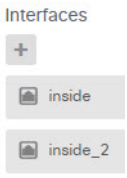
To write policies on an interface, the interface must belong to a security zone. You write policies for the security zones. Thus, as you add and remove interfaces in the zones, you automatically change the policies applied to the interface.

a) Click **Objects** in the main menu.

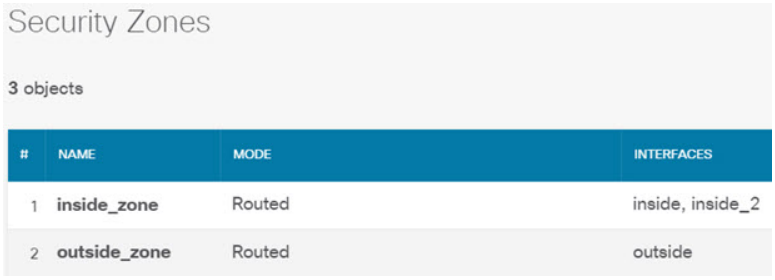
b) Select **Security Zones** from the objects table of contents.

c) Hover over the **Actions** cell on the right side of the row for the **inside_zone** object, and click the edit icon ( ).

d) Click + under **Interfaces**, select the inside_2 interface, and click **OK** in the interfaces list.



e) Click **Save**.



Step 4 Create an access control rule that allows traffic between the inside networks.

Traffic is not automatically allowed between any interfaces. You must create access control rules to allow the traffic that you want. The only exception is if you allow traffic in the access control rule's default action. For the purposes of this example, we will assume you retained the block default action that the device setup wizard configures. Thus, you need to create a rule that will allow traffic between the inside interfaces. If you have already created a rule like this, skip this step.

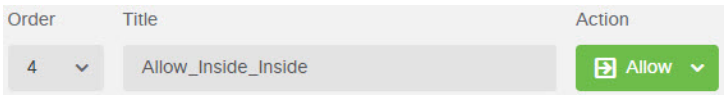
a) Click **Policies** in the main menu.

Ensure that the **Access Control** policy is displayed.

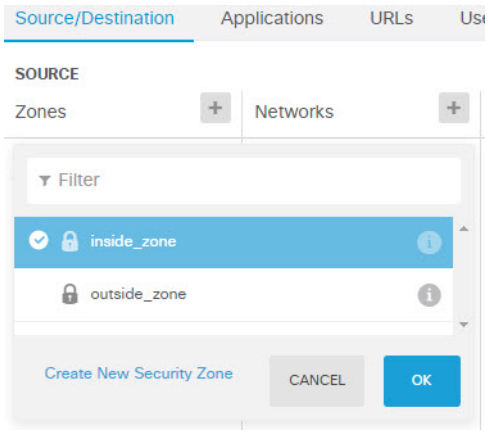
b) Click + to add a new rule.

c) Configure the order, title, and action.

- **Order**—The default is to add new rules to the end of the access control policy. However, you must place this rule ahead of (above) any rule that would match the same Source/Destination and other criteria, or the rule will never be matched (a connection matches one rule only, and that is the first rule it matches in the table). For this rule, we will use unique Source/Destination criteria, so adding the rule to the end of the list is acceptable.
- **Title**—Give the rule a meaningful name, such as Allow_Inside_Inside.
- **Action**—Select **Allow**.

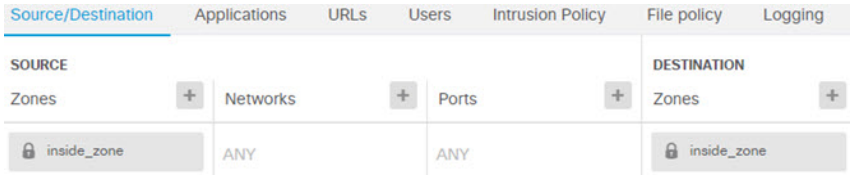


d) On the **Source/Destination** tab, click + for **Source** > **Zones**, select **inside_zone**, then click **OK** in the zones dialog box.



e) Using the same technique, select **inside_zone** for **Destination** > **Zones**.

A security zone must contain at least two interfaces to select the same zone for source and destination.



f) (Optional.) Configure intrusion and malware inspection.

Although the inside interfaces are in a trusted zone, it is typical for users to connect laptops to the network. Thus, a user might unknowingly bring a threat inside your network from an outside network or a Wi-Fi hot spot. Thus, you might want to scan for intrusions and malware in traffic that goes between your inside networks.

Consider doing the following.

- Click the **Intrusion Policy** tab, enable the intrusion policy, and use the slider to select the Balanced Security and Connectivity policy.
- Click the **File Policy** tab, then select the Block Malware All policy.

g) Click the **Logging** tab and select **Select Log Action** > **At Beginning and End of Connection**.

You must enable logging to get information about any connections that match this rule. Logging adds statistics to the dashboard as well as showing events in the event viewer.

h) Click **OK** to save the rule.

Step 5 Verify that required policies are defined for the new subnet.

By adding the interface to the inside_zone security zone, any existing policies for inside_zone automatically apply to the new subnet. However, take the time to inspect your policies and ensure that no additional policies are needed.

If you completed the initial device configuration, the following policies should already apply.

- **Access Control**—The Inside_Outside_Rule should allow all traffic between the new subnet and the outside network. If you followed the previous use cases, the policy also provides intrusion and malware

inspection. You must have a rule that allows some traffic between the new network and the outside network, or users cannot access the Internet or other external networks.

- **NAT**—The InsideOutsideNATrule applies to any interface going to the outside interface, and applies interface PAT. If you kept this rule, traffic from the new network going to the outside will have the IP address translated to a unique port on the outside interface's IP address. If you do not have a rule that applies to all interfaces, or the inside_zone interfaces, when going to the outside interface, you might need to create one now.
- **Identity**—There is no default identity policy. However, if you followed previous use cases, you might have an identity policy that already requires authentication for the new network. If you do not have an identity policy that applies, create one now if you want to have user-based information for the new network.

Step 6 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.



b) Click the **Deploy Now** button.

You can wait until deployment completes, or click **OK** and check the task list or deployment history later.

What to do next

Verify that workstations on the new subnet are getting IP addresses using DHCP, and that they can reach other inside networks and the outside network. Use the monitoring dashboards and the event viewer to evaluate network usage.

How to Passively Monitor the Traffic on a Network

A Firewall Threat Defense device is normally deployed as an active firewall and IPS (intrusion prevention system) security device. The core function of the device is to provide active protection to the network, dropping undesirable connections and threats.

However, you can also deploy the system in a passive mode, where the device simply analyzes the traffic on monitored switch ports. This mode is mainly for demonstration or testing purposes, so that you can become comfortable with the device before deploying it as an active firewall. Using a passive deployment, you can monitor the kinds of threats that appear on the network, the URL categories users are browsing, and so forth.

Although you would normally use passive mode for demonstration or testing purposes only, you can also use passive mode in a production environment if it provides a service that you need, such as IDS (intrusion detection system, without prevention). You can mix passive interfaces with active firewall routed interfaces to provide the exact combination of services required by your organization.

The following procedure explains how to deploy the system passively to analyze the traffic coming through a limited number of switch ports.

Note This example is for a hardware Firewall Threat Defense device. You can also use passive mode for Firewall Threat Defense Virtual, but the network setup is different. For details, see Configure the VLAN for a Firewall Threat Defense Virtual Passive Interface, on page 278. Otherwise, this procedure also applies to Firewall Threat Defense Virtual.

Before you begin

This procedure assumes that you have connected the inside and outside interfaces and completed the initial device setup wizard. Even in a passive deployment, you need a connection to the Internet to download updates for the system databases. You also need to be able to connect to the management interface to open Firewall Device Manager, which you can do through direct connections to the inside or management port.

The example also assumes that you have enabled syslog for intrusion policies on the **Policies** > **Intrusion** page.

Procedure

Step 1 Configure a switch port as a SPAN (Switched Port Analyzer) port and configure a monitoring session for the source interfaces.

The following example sets up a SPAN port and monitoring session for two source interfaces on a Cisco Nexus 5000 series switch. If you are using a different type of switch, the required commands might be different.

```
switch(config)# interface Ethernet1/48
switch(config-if)# switchport monitor
switch(config-if)# exit
switch(config)# monitor session 1
switch(config-monitor)# source interface ethernet 1/7
switch(config-monitor)# source interface ethernet 1/8
switch(config-monitor)# destination interface ethernet 1/48
switch(config-monitor)# no shut
```

To verify:

```
switch# show monitor session 1 brief
   session 1
---------------
type              : local
state             : up
source intf       :
    rx            : Eth1/7        Eth1/8
    tx            : Eth1/7        Eth1/8
    both          : Eth1/7        Eth1/8
source VSANs      :
destination ports : Eth1/48

Legend: f = forwarding enabled, l = learning enabled
```

Step 2 Connect the Firewall Threat Defense interface to the SPAN port on the switch.

It is best to select a currently unused port on the Firewall Threat Defense device. Based on the example switch configuration, you would connect the cable to Ethernet 1/48 on the switch. This is the destination interface for the monitoring session.

Step 3 Configure the Firewall Threat Defense interface in passive mode.

a) Click **Device**, then click the link in the **Interfaces** summary, and then click **Interfaces** or **EtherChannels**.

b) Click the edit icon ( ) for the physical interface or EtherChannel you want to edit.

Pick a currently unused interface. If you intend to convert an in-use interface to a passive interface, you need to first remove the interface from any security zone and remove all other configurations that use the interface.

c) Set the **Status** slider to the enabled setting ( ).

d) Configure the following:

- **Interface Name**—The name for the interface, up to 48 characters. Alphabetic characters must be lower case. For example, **monitor**.
- **Mode**—Select **Passive**.



e) Click **OK**.

Step 4 Create a passive security zone for the interface.

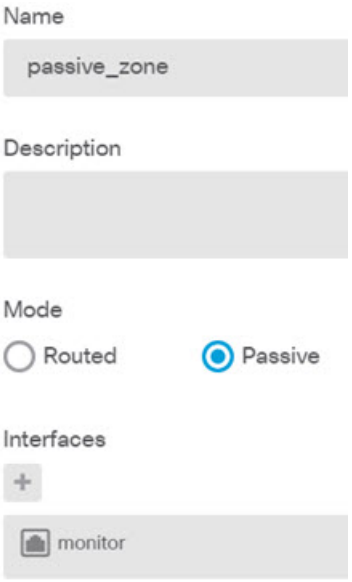
a) Select **Objects**, then select **Security Zones** from the table of contents.

b) Click the + button.

c) Enter a **Name** for the object and optionally, a description. For example, **passive_zone**.

d) For **Mode**, select **Passive**.

e) Click + and select the passive interface.



f) Click **OK**.

Step 5 Configure one or more access control rules for the passive security zone.

The number and type of rules you create depends on the information you want to gather. For example, if you want to configure the system as an IDS (intrusion detection system), you need at least one Allow rule with an assigned intrusion policy. If you want to collect URL category data, you need at least one rule that has a URL category specification.

You can create Block rules to see what connections the system would have blocked on an actively routed interface. These connections are not actually blocked, because the interface is passive, but you will see clearly how the system would have groomed the traffic on the network.

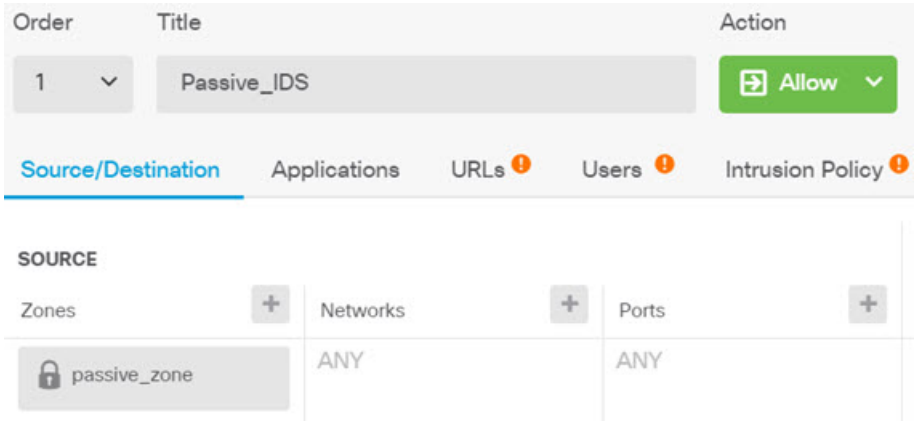
The following use cases cover the main uses for access control rules. These also apply to passive interfaces. Simply select the passive security zone as the source zone for the rules you create.

- How to Block Threats, on page 51
- How to Block Malware, on page 56
- How to Implement an Acceptable Use Policy (URL Filtering), on page 59
- How to Control Application Usage, on page 64

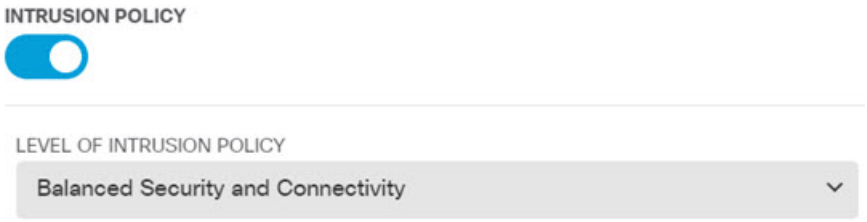
The following procedure creates two Allow rules to apply an intrusion policy and to collect URL category data.

a) Select **Policies** > **Access Control**.
b) Click **+** to add a rule allowing all traffic, but applying an intrusion policy.
c) Select **1** as the rule order. This rule is more specific than the default rule, but does not overlap with it. If you already have custom rules, select an appropriate position so that traffic to the passive interface is not matched to those rules instead.
d) Enter a name for the rule, for example, **Passive_IDS**.
e) Select **Allow** as the **Action**.
f) On the **Source/Destination** tab, select the passive zone under **Source** > **Zones**. Do not configure any other options on the tab.

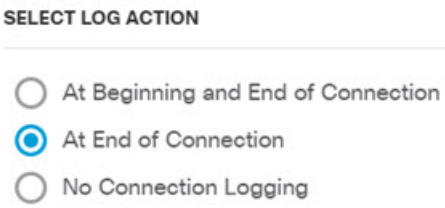
When running in evaluation mode, the rule should look like the following at this point:



g) Click the **Intrusion Policy** tab, click the slider to **On**, and select an intrusion policy such as the **Balanced Security and Connectivity** policy, which is recommended for most networks.



h) Click the **Logging** tab and select **At End of Connection** for the logging option.



i) Click **OK**.

j) Click **+** to add a rule that will require that the system do deep inspection to determine the URL and category for all HTTP requests.

This rule makes it possible for you to see URL category information in the dashboards. To save processing time and improve performance, the system determines URL category only if there is at least one access control rule that specifies a URL category condition.

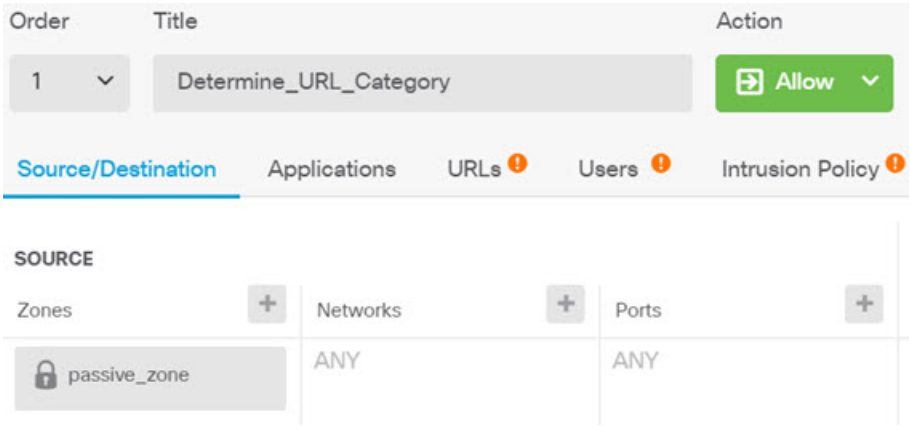
k) Select **1** as the rule order. This will place it above the previous rule (Passive_IDS). If you place it after that rule (which applies to all traffic), the rule you are creating now would never be matched.

l) Enter a name for the rule, for example, **Determine_URL_Category**.

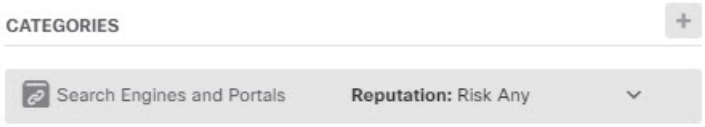
m) Select **Allow** as the **Action**.

Alternatively, you could select **Block**. Either action will accomplish your goal for this rule.

n) On the **Source/Destination** tab, select the passive zone under **Source** > **Zones**. Do not configure any other options on the tab.



o) Click the **URLs** tab, click the **+** next to the **Categories** heading, and select any of the categories. For example, **Search Engines and Portals**. You can optionally select a reputation level, or leave it at the default Any.



p) Click the **Intrusion Policy** tab, click the slider to **On**, and select the same intrusion policy you chose for the first rule.

q) Click the **Logging** tab and select **At End of Connection** for the logging option.

However, if you selected **Block** as the action, select **At Beginning and End of Connection**. Because blocked connections are not ended per se, you get log information at the beginning of the connection only.

r) Click **OK**.

Step 6 (Optional.) Configure other security policies.

You can also configure the following security policies to see how they would impact traffic:

- **Identity**—To collect user information. You can configure a rule in the identity policy to ensure that the user associated with a source IP address is identified. The process for implementing identity policies for passive interfaces is the same as the one for routed interfaces. Please follow the use case described in How to Gain Insight Into Your Network Traffic, on page 44.
- **Security Intelligence**—To block known bad IP addresses and URLs. For details, see How to Block Threats, on page 51.

Note
All encrypted traffic on passive interfaces is classified as undecryptable, so SSL decryption rules are ineffective and do not apply to passive interfaces.

Step 7 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.



b) Click the **Deploy Now** button.

You can wait until deployment completes, or click **OK** and check the task list or deployment history later.

Step 8 Use the monitoring dashboards to analyze the kinds of traffic and threats that are coming across the network. If you decide you want the Firewall Threat Defense device to actively drop unwanted connections, redeploy the device so that you can configure active routed interfaces that provide firewall protection for the monitored network.

More Examples

In addition to the examples in the Use Case chapter, there are example configurations in some of the chapters that explain specific services. You might find the following examples of interest.

Access Control

- How to Control Network Access Using TrustSec Security Group Tags, on page 503

Network Address Translation (NAT)

NAT for IPv4 addresses

- Providing Access to an Inside Web Server (Static Auto NAT), on page 593
- Single Address for FTP, HTTP, and SMTP (Static Auto NAT-with-Port-Translation), on page 596
- Different Translation Depending on the Destination (Dynamic Manual PAT), on page 602
- Different Translation Depending on the Destination Address and Port (Dynamic Manual PAT), on page 608
- DNS Reply Modification, DNS Server on Outside, on page 620
- DNS Reply Modification, DNS Server on Host Network, on page 623
- Exempting Site-to-Site VPN Traffic from NAT, on page 652

NAT for IPv6 addresses

- NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet, on page 579
- NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet and DNS Translation, on page 581
- NAT66 Example, Static Translation between Networks, on page 586
- NAT66 Example, Simple IPv6 Interface PAT, on page 589
- DNS 64 Reply Modification, on page 614

Remote Access Virtual Private Network (RA VPN)

- How to Implement RADIUS Change of Authorization, on page 701
- How to Configure Two-Factor Authentication using Duo LDAP, on page 709
- How to Provide Internet Access on the Outside Interface for Remote Access VPN Users (Hair Pinning), on page 716
- How to Use a Directory Server on an Outside Network with Remote Access VPN, on page 720
- How to Control RA VPN Access By Group, on page 733
- How to Allow RA VPN Access to Internal Networks in Different Virtual Routers, on page 737
- How to Customize the AnyConnect Client Icon and Logo, on page 740

Site-to-Site Virtual Private Network (VPN)

- Exempting Site-to-Site VPN Traffic from NAT, on page 652
- How to Provide Internet Access on the Outside Interface for External Site-to-Site VPN Users (Hair Pinning), on page 658
- How to Secure Traffic from Networks in Multiple Virtual Routers over a Site-to-Site VPN, on page 665

SSL/TLS Decryption

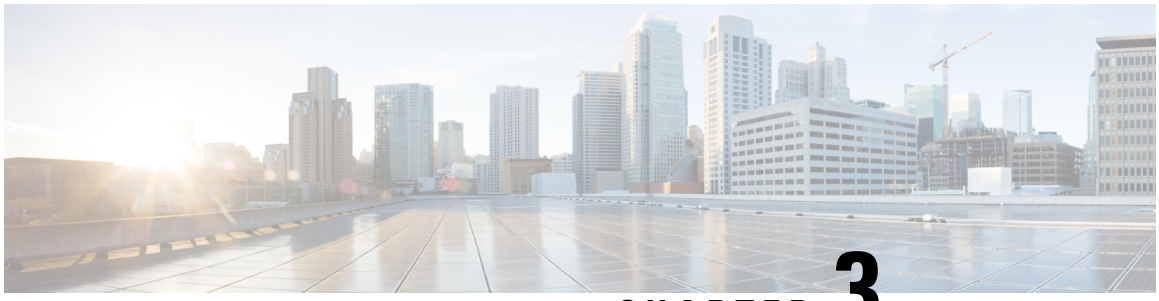
- Example: Blocking Older SSL/TLS Versions from the Network, on page 447

FlexConfig Policy

- How to Enable and Disable Global Default Inspections, on page 854
- How to Undo Your FlexConfig Changes, on page 860
- How to Enable Inspections for Unique Traffic Classes, on page 861

Virtual Routing

- How to Provide Internet Access to Multiple Virtual Routers with Overlapping Address Spaces , on page 332
- How to Route to a Distant Server through Multiple Virtual Routers, on page 327
- How to Allow RA VPN Access to Internal Networks in Different Virtual Routers, on page 737
- How to Secure Traffic from Networks in Multiple Virtual Routers over a Site-to-Site VPN, on page 665



CHAPTER 3

Licensing the System

The following topics explain how to license the Firewall Threat Defense device.

- Smart Licensing for the Firewall System, on page 81
- Managing Smart Licenses, on page 86
- Applying Permanent Licenses, on page 91

Smart Licensing for the Firewall System

Cisco Smart Licensing is a flexible licensing model that provides you with an easier, faster, and more consistent way to purchase and manage software across the Cisco portfolio and across your organization. And it's secure—you control what users can access. With Smart Licensing you get:

- **Easy Activation:** Smart Licensing establishes a pool of software licenses that can be used across the entire organization—no more PAKs (Product Activation Keys).
- **Unified Management:** Cisco License Central provides a complete view into all of your Cisco products and services in an easy-to-use portal, so you always know what you have and what you are using.
- **License Flexibility:** Your license is not node-locked to your hardware, so you can easily use and transfer licenses as needed.

To use Smart Licensing, you must first set up a Smart Account on Cisco License Central.

For a more detailed overview on Cisco Licensing, visit the Licensing Hub.

Cisco Smart Software Manager

When you purchase one or more licenses for the Firewall Threat Defense device, you manage them in the Cisco Smart Software Manager: https://software.cisco.com/#SmartLicensing-Inventory. The Cisco 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.

Licenses and appliances are managed per 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.

When you register a device with Cisco Smart Software Manager, you create a Product Instance Registration Token in the manager, and then enter it in Firewall Device Manager. A registered device becomes associated with a virtual account based on the token that is used.

For more information about the Cisco Smart Software Manager, see the online help for the manager.

Periodic Communication with the License Authority

When you use a Product Instance Registration Token to register the Firewall Threat Defense device, the device registers with the Cisco License Authority. The License Authority issues an ID certificate for communication between the device 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 device reverts to a de-registered state and licensed feature usage is suspended.

The device communicates with the License Authority on a periodic basis. If you make changes in the Cisco Smart Software Manager, you can refresh the authorization on the device so the changes immediately take effect. You also can wait for the device to communicate as scheduled. Normal license communication occurs every 12 hours, but with the grace period, your device will operate for up to 90 days without calling home. You must contact the License Authority before 90 days have passed.

Smart License Types

The following table explains the licenses available for the Firewall Threat Defense device.

Your purchase of a Firewall Threat Defense device automatically includes a Base license. All additional licenses are optional.

Table 4: Smart License Types

License	Duration	Granted Capabilities
Base	Perpetual	All features not covered by the optional term licenses. The Base license is automatically added to your account when you register. The exception is for the Secure Firewall 3100. You obtain a Base license when you purchase the firewall, and the license is managed like other licenses in your account. For example, you need to make sure the license is in the correct virtual account when you register. You must also specify whether to **Allow export-controlled functionality on the products registered with this token**. You can select this option only if your country meets export-control standards. This option controls your use of advanced encryption and the features that require advanced encryption.

License	Duration	Granted Capabilities
Threat	Term-based	Required to use the following policies: • Intrusion • File (the Malware is also required) • Security Intelligence
Malware	Term-based	File policies (the Threat is also required).
URL	Term-based	URL policies—Category and reputation-based URL filtering or DNS lookup request filtering. You can perform URL filtering on individual URLs without this license.
RA VPN: • AnyConnect Plus • AnyConnect Apex • AnyConnect VPN Only	Term-based or perpetual based on license type.	Remote access VPN configuration. Your base license must allow export-controlled functionality to configure RA VPN. You select whether you meet export requirements when you register the device. The Firewall Device Manager can use any valid AnyConnect Client license. The available features do not differ based on license type. If you have not already purchased one, see Licensing Requirements for Remote Access VPN, on page 677. Also see *Cisco AnyConnect Ordering Guide*, http://www.cisco.com/c/dam/en/us/products/collateral/security/anyconnect-og.pdf.

Firewall Threat Defense Virtual Licensing

Any Firewall Threat Defense Virtual license can be used on any supported Firewall Threat Defense Virtual vCPU/memory configuration. This allows Firewall Threat Defense Virtual customers to run on a wide variety of VM resource footprints. This also increases the number of supported AWS and Azure instances types. When configuring the Firewall Threat Defense Virtual VM, the maximum supported number of cores (vCPUs) is 16 ; and the maximum supported memory is 32 GB RAM .

Performance Tiers for Firewall Threat Defense Virtual Smart Licensing

Session limits for RA VPNs are determined by the installed Firewall Threat Defense Virtual platform entitlement tier, and enforced via a rate limiter. The following table summarizes the session limits based on the entitlement tier and rate limiter.

Table 5: Firewall Threat Defense Virtual Licensed Feature Limits Based on Entitlement

Performance Tier	Device Specifications (Core/RAM)	Rate Limit	RA VPN Session Limit
FTDv5, 100Mbps	4 core/8 GB	100Mbps	50
FTDv10, 1Gbps	4 core/8 GB	1Gbps	250

Performance Tier	Device Specifications (Core/RAM)	Rate Limit	RA VPN Session Limit
FTDv20, 3Gbps	4 core/8 GB	3Gbps	250
FTDv30, 5Gbps	8 core/16 GB	5Gbps	250
FTDv50, 10Gbps	12 core/24 GB	10Gbps	750
FTDv100, 16Gbps	16 core/32 GB	16Gbps	10,000

Firewall Threat Defense Virtual Performance Tier Licensing Guidelines and Limitations

Please keep the following guidelines and limitations in mind when licensing your Firewall Threat Defense Virtual device.

- The Firewall Threat Defense Virtual supports performance-tiered licensing that provides different throughput levels and VPN connection limits based on deployment requirements.
- Any Firewall Threat Defense Virtual license can be used on any supported Firewall Threat Defense Virtual core/memory configuration. This allows the Firewall Threat Defense Virtual customers to run on a wide variety of VM resource footprints.
- You can select a performance tier when you deploy the Firewall Threat Defense Virtual, whether your device is in evaluation mode or is already registered with Cisco Smart Software Manager.

Note Make sure your Smart Licensing account contains the available licenses you need. It's important to choose the tier that matches the license you have in your account. If you are upgrading your Firewall Threat Defense Virtual to Version 7.0, you can choose **FTDv - Variable** to maintain your current license compliance. Your Firewall Threat Defense Virtual continues to perform with session limits based on your device capabilities (number of cores/RAM).

- The default performance tier is FTDv50 when deploying a new Firewall Threat Defense Virtual device, or when provisioning the Firewall Threat Defense Virtual using the REST API.
- Base licenses are subscription-based and mapped to performance tiers. Your virtual account needs to have the Base license entitlements for the Firewall Threat Defense Virtual devices, as well as for Threat, Malware, and URL Filtering licenses.
- Each HA peer consumes one entitlement, and the entitlements on each HA peer must match, including Base license.
- A change in performance tier for an HA pair should be applied to the primary peer.
- Universal PLR licensing is applied to each device in an HA pair separately. The secondary device will not automatically mirror the performance tier of the primary device. It must be updated manually.

Impact of Export Control Setting on Encryption Features

When you register a device, you must also specify whether to **Allow export-controlled functionality on the products registered with this token**. You can select this option only if your country meets export-control standards. This option controls your use of advanced encryption and the features that require advanced encryption.

Evaluation mode is treated the same as registering using a non-export-compliant account. That means that you cannot configure remote access VPN, or use advanced encryption algorithms, when running in evaluation mode.

Most particularly, the DES standard is available only in evaluation or non-export-compliant mode.

Thus, if you configure encrypted features, such as site-to-site VPN, or encrypt the failover connection in a high availability group, you might end up with connection problems after registering in an export-compliant account. If the feature was using DES in evaluation mode, that configuration will be broken after you register the account.

Consider the following recommendations for avoiding encryption-related problems:

- Avoid configuring encrypted features, such as site-to-site VPN and encrypted failover connections, until after you register the device.
- After registering the device using an export-compliant account, edit all encrypted features that you configured in evaluation mode and select more secure encryption algorithms. Test and verify each of these features to ensure they are functioning correctly.



Note If you configured HA failover encryption in evaluation mode, you will also need to reboot both devices in the HA group to start using stronger encryption. We recommend you remove the encryption first to avoid a split-brain situation, where both devices consider themselves the active unit.

Impact of Expired or Disabled Optional Licenses

If one of the following optional licenses expires, you can continue using features that require the license. However, the license is marked out of compliance and you need to purchase the license and add it to your account to bring the license back into compliance.

If you disable an optional license, the system reacts as follows:

- Malware—The system stops querying the Secure Malware Analytics Cloud, and also stops acknowledging retrospective events sent from the Secure Malware Analytics Cloud. You cannot re-deploy existing access control policies if they include file policies. 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.
- Threat—The system no longer applies intrusion or file policies. For Security Intelligence policies, the system no longer applies the policy and stops downloading feed updates. You cannot re-deploy existing policies that require the license.
- URL—Access control rules with URL category conditions immediately stop filtering URLs or DNS lookup requests, and the system no longer downloads updates to URL data. You cannot re-deploy existing access control policies if they include rules with category and reputation-based URL conditions.

- RA VPN—You cannot edit the remote access VPN configuration, but you can remove it. Users can still connect using the RA VPN configuration. However, if you change the device registration so that the system is no longer export compliant, the remote access VPN configuration stops immediately and no remote users can connect through the VPN.

Managing Smart Licenses

Use the Smart License page to view the current license status for the system. The system must be licensed.

The page shows you whether you are using the 90-day evaluation license, or if you have registered with the Cisco Smart Software Manager. Once registered, you can see the status of the connection to the Cisco Smart Software Manager as well as the status for each type of license.

Usage Authorization identifies the Smart License Agent status:

- Authorized ("Connected," "Sufficient Licenses")—The device has contacted and registered successfully with the License Authority, which has authorized the license entitlements for the appliance. The device is now In-Compliance.
- Out-of-Compliance—There is no available license entitlement for the device. Licensed features continue to work. However, you must either purchase or free up additional entitlements to become In-Compliance.
- Authorization Expired—The device has not communicated with the Licensing Authority in 90 or more days. Licensed features continue to work. In this state, the Smart License Agent retries its authorization requests. If a retry succeeds, the agent enters either an Out-of-Compliance or Authorized state, and begins a new Authorization Period. Try manually synchronizing the device.



Note Click the **i** button next to the Smart License status to view the virtual account, export-controlled features, and get a link to open the Cisco Smart Software Manager. Export-Controlled Features control software that is subject to national security, foreign policy, and anti-terrorism laws and regulations.

The following procedure provides an overview of how to manage licenses for the system.

Before you begin

If you do not have a path to the internet for the system, you cannot use Smart Licensing. Instead, switch to Permanent License Reservation (PLR) mode. For detailed information, see Applying Permanent Licenses, on page 91.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Smart License summary.

Step 2 Register the device.

You must register with the Cisco Smart Software Manager before you can assign the optional licenses. Register before the end of the evaluation period.

See Registering the Device, on page 87.

Note
When you register, you elect whether to send usage data to Cisco. You can change your election by clicking the **Go To Cisco Success Network** link next to the gear icon.

Step 3 Request and manage the optional feature licenses.

You must register the optional licenses to use the features controlled by the license. See Enabling or Disabling Optional Licenses, on page 89.

Step 4 Maintain system licensing.

You can do the following tasks:

- Synchronizing with the Cisco Smart Software Manager, on page 90
- Unregistering the Device, on page 90

Registering the Device

Your purchase of the Firewall Threat Defense device automatically includes the Base license. The Base license covers all features not covered by the optional licenses. It is a perpetual license.

During initial system setup, you are prompted to register the device with Cisco Smart Software Manager. If you instead elected to use the 90-day evaluation license, you must register the device before the end of the evaluation period.

When you register the device, your virtual account allocates the license to the device. Registering the device also registers any optional licenses that you have enabled.

Before you begin

When you register a device, only that device is registered. If the device is configured for high availability, you must log into the other unit in the high availability pair to register that unit.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Smart License summary.

Step 2 Click **Register Device** and follow the instructions.

a) Click the link to open the Cisco Smart Software Manager and log into your account, or create a new one if necessary.

b) Generate a new token.

When you create the token, you specify the amount of time the token is valid for use. The recommended expiration period is 30 days. This period defines the expiration date of the token itself, and has no impact on the device that you register using the token. If the token expires before you can use it, you can simply generate a new token.

You must also specify whether to **Allow export-controlled functionality on the products registered with this token**. You can select this option only if your country meets export-control standards. This option controls your use of advanced encryption and the features that require advanced encryption.

c) Copy and paste the token into the edit box on the Smart License Registration dialog box.

d) (**Firewall Threat Defense Virtual only**) Select a performance tier for your Firewall Threat Defense Virtual device, or leave the default selection.

When no performance tier is selected, your Firewall Threat Defense Virtual device runs in legacy mode with default settings of 4 cores/8 GB; see Changing the Firewall Threat Defense Virtual Performance Tier, on page 88 for more information.

e) Select your region for Cisco Cloud Services registration.

After registration, if you need to change this region, you must unregister the device, then register it again and select the new region.

f) Decide whether to send usage data to Cisco.

Read the information in the Cisco Success Network step, click the **Sample Data** link to view the actual data that is collected, then decide whether to leave the **Enable Cisco Success Network** option selected.

g) Click **Register Device**.

Changing the Firewall Threat Defense Virtual Performance Tier

The Firewall Threat Defense Virtual supports performance-tiered licensing that provides different throughput levels and VPN connection limits based on deployment requirements. Any license can be used on any supported core/memory configuration. This allows the Firewall Threat Defense Virtual customers to run on a wide variety of VM resource footprints; see Performance Tiers for Firewall Threat Defense Virtual Smart Licensing, on page 83.

When you upgrade the Firewall Threat Defense Virtual to version 7.0+, the device is automatically moved to the "FTDv Variable" tier state and continues to consume non-tiered entitlements until you choose the entitlement level.

Keep the following in mind:

- You can change the performance tier to meet your deployment needs based on throughput or RA VPN requirements. Keep in mind that the Firewall Threat Defense Virtual deploys with adjustable core and memory resources. Your selected performance tier should not exceed your device specifications.
- Changing the performance tier is not supported for AWS.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Smart License summary.

Step 2 Select the desired option from the **Performance Tier** drop-down list.

- FTDv5 (4 core/8 GB)
- FTDv10 (8 core/8 GB)
- FTDv20 (8 core/8 GB)
- FTDv30 (8 core/16 GB)
- FTDv50 (12 core/24 GB)
- FTDv100 (16 core/24 GB)

Note
The system highlights the optimal tier based on the current device specifications.

Step 3 Review your selection and device specifications.

Note
When configuring the Firewall Threat Defense Virtual VM, the maximum supported number of cores (vCPUs) is 12 (16 for FTDv100 on VMware and KVM); and the maximum supported memory is 24 GB RAM. Your selected performance tier should not exceed your device specifications.

Step 4 Click **YES** to change the performance tier.

Enabling or Disabling Optional Licenses

You can enable (register) or disable (release) optional licenses. You must enable a license to use the features controlled by the license.

If you no longer want to use the features covered by an optional term license, you can disable the license. Disabling the license releases it in your Cisco Smart Software Manager account, so that you can apply it to another device.

You can also enable evaluation versions of these licenses when running in evaluation mode. In evaluation mode, the licenses are not registered with Cisco Smart Software Manager until you register the device. However, you cannot enable the RA VPN license in evaluation mode.

Before you begin

Before disabling a license, ensure that you are not using it. Rewrite or delete any policies that require the license.

For units operating in a high availability configuration, you enable or disable licenses on the active unit only. The change is reflected on the standby unit the next time you deploy the configuration, when the standby unit requests (or frees) the necessary licenses. When enabling licenses, you must ensure that your Cisco Smart Software Manager account has sufficient licenses available, or you could have one unit compliant while the other unit is non-compliant.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Smart License summary.

Step 2 Click the **Enable/Disable** control for each optional license as desired.

- **Enable**—Registers the license with your Cisco Smart Software Manager account and enables the controlled features. You can now configure and deploy policies controlled by the license.
- **Disable**—Unregisters the license with your Cisco Smart Software Manager account and disables the controlled features. You cannot configure the features in new policies, nor can you deploy policies that use the feature.

Step 3 If you enabled the **RA VPN** license, select the type of license you have available in your account.

You can use any of the AnyConnect Client licenses: **Advantage**, **Premier**, or **VPN Only**. You can select **Advantage and Premier** if you have both licenses and you want to use them both.

Synchronizing with the Cisco Smart Software Manager

The system periodically synchronizes license information with Cisco Smart Software Manager. Normal license communication occurs every 30 days, but with the grace period, your appliance will operate for up to 90 days without calling home.

However, if you make changes in the Cisco Smart Software Manager, you can refresh the authorization on the device so the changes immediately take effect.

Synchronization gets the current status of licenses, and renews authorization and the ID certificate.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Smart License summary.

Step 2 Select **Resync Connection** from the gear drop-down list.

Unregistering the Device

If you no longer want to use the device, you can unregister it from the Cisco Smart Software Manager. When you unregister, the Base license and all optional licenses associated with the device are freed in your virtual account. Optional licenses are available to be assigned to other devices. In addition, the device is unregistered from the cloud and cloud services.

After unregistering the device, the current configuration and policies on the device continue to work as-is, but you cannot make or deploy any changes.



Caution You must use this procedure to unregister the device. If you instead unregister from your Cisco Smart Software Manager account, there will be an inconsistency between the license state on the device and the one in Cisco Smart Software Manager. This will cause errors when you try to license the device again, such as when you form an HA pair. Messages symptomatic of this failure are "Failed to generate token to enroll with Cisco Cloud using Smart License" and "Could not return the certificate for the given sn (*serial_number*) since it is REVOKED." If this happens, unregister the units using this procedure and try again.

Before you begin

When you unregister a device, only that device is unregistered. If the device is configured for high availability, you must log into the other unit in the high availability pair to unregister that unit.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Smart License summary.

Step 2 Select **Unregister Device** from the gear drop-down list.

Step 3 Read the warning and click **Unregister** if you really want to unregister the device.

Applying Permanent Licenses

You can apply permanent licenses in both non-air-gapped and air-gapped networks. An air-gapped network is one in which there is no path to the Internet. These are high-security networks where you want to prevent any possibility of external entry and attack. Because there is no path to the Internet, you cannot register the device directly with the Cisco Smart Software Manager. Instead, you can use Permanent License Reservation (PLR) mode to obtain a license you can apply to the device.

If you need to use PLR mode, please keep the following in mind:

- Features that require access to the internet, such as file policies, URL Lookups, or contextual cross-launch to public web sites, will not work.
- Even if you enable Web Analytics and Cisco Success Network, Cisco does not collect the associated data due to the lack of internet access.
- You will need to manually upload updates to the Geolocation Database, Intrusion Rules, and Vulnerability Database (VDB). For example, you can download the updates to a flash drive, then take the drive into your secured building and upload them from a secured workstation.



Note Cisco Smart Software Manager uses the device's serial number to assign the permanent license. If you need to unregister the device, and the normal unregistration or cancellation processes fail to remove the license assignment, you will need to contact Cisco Technical Support to remove the registration from Cisco Smart Software Manager. Reimaging the device will not remove the license registration.

The following topics explain more about the different types of permanent license, how to apply them, and how to cancel registration or unregister the device.

Universal Permanent vs. Specific License Reservation

There are two separate types of license reservation:

- Universal Permanent License Reservation (Universal PLR or UPLR)—The Universal Permanent License permits perpetual, unlimited use of supported firewall products, including all optional licenses. Once you purchase and apply a Universal Permanent License, any applied feature licenses, which are normally time-based, are permanently applicable. However, you are still responsible for purchasing replacement licenses as they expire in your Smart License account. ISA 3000 does not support Universal PLR.
- Specific License Reservation—Specific License Reservation requires the same number and types of licenses as standard Smart Licensing. When you obtain this license, you select which optional feature

licenses you want in addition to the base license. You must periodically update your licenses as they expire.

Firewall Device Manager supports Universal PLR only.

You must work with your Cisco representative to enable Universal Permanent License Reservation (PLR) mode in your Cisco Smart Software Manager (CSSM) account.

Verify That Your Smart Account Can Provide a Universal License

To verify that you can obtain and apply a permanent license, log into your CSSM account and go to the **Smart Software Licensing** > **Inventory** page, then click the **Licenses** tab. If you can see the **License Reservation** button, then you are authorized to obtain permanent license reservations.

However, this button starts a wizard that works for both Universal and Specific Permanent Licenses.

You must also look through your list of available licenses to verify that there is a Universal License for the device. This license will appear as a selectable item in step 2 of the wizard launched by the **License Reservation** button.

If you can see the **License Reservation** button and you can get a Universal License, then you can proceed with converting the system to use a permanent license. If the button does not appear, or you can reserve Specific licenses only, call your Cisco representative and request that Universal PLR mode be enabled for your account.

Switch to PLR Mode and Apply a Universal License

Once you verify that you can obtain a permanent license, as explained in Verify That Your Smart Account Can Provide a Universal License, on page 92, and you have purchased the required Universal License, you can switch to Permanent License Reservation (PLR) mode and apply the license.

Caution If you are currently in evaluation mode, once you switch to PLR mode, you cannot switch back to evaluation mode.

Before you begin

If the device is configured for High Availability, you must complete this task separately for both devices in the HA group.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Smart License summary.

Step 2 If you have already registered the device using Smart Licensing, select **Unregister Device** from the gear drop-down list and then confirm unregistration. Wait for the unregistration task to complete before proceeding.

Step 3 Select **Switch to Universal PLR** from the gear drop-down list to switch to Universal Permanent License Reservation (PLR) mode.

Read the warning and click **Yes** to confirm the switch.

The system converts to PLR mode and then starts the PLR registration process.

Step 4 Complete the PLR registration.

a) When the system opens the Universal Permanent License Reservation dialog box, the first step includes the request code you will need. You can click **Save As TXT** to save it in a text file, or **Print** to print it out. You can also highlight the string and press Ctrl+C to copy it to the clipboard.

If you canceled out of the process after switching modes, you can restart at this point by clicking the **Continue Reservation** button on the Licensing page.

b) Log into your CSSM account, go to the **Smart Software Licensing** > **Inventory** page, and click the **Licenses** tab.

c) Click the **License Reservation** button and follow the instructions in the wizard. You will be prompted to enter the request code you generated, and in return, you will get an authorization code.

The wizard includes these steps:

1. Enter the license request code, or upload the text file that contains the code, and click **Next**.
2. In step 2, you are presented with the product details for the system you are licensing, and a bullet list of the available licenses. Select the Universal License for a locally-managed the Firewall Threat Defense device, and click **Next**.
3. In step 3, verify you have the correct license selected, and click **Generate Authorization Code**.
4. In step 4, you are presented with the authorization code. Either click **Download as File** or **Copy to Clipboard**, as appropriate, to save the code.
5. Click **Close** to exit the wizard.

d) Back in the Firewall Device Manager, paste the authorization code into the appropriate field.

A valid authorization code for a Universal License has the following format: XXXXXX-XXXXXX-XXXXXX-XXXXXX-XXXXXX-XXXX, where X is an alpha-numeric character. If your authorization code is instead an XML file, you have a Specific License and you cannot use it on this system. Please cancel the registration as described in Cancel PLR Registration, on page 93, ensuring that you release the reserved licenses in CSSM. Then, work with your Cisco representative to get your Smart Account converted to Universal PLR.

e) Click **Register**.

The system will start the registration process. Refresh the Licensing page to check the registration status.

Step 5 Enable the optional feature licenses as required.

The Universal License registers the device for the Base license only. You can now click **Enable** for each of the feature licenses you need.

Cancel PLR Registration

You can cancel a Universal Permanent License Reservation (PLR) request before it is completed. For example, if you start the PLR registration process, and discover that your Smart Software Manager account is not set

up for PLR, you can cancel the process while you get the authorization for PLR mode and your Smart License account is set up appropriately.

If you have completed the PLR registration process, you cannot cancel it. Instead, see Unregister the Device in PLR Mode, on page 94.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Smart License summary.

Step 2 Select **Cancel PLR** from the gear drop-down list to start the cancellation process.

Step 3 Select the option that fits your situation:

- **I have a license in CSSM**—Use this option if you have gone through the License Registration wizard in the Cisco Smart Software Manager (CSSM) and you have obtained an authorization code. At this point, there are licenses reserved in CSSM, and you need to release them.
- **I do not have a license in CSSM**—Use this option if you have not completed the CSSM wizard to the point where you obtained an authorization code. For example, if you started PLR registration in Firewall Device Manager but then discovered the **License Reservation** button was not available in your Smart Account.

Step 4 (If you selected **I have a license in CSSM**.) You need to obtain a release code from CSSM to ensure that your licenses are no longer marked as in-use. Otherwise, those licenses will not be usable by other devices.

a) Paste the authorization code you obtained from CSSM (when registering) into the cancellation dialog box and click **Generate Release Code**.

b) When there is a code in the **Release License Code** field, click **Save As TXT** to save it to a text file, or **Print** to print it. You can also select the code and press Ctrl+C to copy it to the clipboard.

c) In CSSM, find the device in the **Smart Software Licensing** > **Inventory** page (the Name is the device serial number), click **Action** > **Remove**, and enter the release code.

Wait for CSSM to indicate that the product was successfully removed.

Step 5 Click **OK** to complete the cancellation process.

The system returns to Smart License mode. However, the device will be unregistered, and you cannot restart evaluation mode. At this point, you must register the device using a Smart License, or switch back to PLR mode and register again, to use it.

Unregister the Device in PLR Mode

If you no longer need to license the device, for example, because you are decommissioning it or moving it to a different facility, where you will license it separately, you can unregister the device.

Unregistering the device returns the license to an unused state. If you do not unregister the device, the license remains marked as in-use and you cannot use it for other purposes.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Smart License summary.

Step 2 Select **Unregister Universal PLR** from the gear drop-down list, read the warning, and click **Yes** to start the process.

Step 3 When the Unregister Universal Permanent License Reservation dialog box opens, the **Release License Code** field is populated with the code you need to release the licenses currently assigned in your CSSM account. Click **Save as TXT** or **Print** to retain a copy of this code. You can also select it and use Ctrl+C to copy it to the clipboard.

Step 4 Go to your CSSM account, find the device in the **Smart Software Licensing** > **Inventory** page (the Name is the device serial number), click **Action** > **Remove**, and enter the release code.

Wait for CSSM to indicate that the product was successfully removed.

Step 5 Back in the Firewall Device Manager, click **Unregister** in the Unregister Device dialog box.

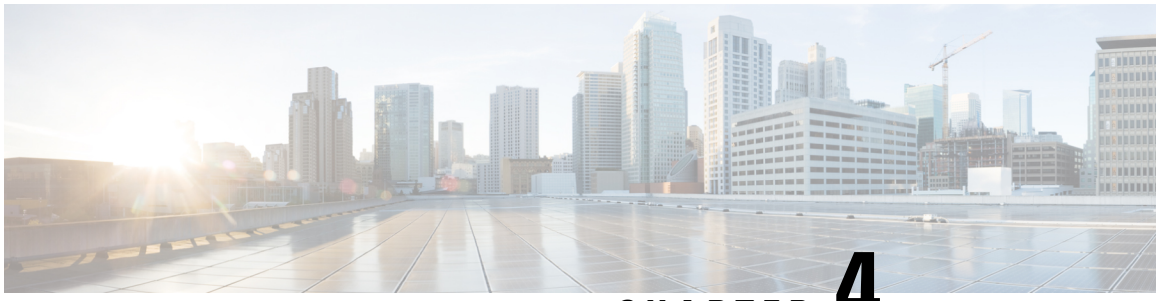
This completes the process. At this point, the licenses in CSSM are free to assign to another device, and the Firewall Threat Defense device is unlicensed.



PART I

System Monitoring

- Monitoring the Device, on page 99
- Alarms for the Cisco ISA 3000, on page 119



CHAPTER 4

Monitoring the Device

The system includes dashboards and an Event Viewer that you can use to monitor the device and traffic that is passing through the device.

- Enable Logging to Obtain Traffic Statistics, on page 99
- Monitoring Traffic and System Dashboards, on page 102
- Monitoring Additional Statistics Using the Command Line, on page 105
- Viewing Events, on page 105

Enable Logging to Obtain Traffic Statistics

You can monitor a wide range of traffic statistics using the monitoring dashboards and the Event Viewer. However, you must enable logging to tell the system which statistics to collect. Logging generates various types of events that provide insight into the connections going through the system.

The following topics explain more about events and the information they provide, with special emphasis on connection logging.

Event Types

The system can generate the following types of events. You must generate these events to see related statistics in the monitoring dashboards.

Connection Events

You can generate events for connections as users generate traffic that passes through the system. Enable connection logging on access rules to generate these events. You can also enable logging on Security Intelligence policies and SSL decryption rules to generate connection events.

Connection events include a wide variety of information about a connection, including source and destination IP addresses and ports, URLs and applications used, and the number of bytes or packets transmitted. The information also includes the action taken (for example, allowing or blocking the connection), and the policies applied to the connection.

Intrusion Events

The system examines the packets that traverse your network for malicious activity that could affect the availability, integrity, and confidentiality of a host and its data. When the system identifies a possible intrusion, it generates an intrusion event, which is a record of the date, time, type of exploit, and contextual

information about the source of the attack and its target. Intrusion events are generated for any intrusion rule set to block or alert, regardless of the logging configuration of the invoking access control rule.

File Events

File events represent files that the system detected, and optionally blocked, in network traffic based on your file policies. You must enable file logging on the access rule that applies the file policy to generate these events.

When the system generates a file event, the system also logs the end of the associated connection regardless of the logging configuration of the invoking access control rule.

Malware Events

The system can detect malware in network traffic as part of your overall access control configuration. The Malware Defense can generate a malware event, containing the disposition of the resulting event, and contextual data about how, where, and when the malware was detected. You must enable file logging on the access rule that applies the file policy to generate these events.

The disposition of a file can change, for example, from clean to malware or from malware to clean. If Malware Defense queries the AMP Cloud about a file, and the cloud determines the disposition has changed within a week of the query, the system generates retrospective malware events.

Security Intelligence Events

Security Intelligence events are a type of connection event generated by the Security Intelligence policy for each connection blocked or monitored by the policy. All Security Intelligence events have a populated Security Intelligence Category field.

For each of these events, there is a corresponding “regular” connection event. Because the Security Intelligence policy is evaluated before many other security policies, including access control, 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.

Configurable Connection Logging

You should log connections according to the security and compliance needs of your organization. If your goal is to limit the number of events you generate and improve performance, only enable logging for the connections critical to your analysis. However, if you want a broad view of your network traffic for profiling purposes, you can enable logging for additional connections.

Because the system can log a connection for multiple reasons, disabling logging in one place does not ensure that matching connections will not be logged.

You can configure connection logging in the following places.

- Access control rules and default action—Logging at the end of a connection provides the most information about the connection. You can also log the beginning of the connection, but these events have incomplete information. Connection logging is disabled by default, so you must enable it for each rule (and the default action) that targets traffic that you want to track.
- Security Intelligence policy—You can enable logging to generate Security Intelligence connection events for each blocked connection. 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.

- SSL Decryption rules and default action—You can configure logging at the end of a connection. 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.

Automatic Connection Logging

The system automatically saves the following end-of-connection events, regardless of any other logging configurations.

- The system automatically logs connections associated with intrusion events, unless the connection is handled by the access control policy's default action. You must enable logging on the default action to get intrusion events for matching traffic.
- The system automatically logs connections associated with file and malware events. This is for connection events only: you can optionally disable the generation of file and malware events.

Tips for Connection Logging

Keep the following tips in mind when considering your logging configuration and the evaluation of related statistics:

- When you allow traffic with an access control rule, you can use an associated intrusion or file policy (or both) to further inspect traffic and block intrusions, prohibited files, and malware before the traffic can reach its final destination. Note, however, that by default file and intrusion inspection is disabled for encrypted payloads. If the intrusion or file policies find reason to block a connection, the system immediately logs an end-of-connection event regardless of your connection log settings. Logging allowed connections provides the most statistical information on the traffic in your network.
- A trusted connection is one that is handled by a Trust access control rule or the default action in an access control policy. However, trusted connections are not inspected for discovery data, intrusions, or prohibited files and malware. Therefore, connection events for trusted connections contain limited information.
- For access control rules and access control policy default actions that block traffic, the system logs beginning-of-connection events. Matching traffic is denied without further inspection.
- 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 a Block rule, consider whether the rule monitors traffic on an Internet-facing interface or other interface vulnerable to DoS attack.
- If you select the **Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)** option when you configure remote access VPN connection profiles, or you otherwise enable the **sysopt connection permit-vpn** command, all site-to-site or remote access VPN traffic bypasses inspection and the access control policy. Thus, you will get no connection events for this traffic, and the traffic will not be reflected in any statistical dashboards.

Sending Events to an External Syslog Server

Besides viewing events through the Firewall Device Manager, which has a limited capacity to store events, you can selectively configure rules and policies to send events to an external syslog server. You can then use the features and additional storage of your selected syslog server platform to view and analyze event data.

To send events to an external syslog server, edit each rule, default action, or policy that enables connection logging and select a syslog server object in the log settings. To send intrusion events to a syslog server, configure the server in the intrusion policy settings. To send file/malware events to a syslog server, configure the server on **Device** > **System Settings** > **Logging Settings**.

For more information, see the help for each rule and policy type and also see Configuring Syslog Servers, on page 143.

Evaluating Events Using Cisco Cloud-Based Services

In addition to using Event Viewer and your own syslog servers, you can send connection events, and high-priority intrusion, file, and malware events, to a Cisco cloud-based server. Cisco cloud-based services, such as Threat Response, can pull the events from that cloud server and you can then use those services to evaluate these events.

These cloud-based services are separate from the Firewall Threat Defense device and Firewall Device Manager. If you elect to use a service that requires you to send these events to the Cisco cloud, you must enable the connection on the **Device** > **System Settings** > **Cloud Services** page. See Sending Events to the Cisco Cloud, on page 775.

Monitoring Traffic and System Dashboards

The system includes several dashboards that you can use to analyze the traffic going through the device and the results of your security policy. Use the information to evaluate the overall efficacy of your configuration and to identify and resolve network problems.

The dashboards for units in a high availability group show statistics for that device only. Statistics are not synchronized among the units.



Note The data used in traffic-related dashboards is collected from access control rules that enable connection or file logging, and other security policies that allow logging. The dashboards do not reflect traffic that matches rules for which no logging is enabled. Ensure that you configure your rules to log the information that matters to you. In addition, user information is available only if you configure identity rules to collect user identity. And finally, intrusion, file, malware, and URL category information is available only if you have a license for those features and configure rules that use the features.

Procedure

Step 1 Click **Monitoring** in the main menu to open the Dashboards page.

You can select predefined time ranges, such as the last hour or week, or define a custom time range with specific start and end times, to control the data shown in the dashboard graphs and tables.

Traffic-related dashboards include the following types of display:

- Top 5 bar graphs—These are shown in the **Network Overview** dashboard, and in the per-item summary dashboards you see if you click on an item in a dashboard table. You can toggle the information between a count of **Transactions** or **Data Usage** (total bytes sent and received). You can also toggle the display to show all transactions, allowed transactions, or denied transactions. Click the **View More** link to see the table associated with the graph.
- Tables—Tables show items of a particular type (for example, applications or URL categories) with that item's total transactions, allowed transactions, blocked transactions, data usage, and bytes sent and received. You can toggle the numbers between raw **Values** and **Percentages**, and show the top 10, 100, or 1000 entries. If the item is a link, click it to see a summary dashboard with more detailed information.

Step 2 Click the **Dashboard** links in the table of contents to see dashboards for the following data:

- **Network Overview**—Shows summary information about the traffic in the network, including the access rules (policies) matched, users initiating traffic, applications used in connections, intrusion threats (signatures) matched, URL categories for URLs accessed, and the most frequent destinations for connections.
- **Users**—Shows the top users of your network. You must configure identity policies to see user information. If there is no user identity, the source IP address is included. You might see the following special entities:
  - **Failed Authentication**—The user was prompted to authenticate, but failed to enter a valid username/password pair within the maximum number of allowed attempts. Failure to authenticate does not itself prevent the user from accessing the network, but you can write an access rule to limit network access for these users.
  - **Guest**—Guest users are like Failed Authentication users, except that your identity rule is configured to call these users Guest. Guest users were prompted to authenticate and failed to do so within the maximum number of attempts.
  - **No Authentication Required**—The user was not prompted to authentication, because the user's connections matched identity rules that specified no authentication.
  - **Unknown**—There is no user mapping for the IP address, and there is no record of failed authentication yet. Typically, this means that no HTTP traffic has yet been seen from that address.
- **Applications**—Shows the top applications, such as HTTP, that are being used in the network. The information is available only for connections that are inspected. Connections are inspected if they match an “allow” rule, or a block rule that uses criteria other than zone, address, and port. Thus, application information is not available if the connection is trusted or blocked prior to hitting any rule that requires inspection.
- **Web Applications**—Shows the top web applications, such as Google, that are being used in the network. The conditions for collecting web application information are the same as those for the Application dashboard.
- **URL Categories**—Shows the top categories of web sites, such as Gambling or Educational Institutions, that are being used in the network based on the categorization of web sites visited. You must have at least one access control rule that uses URL category as a traffic matching criteria to get this information. The information will be available for traffic that matches the rule, or for traffic that has to be inspected

to determine if it matches the rule. You will not see category (or reputation) information for connections that match rules that come before the first web-category access control rule.

- **Access And SI Rules**—Shows the top access rules and Security Intelligence rule-equivalents matched by network traffic.
- **Zones**—Shows the top security zone pairs for traffic entering and then exiting the device.
- **Destinations**—Shows the top destinations for network traffic.
- **Attackers**—Shows the top attackers, which are the source of connections that trigger intrusion events. You must configure intrusion policies on access rules to see this information.
- **Targets**—Shows the top targets of intrusion events, which are the victims of an attack. You must configure intrusion policies on access rules to see this information.
- **Threats**—Shows the top intrusion rules that have been triggered. You must configure intrusion policies on access rules to see this information.
- **File Logs**—Shows the top file types seen in network traffic. You must configure file policies on access rules to see this information.
- **Malware**—Shows the top Malware action and disposition combinations. You can drill down to see information on the associated file types. You must configure file policies on access rules to see this information.
  - Possible actions are: Malware Cloud Lookup, Block, Archive Block (Encrypted), Detect, Custom Detection, Cloud Lookup Timeout, Malware Block, Archive Block (Depth Exceeded), Custom Detection Block, TID block, Archive Block (Failed to Inspect).
  - Possible dispositions are: Malware, Unknown, Clean, Custom Detection, Unavailable.
- **SSL Decryption**—Shows the breakdown of encrypted vs. plain text traffic through the device, plus the breakdown of how encrypted traffic was decrypted according to SSL decryption rules.
- **System**— Shows an overall system view, including a display of interfaces and their status (mouse over an interface to see its IP addresses), overall average system throughput (in 5 minute buckets for up to one hour, and one hour buckets for longer periods), and summary information on system events, CPU usage, memory usage, and disk usage. You can restrict the throughput graph to show a specific interface rather than all interfaces.

Note
The information shown on the System dashboard is at the overall system level. If you log into the device CLI, you can use various commands to see more detailed information. For example, the **show cpu** and **show memory** commands include parameters for showing other details, whereas these dashboards show data from the **show cpu system** and **show memory system** commands.

Step 3 You can also click these links in the table of contents:

- **Events**—To view events as they occur. You must enable connection logging in individual access rules to see connection events related to those rules. Also, enable logging in the Security Intelligence policy and SSL decryption rules to see Security Intelligence events and additional connection event data. These events can help you resolve connection problems for your users.

- **Sessions**—To view and manage the Firewall Device Manager user sessions. For more information, see Managing the Firewall Device Manager User Sessions, on page 812.

Monitoring Additional Statistics Using the Command Line

The Firewall Device Manager dashboards provide a wide variety of statistics related to the traffic going through the device and general system usage. However, you can get additional information on areas not covered by the dashboards using the CLI Console or by logging into the device CLI (see Logging Into the Command Line Interface (CLI), on page 15).

The CLI includes a variety of **show** commands to provide these statistics. You can also use the CLI for general troubleshooting, including commands such as **ping** and **traceroute**. Most **show** commands have companion **clear** commands to reset statistics to 0. (You cannot clear statistics from the CLI Console.)

You can find documentation for the commands in *Cisco Secure Firewall Threat Defense Command Reference*, http://www.cisco.com/c/en/us/td/docs/security/firepower/command_ref/b_Command_Reference_for_Firepower_Threat_Defense.html.

For example, you might find the following commands generally useful.

- **show nat** displays hit counts for your NAT rules.
- **show xlate** displays the actual NAT translations that are active.
- **show conn** provides information about current connections going through the device.
- **show dhcpd** provides information about the DHCP servers you configure on the interfaces.
- **show interface** provides usage statistics for each interface.

Viewing Events

You can view events that are generated from your security policies that enable logging. Events are also generated for intrusion and file policies that are triggered.

The event viewer table shows the events generated in real time. As new events are generated, older events are rolled out of the table.

Before you begin

Whether events of particular types are generated depends on the following in addition to connections that match the related policy:

- Connection events—An access rule must enable connection logging. You can also enable connection logging in the Security Intelligence policy and in SSL decryption rules.
- Intrusion events—An access rule must apply an intrusion policy.
- File and Malware events—An access rule must apply a file policy and enable file logging.
- Security Intelligence events—You must enable and configure the Security Intelligence policy, and enable logging.

Procedure

Step 1 Click **Monitoring** in the main menu.

Step 2 Select **Events** from the table of contents.

The event viewer organizes events on tabs based on event types. For more information, see Event Types, on page 99.

Step 3 Click the tab that shows the type of event you want to view.

You can do the following with the event list:

- Click **Pause** to stop the addition of new events so that you can more easily find and analyze an event. Click **Resume** to allow new events to appear.
- Select a different refresh rate (5, 10, 20, or 60 seconds) to control how fast new events are shown.
- Create a custom view that includes the columns you want. To create a custom view, either click the **+** button in the tab bar, or click **Add/Remove Columns**. You cannot change the pre-set tabs, so adding or removing columns creates a new view. For more information, see Configuring Custom Views, on page 106.
- To change the width of a column, click and drag the column heading divider to the desired width.
- Mouse over an event and click **View Details** to see complete information on an event. For a description of the various fields in an event, see Event Field Descriptions, on page 108.

Step 4 If necessary, apply a filter to the table to help you locate the desired events based on various event attributes.

To create a new filter, either manually type in the filter by selecting atomic elements from the drop-down list and entering the filter value, or build a filter by clicking a cell in the events table that includes a value on which you want to filter. You can click multiple cells in the same column to create an OR condition among the values, or click cells in different columns to create an AND condition among the columns. If you build the filter by clicking cells, you can also edit the resulting filter to fine-tune it. For detailed information about creating filter rules, see Filtering Events, on page 107.

Once you build the filter, do any of the following:

- To apply the filter and update the table to show only those events that match the filter, click the **Filter** button.
- To clear an entire filter that you have applied and return the table to a non-filtered state, click **Reset Filters** in the **Filter** box.
- To clear one of the atomic elements of a filter, mouse over the element and click the **X** for the element. Then, click the **Filter** button.

Configuring Custom Views

You can create your own custom views so that you can easily see the columns you want when viewing events. You can also edit or delete custom views, although you cannot edit or delete the pre-defined views.

Procedure

Step 1 Select **Monitoring** > **Events**.

Step 2 Do one of the following:

- To create a new view based on an existing custom (or pre-defined) view, click the tab for the view, then click the + button to the left of the tabs.
- To edit an existing custom view, click the tab for the view.

Note
To delete a custom view, simply click the **X** button in the view's tab. You cannot undo a delete.

Step 3 Click the **Add/Remove Columns** link above the events table on the right, and select or deselect columns until the selected list includes only those columns to include in the view.

Click and drag columns between the available (but not used) and selected lists. You can also click and drag columns in the selected list to change the left-to-right order of the columns in the table. For a description of the columns, see Event Field Descriptions, on page 108.

When finished, click **OK** to save your column changes.

Note
If you change column selection while viewing a pre-defined view, a new view is created.

Step 4 If necessary, change column widths by clicking and dragging the column separators.

Filtering Events

You can create complex filters to limit the events table to the events that currently interest you. You can use the following techniques, alone or in combination, to build a filter:

Clicking columns

The easiest way to build a filter is to click on cells in the events table that contain the values on which you intend to filter. Clicking a cell updates the **Filter** field with a correctly-formulated rule for that value and field combination. However, using this technique requires that the existing list of events contains the desired values.

You cannot filter on all columns. If you can filter on the contents of a cell, it is underlined when you mouse over it.

Selecting atomic elements

You can also build a filter by clicking in the **Filter** field and selecting the desired atomic element from the drop-down list, then typing in the match value. These elements include event fields that are not shown as columns in the events table. They also include operators to define the relationship between the value you type in and the events to display. Whereas clicking columns always results in an "equals (=)" filter, when you select an element, you can also select "greater than (>)" or "less than (<)" for numeric fields.

Regardless of how you add an element to the **Filter** field, you can type into the field to adjust the operator or value. Click **Filter** to apply the filter to the table.

Operators for Event Filters

You can use the following operators in an event filter:

=	Equals. The event matches the specified value. You cannot use wildcards.
!=	Not equals. The event does not match the specified value. You must type in the ! (exclamation point) to build a not-equals expression.
>	Greater than. The event contains a value that is greater than the specified value. This operator is available for numeric values only, such as port and IP address.
<	Less than. The event contains a value that is less than the specified value. This operator is available for numeric values only.

Rules for Complex Event Filters

When building a complex filter that contains more than one atomic element, keep the following rules in mind:

- Elements of the same type have an OR relationship between all values for that type. For example, including Initiator IP=10.100.10.10 and Initiator IP=10.100.10.11 matches events that have either of these addresses as the traffic source.
- Elements of different types have an AND relationship. For example, including Initiator IP=10.100.10.10 and Destination Port/ICMP Type=80 matches events that have this source address AND destination port only. Events from 10.100.10.10 to a different destination port are not shown.
- Numeric elements, including IPv4 and IPv6 addresses, can specify ranges. For example, you could specify Destination Port=50-80 to capture all traffic for ports within this range. Use a hyphen to separate the start and end numbers. Ranges are not allowed for all numeric fields, for example, you cannot specify an IP address range in the Source element.
- You cannot use wildcards or regular expressions.

Event Field Descriptions

Events can contain the following information. You can see this information when you view event details. You can also add columns to the Event Viewer table to show the information that most interests you.

Following is a complete list of the available fields. Not every field applies to every type of event. Keep in mind that the information available for any individual event can vary depending on how, why, and when the system logged the connection.

Action

For connection or security intelligence events, the action associated with the access control rule or default action that logged the connection:

Allow

Explicitly allowed connections.

Trust

Trusted connections. TCP connections detected by a trust rule on the first packet only generate an end-of-connection event. The system generates the event one hour after the final session packet.

Block

Blocked connections. The **Block** action can be associated with Allow access rules under the following conditions:

- Connections where an exploit was blocked by an intrusion policy.
- Connections where a file was blocked by a file policy.
- Connections blocked by Security Intelligence.
- Connections blocked by an SSL policy.

Default Action

The connection was handled by the default action.

For file or malware events, the file rule action associated with the rule action for the rule the file matched, and any associated file rule action options.

Allowed Connection

Whether the system allowed the traffic flow for the event.

Application

The application detected in the connection.

Application 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.

Application Categories, Application 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.

Block Type

The type of block specified in the access control rule matching the traffic flow in the event: block or interactive block.

Client Application, Client Version

The client application and version of that client detected in the connection.

Client Business Relevance

The business relevance associated with the client traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of client detected in the connection has an associated business relevance; this field displays the lowest (least relevant) of those.

Client Category, Client Tag

Criteria that characterize the application to help you understand the application's function.

Client Risk

The risk associated with the client traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of client detected in the connection has an associated risk; this field displays the highest of those.

Connection

The unique ID for the traffic flow, internally generated.

Connection Blocktype Indicator

The type of block specified in the access control rule matching the traffic flow in the event: block or interactive block.

Connection Bytes

The total bytes for the connection.

Connection Time

The time for the beginning of the connection.

Connection Timestamp

The time the connection was detected.

Denied Connection

Whether the system denied the traffic flow for the event.

Destination Country and Continent

The country and continent of the receiving host.

Destination IP

The IP address used by the receiving host in an intrusion, file, or malware event.

Destination Port/ICMP Code; Destination Port; Destination Icode

The port or ICMP code used by the session responder.

Destination Security Group Tag, Destination Security Group Tag Name

The TrustSec security group tag number and name associated with the destination, if any.

Direction

The direction of transmission for a file.

Disposition

The file's disposition:

Malware

Indicates that the AMP Cloud categorized the file as malware or the file’s threat score exceeded the malware threshold defined in the file policy. Local malware analysis can also mark files as malware.

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 that a Detect Files or Block Files rule handled the file and the system did not query the AMP Cloud.

Egress Interface, Egress Security Zone

The interface and zone through which the connection exited the device.

Egress Virtual Router

The name of the virtual router, if any, to which the destination interface belongs.

Event, Event Type

The type of event.

Event Seconds, Event Microseconds

The time, in seconds or microseconds, when the event was detected.

File Category

The general categories of file type, for example: Office Documents, Archive, Multimedia, Executables, PDF files, Encoded, Graphics, or System Files.

File Event Timestamp

The time and date the file or malware file was created.

File Name

The name of the file.

File Rule Action

The action associated with file policy rule that detected the file, and any associated file rule action options.

File SHA-256

The SHA-256 hash value of the file.

File Size (KB)

The size of the file, in kilobytes. File size can be blank in cases where the system blocked the file before it was completely received.

File Type

The type of file, for example, HTML or MSEXE.

File/Malware Policy

The file policy associated with the generation of the event.

Filelog Blocktype Indicator

The type of block specified in the file rule matching the traffic flow in the event: block or interactive block.

Firewall Policy Rule, Firewall Rule

The access control rule or default action that handled the connection.

First Packet

The date and time the first packet of the session was seen.

HTTP Referrer

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

The HTTP status code sent in response to a client's HTTP request over a connection.

IDS Classification

The classification where the rule that generated the event belongs.

Ingress Interface, Ingress Security Zone

The interface and zone through which the connection entered the device.

Ingress Virtual Router

The name of the virtual router, if any, to which the source interface belongs.

Initiator Bytes, Initiator Packets

The total number of bytes or packets transmitted by the session initiator.

Initiator Country and Continent

The country and continent of the host that initiated the session. Available only if the initiator IP address is routable.

Initiator IP

The host IP address (and hostname, if DNS resolution is enabled) that initiated the session in a connection or Security Intelligence event.

Inline Result

Whether the system dropped or would have dropped the packet that triggered an intrusion event if operating in inline mode. Blank indicates that the triggered rule was not set to Drop and Generate Events

Intrusion Policy

The intrusion policy where the rule that generated the event was enabled.

IPS Blocktype Indicator

The action of the intrusion rule matching the traffic flow in the event.

Last Packet

The date and time the last packet of the session was seen.

MPLS Label

The Multiprotocol Label Switching label associated with the packet that triggered this intrusion event.

Malware Blocktype Indicator

The type of block specified in the file rule matching the traffic flow in the event: block or interactive block.

Message

For intrusion events, the explanatory text for the event. For malware or file events, any additional information associated with the malware event.

NAT Destination IP

For packets subject to Network Address Translation (NAT), the translated destination IP address.

NAT Destination Port

For packets subject to Network Address Translation (NAT), the translated destination port.

NAT Source IP

For packets subject to Network Address Translation (NAT), the translated source IP address.

NAT Source Port

For packets subject to Network Address Translation (NAT), the translated source port.

NetBIOS Domain

The NetBIOS domain used in the session.

Original Client Country and Continent

The country and continent of the original client host that initiated the session. Available only if the original client IP address is routable.

Original Client IP

The original IP address of the client that initiated an HTTP connection. This address is derived from the X-Forwarded-For (XFF) or True-Client-IP HTTP header fields or their equivalent.

Policy, Policy Revision

The access control policy, and its revision, that includes the access (firewall) rule associated with the event.

Priority

The event priority as determined by the Cisco Talos Intelligence Group (Talos): high, medium, or low.

Protocol

The transport protocol used in the connection.

Reason

The reason or reasons the connection was logged, in the situations explained in the following table. This field is otherwise empty.

Reason	Description
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.
Elephant Flow	The connection is large enough to be considered an elephant flow, which is a flow that can be large enough to affect overall system performance. By default, elephant flows are those larger than 1GB/10 seconds. You can adjust the byte and time thresholds for identifying elephant flows in the device CLI using the **system support elephant-flow-detection** command.
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.
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.
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.
SSL Block	The system blocked an encrypted connection based on the 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.

Receive Times

The date and time the event was generated.

Referenced Host

If the protocol in the connection is HTTP or HTTPS, this field displays the hostname that the respective protocol was using.

Responder Bytes, Responder Packets

The total number of bytes or packets transmitted by the session responder.

Responder Country and Continent

The country and continent of the host that responded to the session. Available only if the responder IP address is routable.

Responder IP

The host IP address (and hostname, if DNS resolution is enabled) of the session responder in a connection or Security Intelligence event.

SI Category ID (Security Intelligence Category)

The name of the object that contained the blocked item, such as a network or URL object name, or the name of a feed category.

Signature

The signature ID for a file/malware event.

Source Country and Continent

The country and continent of the sending host. Available only if the source IP address is routable.

Source IP

The IP address used by the sending host in an intrusion, file, or malware event.

Source Port/ICMP Type; Source Port; Source Port Itype

The port or ICMP type used by the session initiator.

Source Security Group Tag, Source Security Group Tag Name

The TrustSec security group tag number and name associated with the source, if any.

SSL Actual Action

The actual action that the system applied to the connection. This can differ from the expected action. For example, a connection might match a rule that applies decryption, but could not be decrypted for some reason.

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.

Action	Description
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 Fingerprint

The SHA hash value used to authenticate the certificate.

SSL Certificate Status

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

The cipher suite used in the connection.

SSL Expected Action

The action specified in the SSL rule the connection matched.

SSL Flow Flags

The first ten debugging level flags for an encrypted connection.

SSL Flow Messages

The SSL/TLS messages exchanged between client and server during the SSL handshake, such as HELLO_REQUEST and CLIENT_HELLO. See http://tools.ietf.org/html/rfc5246 for more information about the messages exchanged in TLS connections.

SSL Policy

The name of the SSL Decryption policy applied to the connection.

SSL Rule

The name of the SSL Decryption rule applied to the connection.

SSL Session ID

The hexadecimal Session ID negotiated between the client and server during the SSL handshake.

SSL Ticket ID

A hexadecimal hash value of the session ticket information sent during the SSL handshake.

SSL URL Category

The URL category of the destination web server as determined during SSL decryption processing.

SSL Version

The SSL/TLS version used in the connection.

TCP Flags

The TCP flags detected in the connection.

Total Packets

The total number of packets transmitted in the connection, which is **Initiator Packets** + **Responder Packets**.

URL, URL Category, URL Reputation, URL Reputation Score

The URL requested by the monitored host during the session and its associated category, reputation, and reputation score, if available.

For DNS lookup request filtering, the category and reputation are for the FQDN shown in the DNS Query field. The URL field will be blank, because the category/reputation lookup is being done for a DNS request rather than a web request.

If the system identifies or blocks an SSL application, the requested URL is in encrypted traffic, so the system identifies the traffic based on an SSL certificate. For SSL applications, therefore, the URL indicates the common name contained in the certificate.

User

The user associated with the initiator IP address.

VLAN

The innermost VLAN ID associated with the packet that triggered the event.

Web App Business Relevance

The business relevance associated with the web application traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of web application detected in the connection has an associated business relevance; this field displays the lowest (least relevant) of those.

Web App Categories, Web App Tag

Criteria that characterize the web application to help you understand the web application's function.

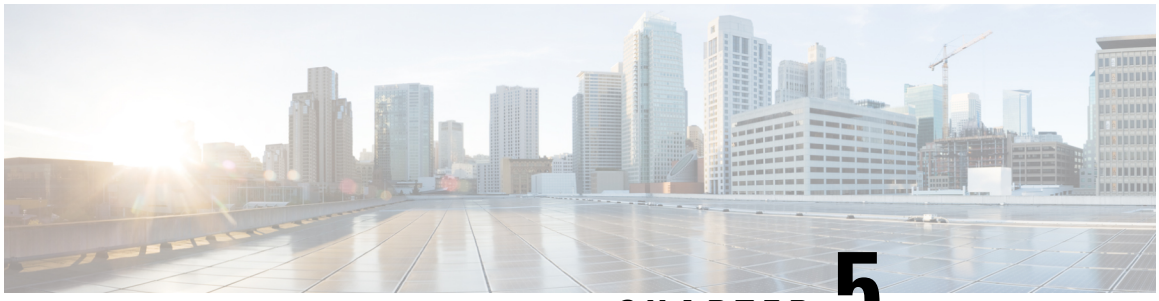
Web App Risk

The risk associated with the web application traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of web application detected in the connection has an associated risk; this field displays the highest of those.

Web Application

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.



CHAPTER 5

Alarms for the Cisco ISA 3000

You can configure the alarm system on a Cisco ISA 3000 device to alert you when undesirable conditions occur.

- About alarms, on page 119
- Defaults for alarms, on page 121
- Configuring Alarms for the ISA 3000, on page 122
- Monitoring alarms, on page 127

About alarms

You can configure the ISA 3000 to issue alarms for a variety of conditions. If any conditions do not match the configured settings, the system triggers an alarm, which is reported by way of LEDs, syslog messages, SNMP traps, and through external devices connected to the alarm output interface. By default, triggered alarms issue syslog messages only.

You can configure the alarm system to monitor the following:

- Power supply.
- Primary and secondary temperature sensors.
- Alarm input interfaces.

The ISA 3000 has internal sensors plus two alarm input interfaces and one alarm output interface. You can connect external sensors, such as door sensors, to the alarm inputs. You can connect external alarm devices, such as buzzers or lights, to the alarm output interface.

The alarm output interface is a relay mechanism. Depending on the alarm conditions, the relay is either energized or de-energized. When it is energized, any device connected to the interface is activated. A de-energized relay results in the inactive state of any connected devices. The relay remains in an energized state as long as alarms are triggered.

For information about connecting external sensors and the alarm relay, Refer to Alarms for the Cisco ISA 3000 chapter in Cisco ASA Series General Operations CLI Configuration Guide. see Cisco ISA 3000 Industrial Security Appliance Hardware Installation Guide.

Alarm input interfaces

You can connect the alarm input interfaces (or contacts) to external sensors, such as one that detects if a door is open.

Each alarm input interface has a corresponding LED. These LEDs convey the alarm status of each alarm input. You can configure the trigger and severity for each alarm input. In addition to the LED, you can configure the contact to trigger the output relay (to activate an external alarm), to send syslog messages, and to send SNMP traps.

The following table explains the statuses of the LEDs in response to alarm conditions for the alarm inputs. It also explains the behavior for the output relay, syslog messages, and SNMP traps, if you enable these responses to the alarm input.

Alarm Status	LED	Output Relay	Syslog	SNMP Trap
Alarm not configured	Off	—	—	—
No alarms triggered	Solid green	—	—	—
Alarm activated	Minor alarm—solid red Major alarm—flashing red	Relay energized	Syslog generated	SNMP trap sent
Alarm end	Solid green	Relay de-energized	Syslog generated	—

Alarm output interface

You can connect an external alarm, such as a buzzer or light, to the alarm output interface.

The alarm output interface functions as a relay and also has a corresponding LED, which conveys the alarm status of an external sensor connected to the input interface, and internal sensors such as the dual power supply and temperature sensors. You configure which alarms should activate the output relay, if any.

The following table explains the statuses of the LEDs and output relay in response to alarm conditions. It also explains the behavior for syslog messages, and SNMP traps, if you enable these responses to the alarm.

Alarm Status	LED	Output Relay	Syslog	SNMP Trap
Alarm not configured	Off	—	—	—
No alarms triggered	Solid green	—	—	—
Alarm activated	Solid red	Relay energized	Syslog generated	SNMP trap sent
Alarm end	Solid green	Relay de-energized	Syslog generated	—

Syslog Alarms

By default, the system sends syslog messages when any alarm is triggered. You can disable syslog messaging if you do not want the messages.

For syslog alarms to work, you must also enable diagnostic logging on **Device** > **System Settings** > **Logging Settings**. Configure a syslog server, console logging, or internal buffer logging.

Without enabling a destination for diagnostic logging, the alarm system has nowhere to send syslog messages.

SNMP Trap Alarms

You can optionally configure the alarms to send SNMP traps to your SNMP server. For SNMP trap alarms to work, you must also configure SNMP settings.

Use the Firewall Threat Defense API to configure SNMP. Click the more options button ( ), and choose API Explorer. Then, look for the SNMP resource and examine the model documentation for information on how to configure the feature. You can use SNMP versions 2c or 3; version 1 is not supported. For complete information on configuring SNMP, see the SNMP chapter of the *CLI Book 1: Cisco ASA Series General Operations CLI Configuration Guide* for the newest version of the ASA software. The guides are available at https://www.cisco.com/c/en/us/support/security/asa-5500-series-next-generation-firewalls/products-installation-and-configuration-guides-list.html.

Defaults for alarms

The following table specifies the defaults for alarm input interfaces (contacts), redundant power supply, and temperature.

	Alarm	Trigger	Severity	SNMP Trap	Output Relay	Syslog Message
Alarm Contact 1	Enabled	Closed State	Minor	Disabled	Disabled	Enabled
Alarm Contact 2	Enabled	Closed State	Minor	Disabled	Disabled	Enabled
Redundant Power Supply (when enabled)	Enabled	—	—	Disabled	Disabled	Enabled

	Alarm	Trigger	Severity	SNMP Trap	Output Relay	Syslog Message
Temperature	Enabled for the primary temperature alarm (default values of 92°C and -40°C for the high and low thresholds respectively) Disabled for the secondary alarm.	—	—	Enabled for primary temperature alarm	Enabled for primary temperature alarm	Enabled for primary temperature alarm

Configuring Alarms for the ISA 3000

You use FlexConfig to configure alarms for the ISA 3000. The following topics explain how to configure the different types of alarms.

Configure Alarm Input Contacts

If you connect the alarm input contacts (interfaces) to external sensors, you can configure the contacts to issue alarms based on the input from the sensor. In fact, the contacts are enabled by default to send syslog messages if the contact is closed, that is, if the electrical current stops flowing through the contact. You need to configure the contact only if the defaults do not meet your requirements.

The alarm contacts are numbered 1 and 2, so you need to understand how you have wired the physical pins to configure the correct settings. You configure the contacts separately.

Procedure

Step 1 Click **View Configuration** in **Device** > **Advanced Configuration**.

Step 2 Click **FlexConfig** > **FlexConfig Objects** in the Advanced Configuration table of contents.

Step 3 Click the + button to create a new object.

Step 4 Enter a name for the object. For example, **Enable_Alarm_Contact**.

Step 5 In the **Template** editor, enter the commands needed to configure the contact.

a) Configure a description for the alarm contact.

alarm contact {**1** | **2**} **description** *string*

For example, to set the description of contact 1 to "Door Open," enter the following:

```
alarm contact 1 description Door Open
```

b) Configure the severity for the alarm contact.

alarm contact {**1** | **2** | **any**} **severity** {**major** | **minor** | **none**}

Instead of configuring one contact, you can specify **any** to change the severity for all contacts. The severity controls the behavior of the LED associated with the contact.

- **major**—The LED blinks red.
- **minor**—The LED is solid red. This is the default.
- **none**—The LED is off.

For example, to set the severity of contact 1 to Major, enter the following:

```
alarm contact 1 severity major
```

c) Configure the trigger for the alarm contact.

alarm contact {**1** | **2** | **any**} **trigger** {**open** | **closed**}

Instead of configuring one contact, you can specify **any** to change the trigger for all contacts. The trigger determines the electrical condition that signals an alert.

- **open**—The normal condition for the contact is closed, that is, the electrical current is running through the contact. An alert is triggered if the contact becomes open, that is, the electrical current stops flowing.
- **closed**—The normal condition for the contact is open, that is, the electrical current does not run through the contact. An alert is triggered if the contact becomes closed, that is, the electrical current starts running through the contact. This is the default.

For example, you connect a door sensor to alarm input contact 1, and its normal state has no electrical current flowing through the alarm contact (it is open). If the door is opened, the contact is closed and electrical current flows through the alarm contact. You would set the alarm trigger to closed so that the alarm goes off when the electrical current starts flowing.

```
alarm contact 1 trigger closed
```

d) Configure the actions to take when the alarm contact is triggered.

alarm facility input-alarm {**1** | **2**} {**relay** | **syslog** | **notifies**}

You can configure more than one action. For example, you can configure the device to activate the external alarm, send syslog messages, and also send SNMP traps.

- **relay**—Energize the alarm output relay, which activates the external alarm that you attached to it, such as a buzzer or a flashing light. The output LED also goes red.
- **syslog**—Send a syslog message. This option is enabled by default.
- **notifies**—Send an SNMP trap.

For example, to enable all actions for the alarm input contact 1, enter the following:

```
alarm facility input-alarm 1 relay
alarm facility input-alarm 1 syslog
alarm facility input-alarm 1 notifies
```

Step 6 In the **Negate Template** editor, enter the lines required to undo this configuration.

All of these commands take the **no** form to disable them and return to default settings. For example, if your template includes all of the command examples shown in this procedure, the negate template would be the following:

```
no alarm contact 1 description Door Open
no alarm contact 1 severity major
no alarm contact 1 trigger closed
no alarm facility input-alarm 1 relay
no alarm facility input-alarm 1 syslog
no alarm facility input-alarm 1 notifies
```

Step 7 Click **OK** to save the object.

Step 8 Add the object to the FlexConfig policy.

a) Click **FlexConfig Policy** in the table of contents.

b) Click **+** in the Group List.

c) Select the Enable_Alarm_Contact object and click **OK**.

The preview should update with the commands in the template. Verify you are seeing the expected commands.

d) Click **Save**.

You can now deploy the policy.

Step 9 After deployment completes, in CLI Console or an SSH session, use the **show running-config** command and verify that the running configuration has the correct changes. Test the external sensor to verify that alarms are getting triggered.

Configure Power Supply Alarms

The ISA 3000 has two power supplies. By default, the system operates in single-power mode. However, you can configure the system to operate in dual mode, where the second power supply automatically provides power if the primary power supply fails. When you enable dual-mode, the power supply alarm is automatically enabled to send syslog alerts, but you can disable the alert altogether, or also enable SNMP traps or the alarm hardware relay.

The following procedure explains how to enable dual mode, and how to configure the power supply alarms.

Procedure

Step 1 Click **View Configuration** in **Device** > **Advanced Configuration**.

Step 2 Click **FlexConfig** > **FlexConfig Objects** in the Advanced Configuration table of contents.

Step 3 Click the **+** button to create a new object.

Step 4 Enter a name for the object. For example, **Enable_Power_Supply_Alarm**.

Step 5 In the **Template** editor, enter the commands needed to configure the power supply alarm.

a) Enable dual power supply mode.

power-supply dual

For example:

```
power-supply dual
```

b) Configure the actions to take when the power supply alarm is triggered.

alarm facility power-supply rps {**relay** | **syslog** | **notifies** | **disable**}

You can configure more than one action. For example, you can configure the device to activate the external alarm, send syslog messages, and also send SNMP traps.

- **relay**—Energize the alarm output relay, which activates the external alarm that you attached to it, such as a buzzer or a flashing light. The output LED also goes red.
- **syslog**—Send a syslog message. This option is enabled by default.
- **notifies**—Send an SNMP trap.
- **disable**—Disable the power supply alarm. Any other actions configured for the power supply alarm are inoperable.

For example, to enable all actions for the power supply alarm, enter the following:

```
alarm facility power-supply rps relay
alarm facility power-supply rps syslog
alarm facility power-supply rps notifies
```

Step 6 In the **Negate Template** editor, enter the lines required to undo this configuration.

All of these commands take the **no** form to disable them and return to default settings. For example, if your template includes all of the command examples shown in this procedure, the negate template would be the following:

```
no power-supply dual
no alarm facility power-supply rps relay
no alarm facility power-supply rps syslog
no alarm facility power-supply rps notifies
```

Step 7 Click **OK** to save the object.

Step 8 Add the object to the FlexConfig policy.

a) Click **FlexConfig Policy** in the table of contents.
b) Click **+** in the Group List.
c) Select the Enable_Power_Supply_Alarm object and click **OK**.

The preview should update with the commands in the template. Verify you are seeing the expected commands.

d) Click **Save**.

You can now deploy the policy.

Step 9 After deployment completes, in CLI Console or an SSH session, use the **show running-config** command and verify that the running configuration has the correct changes.

Configure Temperature Alarms

You can configure alarms based on the temperature of the CPU card in the device.

You can set a primary and secondary temperature range. If the temperature drops below the low threshold, or exceeds the high threshold, the alarm is triggered.

The primary temperature alarm is enabled by default for all alarm actions: output relay, syslog, and SNMP. The default settings for the primary temperature range is -40°C to 92°C.

The secondary temperature alarm is disabled by default. You can set the secondary temperature within the range -35°C to 85°C.

Because the secondary temperature range is more restrictive than the primary range, if you set either the secondary low or high temperature, that setting disables the corresponding primary setting, even if you configure non-default values for the primary setting. You cannot enable two separate high and two separate low temperature alarms.

Thus, in practice, you should configure the primary only, or the secondary only, setting for high and low.

Procedure

Step 1 Click **View Configuration** in **Device** > **Advanced Configuration**.

Step 2 Click **FlexConfig** > **FlexConfig Objects** in the Advanced Configuration table of contents.

Step 3 Click the + button to create a new object.

Step 4 Enter a name for the object. For example, **Enable_Temperature_Alarm**.

Step 5 In the **Template** editor, enter the commands needed to configure the temperature alarm.

a) Configure the acceptable temperature range.

alarm facility temperature {**primary** | **secondary**} {**low** | **high**} *temperature*

The temperature is in Celsius. The allowed range for the primary alarm is -40 to 92, which is also the default range. The allowed range for the secondary alarm is -35 to 85. The low value must be lower than the high value.

For example, to set a more restrictive temperature range of -20 to 80, which falls within the allowed range for the secondary alarm, configure the secondary alarm as follows:

```
alarm facility temperature secondary low -20
alarm facility temperature secondary high 80
```

b) Configure the actions to take when the temperature alarm is triggered.

alarm facility temperature {**primary** | **secondary**} {**relay** | **syslog** | **notifies**}

You can configure more than one action. For example, you can configure the device to activate the external alarm, send syslog messages, and also send SNMP traps.

- **relay**—Energize the alarm output relay, which activates the external alarm that you attached to it, such as a buzzer or a flashing light. The output LED also goes red.
- **syslog**—Send a syslog message.
- **notifies**—Send an SNMP trap.

For example, to enable all actions for the secondary temperature alarm, enter the following:

```
alarm facility temperature secondary relay
alarm facility temperature secondary syslog
alarm facility temperature secondary notifies
```

Step 6 In the **Negate Template** editor, enter the lines required to undo this configuration.

All of these commands take the **no** form to return to default settings (for the primary alarm) or disable them (for the secondary alarm). For example, if your template includes all of the command examples shown in this procedure, the negate template would be the following:

```
no alarm facility temperature secondary low -20
no alarm facility temperature secondary high 80
no alarm facility temperature secondary relay
no alarm facility temperature secondary syslog
no alarm facility temperature secondary notifies
```

Step 7 Click **OK** to save the object.

Step 8 Add the object to the FlexConfig policy.

a) Click **FlexConfig Policy** in the table of contents.

b) Click **+** in the Group List.

c) Select the Enable_Temperature_Alarm object and click **OK**.

The preview should update with the commands in the template. Verify you are seeing the expected commands.

d) Click **Save**.

You can now deploy the policy.

Step 9 After deployment completes, in CLI Console or an SSH session, use the **show running-config** command and verify that the running configuration has the correct changes.

Monitoring alarms

The following topics explain how to monitor and manage alarms.

Monitoring alarm status

You can use the following commands in the CLI to monitor alarms.

- **show alarm settings**

  Shows the current configuration for each possible alarm.

- **show environment alarm-contact**

  Shows information about the physical status of the input alarm contacts.

- **show facility-alarm relay**

  Shows information about the alarms that have triggered the output relay.

- **show facility-alarm status** [**info** | **major** | **minor**]

  Shows information on all alarms that have been triggered. You can limit the view by filtering on **major** or **minor** status. The **info** keyword provides the same output as using no keyword.

Monitoring syslog messages for alarms

Depending on the type of alarms you configure, you might see the following syslog messages.

Dual Power Supply Alarms

- %FTD-1-735005: Power Supply Unit Redundancy OK
- %FTD-1-735006: Power Supply Unit Redundancy Lost

Temperature Alarms

In these alarms, *Celsius* is replaced by the temperature detected on the device, in Celsius.

- %FTD-6-806001: Primary alarm CPU temperature is High *Celsius*
- %FTD-6-806002: Primary alarm for CPU high temperature is cleared
- %FTD-6-806003: Primary alarm CPU temperature is Low *Celsius*
- %FTD-6-806004: Primary alarm for CPU Low temperature is cleared
- %FTD-6-806005: Secondary alarm CPU temperature is High *Celsius*
- %FTD-6-806006: Secondary alarm for CPU high temperature is cleared
- %FTD-6-806007: Secondary alarm CPU temperature is Low *Celsius*
- %FTD-6-806008: Secondary alarm for CPU Low temperature is cleared

Alarm Input Contact Alarms

In these alarms, *description* is the description for the contact that you configured.

- %FTD-6-806009: Alarm asserted for ALARM_IN_1 *alarm_1_description*
- %FTD-6-806010: Alarm cleared for ALARM_IN_1 *alarm_1_description*
- %FTD-6-806011: Alarm asserted for ALARM_IN_2 *alarm_2_description*
- %FTD-6-806012: Alarm cleared for ALARM_IN_2 *alarm_2_description*

Turn off the external alarm

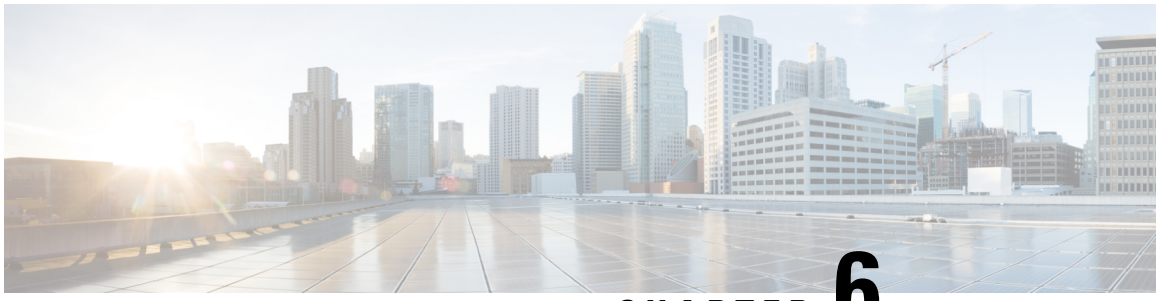
If you are using an external alarm that is attached to the alarm output, and the alarm is triggered, you can turn off the external alarm from the device CLI using the **clear facility-alarm output** command. This command de-energizes the output pin and also turns off the output LED.



PART II

Reusable Objects

- Objects, on page 131
- Certificates, on page 147
- Identity Sources, on page 157



CHAPTER 6

Objects

Objects are reusable containers that define criteria that you want to use in policies or other settings. For example, network objects define host and subnet addresses.

Objects let you define criteria so that you can easily reuse the same criteria in different policies. When you update an object, all policies that use the object are automatically updated.

- Object Types, on page 131
- Guidelines and Limitations for Objects, on page 134
- Managing Objects, on page 134

Object Types

You can create the following types of object. In most cases, if a policy or setting allows an object, you must use an object.

Object Type	Main Use	Description
AnyConnect Client Profile	Remote access VPN.	AnyConnect Client profiles are downloaded to clients along with the AnyConnect Client software. These profiles define many client-related options, such as auto connect on startup and auto reconnect, and whether the end user is allowed to change the option from the AnyConnect Client preferences and advanced settings. See Configure and Upload Client Profiles, on page 680.
Application Filter	Access control rules.	An application filter object defines the applications used in an IP connection, or a filter that defines applications by type, category, tag, risk, or business relevance. You can use these objects in policies to control traffic instead of using port specifications. See Configuring Application Filter Objects, on page 139.

Object Type	Main Use	Description
Certificates	Identity policies. Remote access VPN. SSL decryption rules. Management web server.	Digital certificates provide digital identification for authentication. Certificates are used for SSL (Secure Socket Layer), TLS (Transport Layer Security), and DTLS (Datagram TLS) connections, such as HTTPS and LDAPS. See Configuring Certificates, on page 150.
DNS Groups	DNS settings for the management and data interfaces.	DNS groups define a list of DNS servers and some associated attributes. DNS servers are needed to resolve fully-qualified domain names (FQDN), such as www.example.com, to IP addresses. See Configuring DNS Groups, on page 762.
Event List Filters	System logging settings for select logging destinations.	Event list filters create a custom filter list for syslog messages. You can use them to limit the messages that are sent to a particular logging location, such as a syslog server or the internal log buffer. See Configure Event List Filters, on page 754.
Geolocation	Security policies.	A geolocation object defines countries and continents that host the device that is the source or destination of traffic. You can use these objects in policies to control traffic instead of using IP addresses. See Configuring Geolocation Objects, on page 142.
Identity Sources	Identity policies. Remote access VPN. Firewall Device Manager access.	Identity sources are servers and databases that define user accounts. You can use this information in a variety of ways, such as providing the user identity associated with an IP address, or authenticating remote access VPN connections or access to the Firewall Device Manager. See Identity Sources, on page 157.
IKE Policy	VPN.	Internet Key Exchange (IKE) Policy objects define the IKE proposal used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and automatically establish IPsec security associations (SAs). There are separate objects for IKEv1 and IKEv2. See Configuring the Global IKE Policy, on page 642.
IPsec Proposal	VPN.	IPsec Proposal objects configure the IPsec proposal used during IKE Phase 2 negotiations. The IPsec proposal defines the combination of security protocols and algorithms that secure traffic in an IPsec tunnel. There are separate objects for IKEv1 and IKEv2. See Configuring IPsec Proposals, on page 647.

Object Type	Main Use	Description
Network	Security policies and a wide variety of device settings.	Network groups and network objects (collectively referred to as network objects) define the addresses of hosts or networks. See Configuring Network Objects and Groups, on page 135.
Port	Security policies.	Port groups and port objects (collectively referred to as port objects) define the protocols, ports, or ICMP services for traffic. See Configuring Port Objects and Groups, on page 136.
Secret Keys	Smart CLI and FlexConfig policies.	Secret key objects define passwords or other authentication strings that you want to encrypt and hide. See Configuring Secret Key Objects, on page 852.
Security Zone	Security policies.	A security zone is a grouping of interfaces. Zones divide the network into segments to help you manage and classify traffic. See Configuring Security Zones, on page 138.
SGT Groups	Access control policies.	Trustsec Security Group Tags (SGT) define tags for traffic as defined in Cisco Identity Services Engine (ISE). You must configure ISE before you can create these objects. You can then use the objects as source/destination matching criteria in access control rules. See Configuring Security Group Tag (SGT) Groups, on page 145.
SLA Monitors	Static routes.	An SLA Monitor defines a target IP address to use for monitoring a static route. If the monitor determines the target IP address can no longer be reached, the system can install a backup static route. See Configure SLA Monitor Objects, on page 313.
SSL Ciphers	SSL settings.	An SSL Cipher object defines a combination of security level, TLS/DTLS protocol versions, and encryption algorithms that can be used when establishing an SSL connection to Firewall Threat Defense. Use these objects in the system settings to define the security requirements for users who make TLS/SSL connections to the box. See Configuring TLS/SSL Cipher Settings, on page 784.

Object Type	Main Use	Description
Syslog Servers	Access control rules. Diagnostic logging. Security Intelligence policies. SSL decryption rules. Intrusion policies. File/malware policies	A syslog server object identifies a server that can receive connection-oriented or diagnostic system log (syslog) messages. See Configuring Syslog Servers, on page 143.
URL	Access control rules. Security Intelligence policies.	URL objects and groups (collectively referred to as URL objects) define the URL or IP addresses of web requests. See Configuring URL Objects and Groups, on page 141.
Users	Remote access VPN.	You can create user accounts directly on the device for use with remote access VPN. You can use the local user accounts instead of, or in addition to, an external authentication source. See Configure Local Users, on page 175.

Guidelines and Limitations for Objects

- Objects share the same name space. You cannot use the same name for two objects, even if they are different types of objects.
- The system considers the following names reserved. You cannot use these names as object names: object, network, host, fqdn, no, range, subnet, description, nat, after-auto, source, destination, static, dynamic, any, dns, inactive, interface, service, pat-pool, round-robin, extended, flat, include-reserve, tcp, udp, no-proxy-arp, route-lookup, remark, rule-id, access-list, line, ip, icmp, icmp6, advanced, ifc, object-group, vlan, event-log, flow-start, flow-end, both, permit, trust, deny, group-object.

Managing Objects

You can configure objects directly through the Objects page, or you can configure them while editing policies. Either method yields the same results, a new or updated object, so use the technique that suits your needs at the time.

The following procedure explains how you can create and manage your objects directly through the Objects page.

Note When you edit a policy or setting, if a property requires an object, you are shown a list of the ones that are already defined, and you select the appropriate object. If the desired object does not yet exist, simply click the **Create New Object** link shown in the list.

Procedure

Step 1 Select **Objects**.

The Objects page has a table of contents listing the available types of objects. When you select an object type, you see a list of existing objects, and you can create new ones from here. You can also see the object contents and type.

Step 2 Select the object type from the table of contents and do any of the following:

- To create an object, click the + button. The content of the objects differ based on type; see the configuration topic for each object type for specific information.
- To create a group object, click the **Add Group** ( ) button. Group objects include more than one item.
- To edit an object, click the edit icon ( ) for the object. You cannot edit the contents of a pre-defined object.
- To delete an object, click the delete icon ( ) for the object. You cannot delete an object if it is currently being used in a policy or another object, or if it is a pre-defined object.

Configuring Network Objects and Groups

Use network group and network objects (collectively referred to as network objects) to define the addresses of hosts or networks. You can then use the objects in security policies for purposes of defining traffic matching criteria, or in settings to define the addresses of servers or other resources.

A network object defines a single host or network address, whereas a network group object can define more than one address.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create network objects while editing an address property by clicking the **Create New Network** link shown in the object list.

Procedure

Step 1 Select **Objects**, then select **Network** from the table of contents.

Step 2 Do one of the following:

- To create an object, click the + button.
- To create a group, click the **Add Group** ( ) button.

- To edit an object or group, click the edit icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 3 Enter a Name for the object and optionally, a description, and define the object contents.

We recommend that you do not use an IP address alone for the name so that you can easily tell object names from object contents or standalone IP addresses. If you want to use an IP address in the name, prefix it with something meaningful, such as host-192.168.1.2 or network-192.168.1.0. If you use an IP address as the name, the system adds a vertical bar as a prefix, for example, |192.168.1.2. Firewall Device Manager does not show the bar in the object selectors, but you will see this naming standard if you examine the running configuration using the **show running-config** command in the CLI.

Step 4 Configure the contents of the object.

Network Objects

Select the object **Type** and configure the contents:

- **Network**—Enter a network address using one of the following formats:
  - IPv4 network including subnet mask, for example, 10.100.10.0/24 or 10.100.10.0/255.255.255.0.
  - IPv6 network including prefix, for example, 2001:DB8:0:CD30::/60.
- **Host**—Enter a host IP address using one of the following formats:
  - IPv4 host address, for example, 10.100.10.10.
  - IPv6 host address, for example, 2001:DB8::0DB8:800:200C:417A or 2001:DB8:0:0:0DB8:800:200C:417A.
- **Range**—A range of addresses, with the starting and ending address separated by a hyphen. You can specify IPv4 or IPv6 ranges. Do not include masks or prefixes. For example, 192.168.1.10-192.168.1.250 or 2001:DB8:0:CD30::10-2001:DB8:0:CD30::100.
- **FQDN**—Enter a single fully-qualified domain name, such as www.example.com. You cannot use wildcards. Also, select the **DNS Resolution** to determine whether you want the IPv4, IPv6, or both IPv4 and IPv6 addresses associated with the FQDN. The default is both IPv4 and IPv6. You can use these objects in access control rules only. The rules match the IP address obtained for the FQDN through a DNS lookup.

Network Groups

Click the + button to select network objects or groups to add to the group. You can also create new objects.

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

Configuring Port Objects and Groups

Use port group and port objects (collectively referred to as port objects) to define the protocols, ports, or ICMP services for traffic. You can then use the objects in security policies for purposes of defining traffic matching criteria, for example, to use access rules to allow traffic to specific TCP ports.

A port object defines a single protocol, TCP/UDP port or port range, or ICMP service, whereas a port group object can define more than one service.

The system includes several pre-defined objects for common services. You can use these objects in your policies. However, you cannot edit or delete system-defined objects.

Note When creating port group objects, ensure that the combination of objects makes sense. For example, you cannot have a mixture of protocols in an object if you use it to specify both source and destination ports in an access rule. Exercise care when editing an object that is already being used, or you could invalid (and disable) policies that use the object.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create port objects while editing a service property by clicking the **Create New Port** link shown in the object list.

Procedure

Step 1 Select **Objects**, then select **Ports** from the table of contents.

Step 2 Do one of the following:

- To create an object, click the + button.
- To create a group, click the **Add Group** ( ) button.
- To edit an object or group, click the edit icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 3 Enter a name for the object and optionally, a description, and define the object contents.

Port Objects

Select the **Protocol**, then configure the protocol as follows:

- **TCP**, **UDP**—Enter the single port or port range number, for example, 80 (for HTTP) or 1-65535 (to cover all ports).
- **ICMP**, **IPv6-ICMP**—Select the ICMP **Type** and optionally, the **Code**. Select **Any** for the type to apply to all ICMP messages. For information on the types and codes, see the following pages:
  - ICMP—http://www.iana.org/assignments/icmp-parameters/icmp-parameters.xml
  - ICMPv6—http://www.iana.org/assignments/icmpv6-parameters/icmpv6-parameters.xml
- **Other**—Select the desired protocol.

Port Groups

Click the + button to select port objects to add to the group. You can also create new objects.

Step 4 Click **OK** to save your changes.

Configuring Security Zones

A security zone is a grouping of interfaces. Zones divide the network into segments to help you manage and classify traffic. You can define multiple zones, but a given interface can be in one zone only.

The system creates the following zones during initial configuration. You can edit these zones to add or remove interfaces, or you can delete the zones if you no longer use them.

- **inside_zone**—Includes the inside interface. If the inside interface is a bridge group, this zone includes all the bridge group member interfaces instead of the inside Bridge Virtual Interface (BVI). This zone is intended to represent internal networks.
- **outside_zone**—Includes the outside interface. This zone is intended to represent networks external to your control, such as the Internet.

Typically, you would group interfaces by the role they play in your network. For example, you would place the interface that connects to the Internet in the **outside_zone** security zone, and all of the interfaces for your internal networks in the **inside_zone** security zone. Then, you could apply access control rules to traffic coming from the outside zone and going to the inside zone.

Before creating zones, consider the access rules and other policies you want to apply to your networks. For example, you do not need to put all internal interfaces into the same zone. If you have 4 internal networks, and you want to treat one differently than the other three, you can create two zones rather than one. If you have an interface that should allow outside access to a public web server, you might want to use a separate zone for the interface.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create security zones while editing a security zone property by clicking the **Create New Security Zone** link shown in the object list.

Procedure

Step 1 Select **Objects**, then select **Security Zones** from the table of contents.

Step 2 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 3 Enter a Name for the object and optionally, a description.

Step 4 Select the **Mode** for the zone.

The mode relates directly to the interface mode. The zone can contain a single type of interface.

- **Routed**—Routed interfaces are the normal interfaces for through traffic that can apply security policies.
- **Passive**—Passive interfaces do not affect traffic that flows through the device.

Step 5 In the **Interfaces** list, click + and select the interfaces to add to the zone.

The list shows all named interfaces that are not currently in a zone. You must configure an interface and give it a name before you can add it to a zone.

If all named interfaces are already in zones, the list is empty. If you are trying to move an interface to a different zone, you must first remove it from its current zone.

Note
You cannot add a bridge group interface (BVI) to a zone. Instead, add the member interfaces. You can put the members into different zones.

Step 6 Click **OK** to save your changes.

Configuring Application Filter Objects

An application filter object defines the applications used in an IP connection, or a filter that defines applications by type, category, tag, risk, or business relevance. You can use these objects in policies to control traffic instead of using port specifications.

Although you can specify individual applications, application filters simplify policy creation and administration. For example, you could 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 session is blocked.

You can select applications and application filters directly in a policy without using application filter objects. However, an object is convenient if you want to create several policies for the same group of applications or filters. The system includes several pre-defined application filters, which you cannot edit or delete.



Note Cisco frequently updates and adds additional application detectors via system and vulnerability database (VDB) updates. Thus, a rule blocking high risk applications can automatically apply to new applications without you having to update the rule manually.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create application filter objects while editing an access control rule by clicking the **Save As Filter** link after adding application criteria to the Applications tab.

Before you begin

When editing a filter, if a selected application was removed by a VDB update, "(Deprecated)" appears after the application name. You must remove these applications from the filter, or subsequent deployments and system software upgrades will be blocked.

Procedure

Step 1 Select **Objects**, then select **Application Filters** from the table of contents.

Step 2 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon () for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 3 Enter a Name for the object and optionally, a description.

Step 4 In the **Applications** list, click **Add +** and select the applications and filters to add to the object.

The initial list shows applications in a continually scrolling list. Click **Advanced Filter** to see the filter options and to get an easier view for selecting applications. Click **Add** when you have made your selections. You can repeat the process to add additional applications or filters.

Note
Multiple selections within a single filter criteria have an OR relationship. For example, Risk is High OR Very High. The relationship between filters is AND, so Risk is High OR Very High, AND Business Relevance is Low OR Very Low. As you select filters, the list of applications in the display updates to show only those that meet the criteria. You can use these filters to help you find applications that you want to add individually, or to verify that you are selecting the desired filters to add to the rule.

Risks

The likelihood that the application is used for purposes that might be against your organization's security policy, from very low to very high.

Business Relevance

The likelihood that the application is used within the context of your organization's business operations, as opposed to recreationally, from very low to very high.

Types

The type of application:

- **Application Protocol**—Application protocols such as HTTP and SSH, which represent communications between hosts.
- **Client Protocol**—Clients such as web browsers and email clients, which represent software running on the host.
- **Web Application**—Web applications such as MPEG video and Facebook, which represent the content or requested URL for HTTP traffic.

Categories

A general classification for the application that describes its most essential function.

Tags

Additional information about the application, similar to category.

For encrypted traffic, the system can identify and filter traffic using only the applications tagged **SSL Protocol**. Applications without this tag can only be detected in unencrypted or decrypted traffic. Also, the system assigns the **decrypted traffic** tag to applications that the system can detect in decrypted traffic only, not encrypted or unencrypted.

Applications List (bottom of the display)

This list updates as you select filters from the options above the list, so you can see the applications that currently match the filter. Use this list to verify that your filter is targeting the desired applications when you intend to add filter criteria to the rule. If your intention is to add specific applications, select them from this list.

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

Configuring URL Objects and Groups

Use URL objects and groups (collectively referred to as URL objects) to define the URL or IP addresses of web requests. You can use these objects to implement manual URL filtering in access control policies, or blocking in Security Intelligence policies.

A URL object defines a single URL or IP address, whereas a URL group object can define more than one URL or address.

When creating URL objects, keep the following points in mind:

- If you do not include a path (that is, there are no / characters in the URL), the match is based on the server's hostname only. If you include one or more / character, the entire URL string is used for a substring match. Then, a URL is considered a match if any of the following are true:
  - The string is at the beginning of the URL.
  - The string follows a dot.
  - The string contains a dot in the beginning.
  - The string follows the :// characters.

  For example, ign.com matches ign.com or www.ign.com, but not versign.com.



Note We recommend that you do not use manual URL filtering to block or allow individual web pages or parts of sites (that is, URL strings with / characters), as servers can be reorganized and pages moved to new paths.

- 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.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create URL objects while editing a URL property by clicking the **Create New URL** link shown in the object list.

Procedure

Step 1 Select **Objects**, then select **URL** from the table of contents.

Step 2 Do one of the following:

- To create an object, click the + button.
- To create a group, click the **Add Group** ( ) button.
- To edit an object or group, click the edit icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 3 Enter a Name for the object and optionally, a description.

Step 4 Define the object contents.

URL Objects

Enter a URL or IP address in the **URL** box. You cannot use wildcards in the URL.

URL Groups

Click the + button to select URL objects to add to the group. You can also create new objects.

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

Configuring Geolocation Objects

A geolocation object defines countries and continents that host the device that is the source or destination of traffic. You can use these objects in policies to control traffic instead of using IP addresses. For example, using geographical location, you could easily restrict access to a particular country without needing to know all of the potential IP addresses used there.

You can typically select geographical locations directly in a policy without using geolocation objects. However, an object is convenient if you want to create several policies for the same group of countries and continents.

Note To ensure that you are using up-to-date geographical location data to filter your traffic, Cisco strongly recommends that you regularly update the geolocation database (GeoDB).

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create geolocation objects while editing a network property by clicking the **Create New Geolocation** link shown in the object list.

Procedure

Step 1 Select **Objects**, then select **Geolocation** from the table of contents.

Step 2 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 3 Enter a Name for the object and optionally, a description.

Step 4 In the **Continents/Countries** list, click **Add +** and select the continents and countries to add to the object.

Selecting a continent selects all countries within the continent.

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

Configuring Syslog Servers

A syslog server object identifies a server that can receive connection-oriented or diagnostic system log (syslog) messages. If you have a syslog server set up for log collection and analysis, create objects to define them and use the objects in the related policies.

You can send the following types of events to the syslog server:

- Connection events. Configure the syslog server object on the following types of policy: access control rules and default action, SSL decryption rules and default action, Security Intelligence policy.
- Intrusion events. Configure the syslog server object on the intrusion policy.
- Diagnostic events. See Configure Logging to a Remote Syslog Server, on page 752.
- File/malware events. Configure the syslog server on **Device** > **System Settings** > **Logging Settings**.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create syslog server objects while editing a syslog server property by clicking the **Add Syslog Server** link shown in the object list.

Procedure

Step 1 Select **Objects**, then select **Syslog Servers** from the table of contents.

Step 2 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 3 Configure the syslog server properties:

- **IP Address**—Enter the IP address of the syslog server.
- **Protocol Type**, **Port Number**—Select the protocol and enter the port number to use for syslog. The default is UDP/514. If you select **TCP**, the system can recognize when the syslog server is not available, and stops sending events until the server is available again. The default UDP port is 514, the default TCP port is 1470. If you change the default, the port must be in the range 1025 to 65535.

  Note
  If you use TCP as the transport protocol, the system opens 4 connections to the syslog server to ensure that messages are not lost. If you are using the syslog server to collect messages from a very large number of devices, and the combined connection overhead is too much for the server, use UDP instead.

- **Interface for Device Logs**—Select which interface should be used for sending diagnostic syslog messages. The following types of event always use the management interface: connection, intrusion, file, malware. Your interface selection determines the IP address associated with syslog messages. Select one of the following options:
  - **Data Interface**—Use the data interface you select for diagnostic syslog messages. If the server is accessible through a bridge group member interface, select the bridge group interface (BVI) instead. If it is accessible through the Diagnostic interface (the physical management interface), we recommend that you select **Management Interface** instead of this option. You cannot select a passive interface.

    For connection, intrusion, file, and malware syslog messages, the source IP address will either be for the management interface, or for the gateway interface if you route through data interfaces. Note that there must be appropriate routes in the routing table that direct traffic to the syslog server out the selected interface for these event types.
  - **Management Interface**—Use the virtual Management interface for all types of syslog messages. The source IP address will either be for the Management interface, or for the gateway interface if you route through data interfaces.

Step 4 Click **OK** to save your changes.

Configuring Security Group Tag (SGT) Groups

Use Security Group Tag (SGT) group objects to identify source or destination addresses based on an SGT assigned by Identity Services Engine (ISE). You can then use the objects in access control rules for purposes of defining traffic matching criteria.

You cannot use the information retrieved from ISE directly in an access control rule. Instead, you need to create SGT groups, which refer to the downloaded SGT information. Your SGT groups can refer to more than one SGT, so you can apply policy based on relevant collections of tags if that is appropriate.

For detailed information on how to use SGTs for access control, see How to Control Network Access Using TrustSec Security Group Tags, on page 503.

Before you begin

Before you create SGT groups, you must configure the ISE identity source to subscribe to SXP mappings and deploy changes. The system then retrieves SGT information from the ISE server. Only after SGTs are downloaded will be you able to create SGT groups.

Procedure

Step 1 Select **Objects**, then select **SGT Groups** from the table of contents.

Step 2 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 3 Enter a **Name** for the object and optionally, a description,

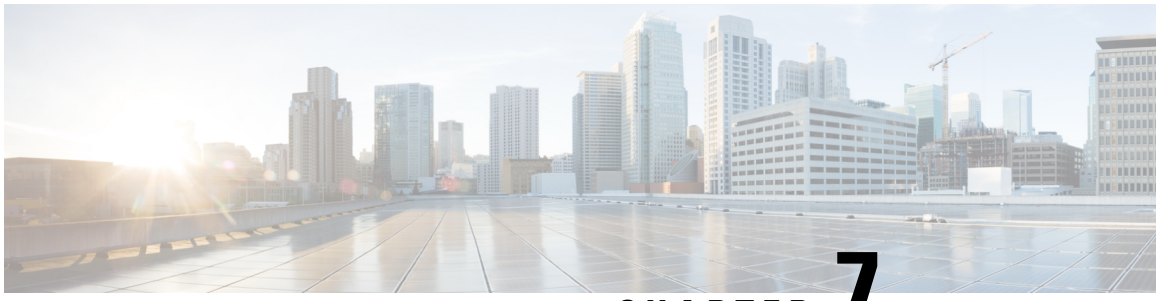
Step 4 Under **Tags**, click + and select the downloaded SGTs to include in the object.

To remove an SGT, click **x** on the right side of the tag name.

If the list is empty, then the system was not able to download any SGT mappings. If this happens:

- Ensure that the ISE identity object is subscribing to the SXP Topic. You must subscribe to SXP to get the mappings.
- Verify that static mappings are defined in ISE, and that ISE is configured to publish these mappings. If no mappings exist, there is simply nothing to download. See Configure security groups and SXP publishing in ISE, on page 504.

Step 5 Click **OK**.



CHAPTER 7

Certificates

Digital certificates provide digital identification for authentication. Certificates are used for SSL (Secure Socket Layer), TLS (Transport Layer Security), and DTLS (Datagram TLS) connections, such as HTTPS and LDAPS. The following topics explain how to create and manage certificates.

- About Certificates, on page 147
- Configuring Certificates, on page 150

About Certificates

Digital certificates provide digital identification for authentication. A digital certificate includes information that identifies a device or user, such as the name, serial number, company, department, or IP address. A digital certificate also includes a copy of the public key for the user or device. Certificates are used for SSL (Secure Socket Layer), TLS (Transport Layer Security), and DTLS (Datagram TLS) connections, such as HTTPS and LDAPS.

You can create the following types of certificate:

- Internal certificates—Internal identity certificates are certificates for specific systems or hosts. You can generate these yourself using the OpenSSL toolkit or get them from a Certificate Authority. You can also generate a self-signed certificate. When internal certificates expire or becomes invalid for any reason, you can regenerate it from the following CLISH CLI command:

```
> system support regenerate-security-keyring
  String  Certificate to be regenerated, default or fdm
```

- Internal Certificate Authority (CA) certificates—Internal CA certificates are certificates that the system can use to sign other certificates. These certificates differ from internal identity certificates with respect to the basic constraints extension and the CA flag, which are enabled for CA certificates but disabled for identity certificates. You can generate these yourself using the OpenSSL toolkit or get them from a Certificate Authority. You can also generate a self-signed internal CA certificate. If you configure self-signed internal CA certificates, the CA runs on the device itself.
- Trusted Certificate Authority (CA) certificates—A trusted CA certificate is used to sign other certificates. It is self-signed and called a root certificate. A certificate that is issued by another CA certificate is called a subordinate certificate.

Certificate Authorities (CAs) are trusted authorities that “sign” certificates to verify their authenticity, thereby guaranteeing the identity of the device or user. CAs issue digital certificates in the context of a PKI, which uses public-key or private-key encryption to ensure security. A CA can be a trusted third party, such as

VeriSign, or a private (in-house) CA that you establish within your organization. CAs are responsible for managing certificate requests and issuing digital certificates. For more information, see Public Key Cryptography, on page 148.

Public Key Cryptography

In public key cryptography, such as the RSA encryption system, each user has a key pair containing both a public and a private key. The keys act as complements, and anything encrypted with one of the keys can be decrypted with the other.

In simple terms, a signature is formed when data is encrypted with a private key. The signature is attached to the data and sent to the receiver. The receiver applies the public key of the sender to the data. If the signature sent with the data matches the result of applying the public key to the data, the validity of the message is established.

This process relies on the receiver having a copy of the public key of the sender and a high degree of certainty that this key belongs to the sender, not to someone pretending to be the sender.

Obtaining the public key of a sender is normally handled externally or through an operation performed at installation. For example, most web browsers are configured with the root certificates of several CAs by default.

You can learn more about digital certificates and public key cryptography through openssl.org, Wikipedia, or other sources. Having a firm understanding of SSL/TLS cryptography will help you establish secure connections to your device.

Certificate Types Used by Feature

You need to create the right type of certificate for each feature. The following features require certificates.

Identity Policies (Captive Portal)—Internal Certificate

(Optional.) Captive portal is used in identity policies. Users must accept this certificate when authenticating to the device for purposes of identifying themselves and getting their IP address associated with their usernames. If you do not supply a certificate, the device uses an automatically generated certificate.

Identity Realms (Identity Policies and Remote Access VPN)—Trusted CA Certificate

(Optional.) If you use an encrypted connection for your directory server, the certificate must be accepted to perform authentication with the directory server. Users must authenticate when prompted by identity and remote access VPN policies. A certificate is not needed if you do not use encryption for the directory server.

Management Web Server (Management Access System Settings)—Internal Certificate

(Optional.) Firewall Device Manager is a web-based application, so it runs on a web server. You can upload a certificate that your browser accepts as valid to avoid getting an Untrusted Authority warning.

Remote Access VPN—Internal Certificate

(Required.) The internal certificate is for the outside interface, which establishes the device identity for AnyConnect Clients when they make a connection to the device. Clients must accept this certificate.

Site-to-Site VPN—Internal and Trusted CA Certificates

If you use certificate authentication for a site-to-site VPN connection, you need to select the internal identity certificate used to authenticate the local peer in the connection. Although it is not part of the

VPN connection definition, you also need to upload the trusted CA certificates that were used to sign the local and remote peer identity certificates, so that the system can authenticate the peers.

SSL Decryption Policy—Internal, Internal Certificate, and Trusted CA Certificates and Certificate Groups

(Required.) The SSL decryption policy uses certificates for the following purposes:

- Internal certificates are used for known key decryption rules.
- Internal CA certificates are used for decrypt re-sign rules when creating the session between the client and Firewall Threat Defense device.
- Trusted CA certificates are used indirectly for decrypt re-sign rules when creating the session between the Firewall Threat Defense device and server. Trusted CA certificates are used to verify the signing authority of the server's certificate. You can configure these certificates directly or in a certificate group in the policy settings. The system includes a large number of trusted CA certificates, which are collected in the Cisco-Trusted-Authorities group, so you might not need to upload any additional certificates.

Example: Generating an Internal Certificate using OpenSSL

The following example uses OpenSSL commands to generate an internal server certificate. You can obtain OpenSSL from openssl.org. Consult OpenSSL documentation for specific information. The commands used in this example might change, and you might have other options available that you might want to use.

This procedure is meant to give you an idea of how to obtain a certificate to upload to Firewall Device Manager.

Note The OpenSSL commands shown here are examples only. Adjust the parameters to fit your security requirements.

Procedure

Step 1 Generate a key.

```
openssl genrsa -out server.key 4096
```

Step 2 Generate a certificate signing request (CSR).

```
openssl req -new -key server.key -out server.csr
```

Step 3 Generate a self-signed certificate with the key and CSR.

```
openssl x509 -req -days 365 -in server.csr -signkey server.key -out server.crt
```

Because the Firewall Device Manager does not support encrypted keys, try to skip the challenge password by just pressing return when generating a self signed certificate.

Step 4 Upload the files into the appropriate fields when creating an internal certificate object in the Firewall Device Manager.

You can also copy/paste the file contents. The sample commands create the following files:

- server.crt—Upload or paste the contents into the Server Certificate field.
- server.key—Upload or paste the contents into the Certificate Key field. If you provided a password when generating the key, you can decrypt it using the following command. The output is sent to stdout, where you can copy it.

```
openssl rsa -in server.key -check
```

Configuring Certificates

Firewall Threat Defense supports X509 certificates in PEM or DER format. Use OpenSSL to generate certificates if needed, obtain them from a trusted Certificate Authority, or create self-signed certificates.

For more information on certificates, see About Certificates, on page 147.

For information on which type is used for each feature, see Certificate Types Used by Feature, on page 148.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create certificate objects while editing a certificate property by clicking the **Create New Certificate** link shown in the object list.

Procedure

Step 1 Select **Objects**, then select **Certificates** from the table of contents.

The system comes with the following pre-defined certificates, which you can use as is or replace.

- DefaultInternalCertificate
- DefaultWebserverCertificate
- NGFW-Default-InternalCA

The system also includes many trusted CA certificates from third party Certificate Authorities. These are used by SSL decryption policies for Decrypt Re-Sign actions. The Cisco-Trusted-Authorities group includes all of these certificates, and is the default group used by the SSL decryption policy.

You can click the pre-defined search filters to limit the list to just **System-defined** or **User-defined** certificates.

You can also use the **Weak Key** filter to find certificates whose keys are different from supported key lengths. The keys are considered weak even if they are longer than the supported lengths. We recommend that you replace these certificates with ones that have supported key lengths. The following are the supported lengths:

- RSA keys can be 2048, 3072, or 4096 bits.
- ECDSA keys can be 256, 384, or 521 bits.

- EDDSA keys can be 256 bits.

Step 2 Do one of the following:

- To create a new certificate object, use the command for the type of certificate from the + menu.
- To create a new certificate group, click  and select **Add Certificate Group**.
- To view or edit a certificate or group, click either the edit icon () or the view icon () for the certificate.
- To delete an unreferenced certificate or group, click the trash can icon () for the certificate.

For detailed information on creating or editing certificates, see the following topics:

- Uploading Internal and Internal CA Certificates, on page 151
- Generating Self-Signed Internal and Internal CA Certificates, on page 153
- Uploading Trusted CA Certificates, on page 154
- Configuring Trusted CA Certificate Groups, on page 156

Uploading Internal and Internal CA Certificates

Internal identity certificates are certificates for specific systems or hosts.

Internal CA certificates are certificates that the system can use to sign other certificates. These certificates differ from internal identity certificates with respect to the basic constraints extension and the CA flag, which are enabled for CA certificates but disabled for identity certificates.

You can generate these certificates yourself using the OpenSSL toolkit or get them from a Certificate Authority, and then upload them using the following procedure. For an example of generating a key, see Example: Generating an Internal Certificate using OpenSSL, on page 149.

You can also generate a self-signed internal identity and internal CA certificates. If you configure self-signed internal CA certificates, the CA runs on the device itself. For information on creating self-signed certificates, see Generating Self-Signed Internal and Internal CA Certificates, on page 153.

For information on the features that use these certificates, see Certificate Types Used by Feature, on page 148.

Procedure

Step 1 Select **Objects**, then select **Certificates** from the table of contents.

Step 2 Do one of the following:

- Click + > **Add Internal Certificate**, then click **Upload Certificate and Key**.
- Click + > **Add Internal CA Certificate**, then click **Upload Certificate and Key**.

- To edit or view a certificate, click the information icon ( ). The dialog box shows the certificate subject, issuer, and valid time range. Click Replace Certificate to upload a new certificate and key. You can also paste the certificate and key in the dialog box.

Step 3 Enter a **Name** for the certificate.

The name is used in the configuration as an object name only, it does not become part of the certificate itself.

Step 4 Click **Upload Certificate** (or **Replace Certificate** when editing) and select the certificate file (for example, *.crt). Allowed file extensions are .pem, .cert, .cer, .crt, and .der. Alternatively, paste in the certificate.

The certificate must be an X509 certificate in PEM or DER format.

The certificate you paste must include the BEGIN CERTIFICATE and END CERTIFICATE lines. For example:

```
-----BEGIN CERTIFICATE-----
MIICMTCCAZoCCQDdUV3NGK/cUjANBgkqhkiG9w0BAQsFADBdMQswCQYDVQQGEwJV
UzETMBEGA1UECAwKU29tZS1TdGF0ZTEhMB8GA1UECgwYSW50ZXJuZXQgV2lkZ2l0
(...5 lines removed...)
shGJDReRYJQqilhHZrYTWZAYTrD7NQPHutK+ZiJng67cPgnNDuXEn55UwMOQoHBp
HMUwmhiGZlzJM8BpX2Js2yQ3ms30pr8rO+gPCPMCAwEAATANBgkqhkiG9w0BAQsF
AAOBgQCB02CebA6YjJCGr2CJZrQSeUwSveRBpmOuoqm98o2Z+5gJM5CkqgfxwCUn
RV7LRfQGFYd76V/5uor4Wx2ZCjqy6+zuQEm4ZxWNSZpA9UBixFXJCs9MBO4qkG5D
vlk3WYJfcgyJ10h4E4b0W2xiixBU+xoOTLRATnbKY36EWAG5cw==
-----END CERTIFICATE-----
```

Step 5 Click **Upload Key** (or **Replace Key** when editing) and select the certificate file (for example, *.key). The file extension must be .key. Alternatively, paste in the key for the certificate.

The key cannot be encrypted and it must be an RSA key.

For example:

```
-----BEGIN RSA PRIVATE KEY-----
MIICXQIBAAKBgQClSu1BknrMjzw/5FZ9YgdMLDUGJlbYgkjN7mVrkjyLQx2TYsem
r8iTiKB6iyTKbuS4iPeyEYkNF5FglCqKWEdmthNZkBhOsPs1A8e60r5mImeDrtw+
Cc0O5cSfnlTAw5CgcGkcxTCaGIZmXMkzwGlfYmzbJDeazfSmvys76A8I8wIDAQAB
AoGAUVDgEX8vXE0m9cOubPZ54pZo64KW/OJzUKP0TwxdLqGw/h39XFpkEXiIgmDL
(...5 lines removed...)
DSWvzekRDH83dmP66+MIbWePhbhty+D1OxbiuVuHV0/ZhxOhCG8tig3R8QJBAJmj
fId05+1dNI4tGbWv6hHh/H/dTP2STlZ3jERMZd29fjIRuJ9jpFC2lIDjvs8YGeAe
0YHkfSOULJn8/jOCf6kCQQDIJiHfGF/31Dk/8/5MGrg+3zau6oKXiuv6db8Rh+7l
MUOx09tvbBUy9REJq1YJWTKpeKD+E0QL+FX0bqvz4tHA
-----END RSA PRIVATE KEY-----
```

Step 6 Click **OK**.

If the key size is smaller than the minimum size we allow for generated self-signed certificates, you are warned that the certificate does not meet the recommended minimum requirements. Click **Proceed** to upload the certificate anyway, but we recommend that you create a new, stronger certificate.

Generating Self-Signed Internal and Internal CA Certificates

Internal identity certificates are certificates for specific systems or hosts.

Internal CA certificates are certificates that the system can use to sign other certificates. These certificates differ from internal identity certificates with respect to the basic constraints extension and the CA flag, which are enabled for CA certificates but disabled for identity certificates.

You can generate a self-signed internal identity and internal CA certificates, that is, the certificates are signed by the device itself. If you configure self-signed internal CA certificates, the CA runs on the device. The system generates both the certificate and the key.

You can also create these certificates using OpenSSL, or obtain them from a trusted CA, and upload them. For more information, see Uploading Internal and Internal CA Certificates, on page 151.

For information on the features that use these certificates, see Certificate Types Used by Feature, on page 148.

Procedure

Step 1 Select **Objects**, then select **Certificates** from the table of contents.

Step 2 Do one of the following:

- Click + > **Add Internal Certificate**, then click **Self-Signed Certificate**.
- Click + > **Add Internal CA Certificate**, then click **Self-Signed Certificate**.

Note

To edit or view a certificate, click the information icon (). The dialog box shows the certificate subject, issuer, and valid time range. Click **Replace Certificate** to upload a new certificate and key. When replacing a certificate, you cannot redo the self-signed characteristics explained in the following steps. Instead, you must paste or upload a new certificate as described in Uploading Internal and Internal CA Certificates, on page 151. The remaining steps apply to new self-signed certificates only.

Step 3 Enter a **Name** for the certificate.

The name is used in the configuration as an object name only, it does not become part of the certificate itself.

Step 4 Configure at least one of the following for the certificate subject and issuer information.

- **Country** (C)—The two-character ISO 3166 country code to include in the certificate. For example, the country code for the United States is US. Select the country code from the drop-down list.
- **State or Province** (ST)—The state or province to include in the certificate.
- **Locality or City** (L)—The locality to include in the certificate, such as the name of the city.
- **Organization** (O)—The organization or company name to include in the certificate.
- **Organizational Unit (Department)** (OU)—The name of the organization unit (for example, a department name) to include in the certificate.
- **Common Name** (CN)—The X.500 common name to include in the certificate. This could be the name of the device, web site, or another text string. This element is usually required for successful connections. For example, you must include a CN in the internal certificate used for remote access VPN.

- **Key Type**—The type of key to generate for this certificate: RSA, ECDSA (Elliptic Curve Digital Signature Algorithm), or EDDSA (Edwards-curve Digital Signature Algorithm).
- **Key Size**—The size of the key to generate. In general, longer keys are more secure. However, keys with larger modulus sizes take longer to generate and longer to process when exchanged. The allowed sizes differ based on the key type.
  - RSA keys can be 2048, 3072, or 4096 bits.
  - ECDSA keys can be 256, 384, or 521 bits.

    EDDSA keys can be 256 bits.
- **Validity Period**—How long the certificate will be considered valid. The default is 825 days from today, regardless of how you set the expiration date. Click **Set default** to return to the default. You can configure the period using one of the following methods. Be sure to replace certificates before they expire.
  - **By Date**—Click the **Expiration Date** and select the last day the certificate should be considered valid.
  - **By Number of Days**—Enter the number of days, starting today, that the certificate should be considered valid. After entering the number, you can click **By Date** to see the calculated expiration date.

Step 5 Click **Save**.

Uploading Trusted CA Certificates

A trusted Certificate Authority (CA) certificate is used to sign other certificates. It is self-signed and called a root certificate. A certificate that is issued by another CA certificate is called a subordinate certificate.

For information on the features that use these certificates, see Certificate Types Used by Feature, on page 148.

Obtain a trusted CA certificate from an external Certificate Authority, or create one using your own internal CA, for example, with OpenSSL tools. Then, use the following procedure to upload the certificate.

Procedure

Step 1 Select **Objects**, then select **Certificates** from the table of contents.

Step 2 Do one of the following:

- Click + > **Add Trusted CA Certificate**.
- To edit a certificate, click the edit icon () for the certificate.

Step 3 Enter a **Name** for the certificate.

The name is used in the configuration as an object name only, it does not become part of the certificate itself.

Step 4 Click **Upload Certificate** (or **Replace Certificate** when editing) and select the trusted CA certificate file (for example *.pem). Allowed file extensions are .pem, .cert, .cer, .crt, and .der. Alternatively, paste in the trusted CA certificate.

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 ad.example.com in the certificate, the connection fails.

The certificate must be an X509 certificate in PEM or DER format.

The certificate you paste must include the BEGIN CERTIFICATE and END CERTIFICATE lines. For example:

```
-----BEGIN CERTIFICATE-----
MIIFgTCCA2mgAwIBAgIJANvdcLnabFGYMA0GCSqGSIb3DQEBCwUAMFcxCzAJBgNV
BAYTAlVTMQswCQYDVQQIDAJUWDEPMA0GA1UEBwwGYXVzdGluMRQwEgYDVQQKDAsx
OTIuMTY4LjEuMTEUMBIGA1UEAwwLMTkyLjE2OC4xLjEwHhcNMTYxMDI3MjIzNDE3
WhcNMTcxMDI3MjIzNDE3WjBXMQswCQYDVQQGEwJVUzELMAkGA1UECAwCVFgxDzAN
BgNVBAcMBmF1c3RpbjEUMBIGA1UECgwLMTkyLjE2OC4xLjExFDASBgNVBAMMCzE5
Mi4xNjguMS4xMIICIjANBgkqhkiG9w0BAQEFAAOCAg8AMIICCgKCAgEA5NceYwtP
ES6Ve+S9z7WLKGX5JlF58AvH82GPkOQdrixn3FZeWLQapTpJZt/vgtAI2FZIK31h
(...20 lines removed...)
hbr6HOgKlOwXbRvOdksTzTEzVUqbgxt5Lwupg3b2ebQhWJz4BZvMsZX9etveEXDh
PY184V3yeSeYjbSCF5rP71fObG9Iu6+u4EfHp/NQv9s9dN5PMffXKieqpuN20Ojv
2b1sfOydf4GMUKLBUMkhQnip6+3W
-----END CERTIFICATE-----
```

Step 5 Set the **Validation Usage** to limit the use of the certificate.

Some features let you select whether a connection can be validated against a specific certificate. You must indicate in the certificate that these features can validly use the certificate, or a connection is refused.

Any feature not included in these options can validate against this certificate without an explicit usage allowance. For example, SSL decryption policies, and the web server that hosts the Firewall Device Manager, ignore the Validation Usage option. If you select any option in this field, the certificate will be downloaded to the running configuration that is displayed using the **show running-config** command.

The primary purpose of these options is to let you prevent VPN connections from getting established because they can be validated against a particular certificate.

- **SSL Server**—Validate the certificate on the remote SSL server. Used for dynamic DNS.
- **SSL Client**—Validate the certificate of the incoming remote access VPN connection.
- **IPsec Client**—Validate the certificate of the incoming IPsec site-to-site VPN connection.
- **Other**—Validate features, such as LDAPS, that are not managed by the Snort inspection engine. Select this option only if you are having problems with a particular feature. This option is mutually exclusive with all other options: you must deselect it before you can select any other options, and deselect all options before you can select this option.

Step 6 Click **OK**.

Configuring Trusted CA Certificate Groups

Use external trusted CA certificate groups in the SSL decryption policy settings to specify which certificates the SSL decryption policy should trust. If an end user attempts a connection to a site whose certificate's issuer's certificate is not among the trusted certificates, the user gets a message asking to trust the certificate. So, not having the certificate in the trusted list is an inconvenience to the end user, but does not itself prevent the connection (which you might accomplish with access control rules).

The default group is Cisco-Trusted-Authorities. You would need to create your own group only if:

- You want to trust certificates that are not in the default group. You would then select both the default group and your new group in the SSL decryption policy settings.
- You want to trust a more limited list of certificates than the default group. You would then create a group that has a complete list of trusted certificates, not just your delta, and select it as the sole group in the SSL decryption policy settings.

Before you begin

Upload all of the trusted CA certificates that you will add to the group if they are not already in the system.

Procedure

Step 1 Select **Objects**, then select **Certificates** from the table of contents.

Step 2 Do one of the following:

- To create a new certificate group, click  and select **Add Certificate Group**.
- To edit a certificate group, click the edit icon () for the group.

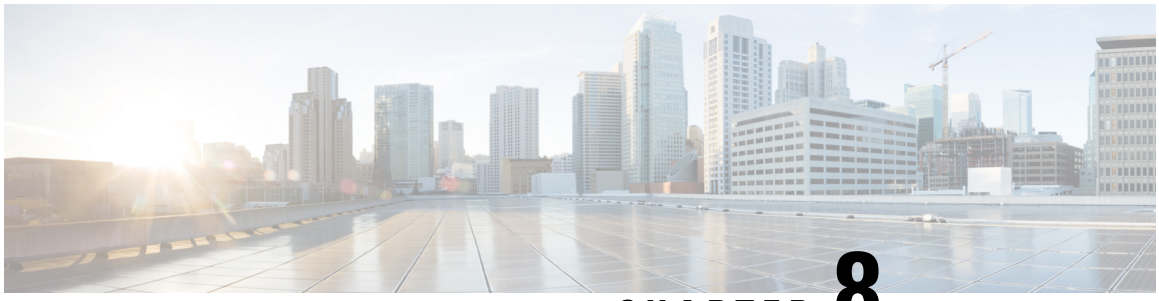
Step 3 Enter a **Name** for the certificate group and optionally, a description.

Step 4 Click **+** to add certificates to the group.

Add all the certificates you need in the group. You can click **Create New Trusted CA Certificate** to upload new ones while you build the group.

If you no longer need a certificate in the group, click the X icon (on the right) for the certificate.

Step 5 Click **OK**.



CHAPTER 8

Identity Sources

Identity sources are servers and databases that define user accounts. You can use this information in a variety of ways, such as providing the user identity associated with an IP address, or authenticating remote access VPN connections or access to the Firewall Device Manager.

The following topics explain how to define the identity sources. You would then use these objects when you configure the services that require an identity source.

- About Identity Sources, on page 157
- Active Directory (AD) Identity Realms, on page 159
- RADIUS Servers and Groups, on page 165
- Identity Services Engine (ISE), on page 169
- SAML Servers, on page 172
- Local Users, on page 175

About Identity Sources

Identity sources are the AAA servers and databases that define user accounts for the people in your organization. You can use this information in a variety of ways, such as providing the user identity associated with an IP address, or authenticating remote access VPN connections or access to the Firewall Device Manager.

Use the **Objects** > **Identity Sources** page to create and manage your sources. You would then use these objects when you configure the services that require an identity source

Following are the supported identity sources and their uses:

Active Directory (AD) Identity Realm

Active Directory provides user account and authentication information. See Active Directory (AD) Identity Realms, on page 159.

You can use this source for the following purposes:

- Remote Access VPN, as a primary identity source. You can use AD in conjunction with a RADIUS server.
- Identity policy, for active authentication and as the user identity source used with passive authentication.

AD (Active Directory) Realm Sequence

An AD realm sequence is an ordered list of AD realm objects. Realm sequences are useful if you manage more than one AD domain in your network. See Configuring an AD Realm Sequence, on page 163.

You can use this source for the following purposes:

- Identity policy, as the user identity source used with passive authentication. The order of realms in the sequence determines how the system determines user identity in the rare cases there is a conflict.

Cisco Identity Services Engine (ISE) or Cisco Identity Services Engine Passive Identity Connector (ISE PIC)

If you are using ISE, you can integrate the Firewall Threat Defense device with your ISE deployment. See Identity Services Engine (ISE), on page 169.

You can use this source for the following purposes:

- Identity policy, as a passive identity source to collect user identity from ISE.

RADIUS Server, RADIUS Server Group

If you are using RADIUS servers, you can also use them with the Firewall Device Manager. You must define each server as a separate object, then put them in server groups (where the servers in a given group are copies of each other). You assign the server group to features, you do not assign individual servers. See RADIUS Servers and Groups, on page 165.

You can use this source for the following purposes:

- Remote Access VPN, as an identity source for authentication, and for authorization and accounting. You can use AD in conjunction with a RADIUS server.
- Identity policy, as a passive identity source to collect user identity from remote access VPN logins.
- External authentication for the Firewall Device Manager or the Firewall Threat Defense CLI management users. You can support multiple management users with different authorization levels. These users can log into the system for device configuration and monitoring purposes.

SAML Server

Security Assertion Markup Language 2.0 (SAML 2.0) is an open standard for exchanging authentication and authorization data between parties, specifically an Identity Provider (IdP) and Service Provider (SP).

You can use this source for the following purposes:

- Remote access VPN, as a single sign-on (SSO) authentication source.

LocalIdentitySource

This is the local user database, which includes users that you have defined in the Firewall Device Manager. Select **Objects** > **Users** to manage the user accounts in this database. See Local Users, on page 175.



Note The local identity source database does not include users you configure in the CLI for CLI access (using the **configure user add** command). CLI users are completely separate from those you create in the Firewall Device Manager.

You can use this source for the following purposes:

- Remote Access VPN, as a primary or fallback identity source.
- Identity policy, as a passive identity source to collect user identity from remote access VPN logins.

Active Directory (AD) Identity Realms

Microsoft Active Directory (AD) defines user accounts. You can create an AD identity realm for an Active Directory domain. The following topics explain how to define an AD identity realm.

Supported Directory Servers

You can use Microsoft Active Directory (AD) on Windows Server 2012, 2016, and 2019.

Note the following about your server configuration:

- If you want to perform user control on user groups or on users within groups, you must configure user groups on the directory server. The system cannot perform user group control if the server organizes the users in basic object hierarchy.
- If you perform access control based on user groups, note that Active Directory servers limit the number of users they report to 5000 per group. If necessary, you can modify your Active Directory server configuration to increase this default limit and accommodate more users.
- The directory server must use the field names listed in the following table in order for the system to retrieve user metadata from the servers for that field.

Metadata	Active Directory Field
LDAP user name	samaccountname
first name	givenname
last name	sn
email address	mail userprincipalname (if mail has no value)
department	department distinguishedname (if department has no value)
telephone number	telephonenumber

Limitations on Number of Users

Firewall Device Manager can download information on up to 50,000 users from the directory server.

If your directory server includes more than 50,000 user accounts, you will not see all possible names when selecting users in an access rule or when viewing user-based dashboard information. You can write rules on only those names that were downloaded.

The limit also applies to the names associated with groups. If a group has more than 50,000 members, only the 50,000 names that were downloaded can be matched against the group membership.

Note The Secure Firewall 200 can download a cumulative total of 10,000 user IPs, SXP/SGT mappings, endpoint profiles, and dynamic objects. After a total of 10,000 have been downloaded, the device stops downloading any objects until some previously downloaded objects have been removed. (For example, if users log out, the memory is free for other objects.)

Determining the Directory Base DN

When you configure directory properties, you need to specify the common base distinguished name (DN) for users and groups. The base is defined in your directory server, and differs from network to network. You must enter the correct bases for identity policies to work. If the base is wrong, the system cannot determine user or group names, and thus identity-based policies will be inoperable.

Tip To get the correct bases, consult the administrator who is responsible for the directory servers.

For active directory, you can determine the correct bases by logging into the Active Directory server as domain administrator, and using the **dsquery** command at a command prompt as follows to determine the bases:

User search base

Enter the **dsquery user** command with a known username (partial or complete) to determine the base distinguished name. For example, the following command uses the partial name "John*" to return information for all users that start with "John."

```
C:\Users\Administrator>dsquery user -name "John*"
"CN=John Doe,CN=Users,DC=csc-lab,DC=example,DC=com"
```

The base DN would be "DC=csc-lab,DC=example,DC=com."

Group search base

Enter the **dsquery group** command with a known group name to determine the base distinguished name. For example, the following command uses the group name Employees to return the distinguished name:

```
C:\>dsquery group -name "Employees"
"CN=Employees,CN=Users,DC=csc-lab,DC=example,DC=com"
```

The group base DN would be "DC=csc-lab,DC=example,DC=com."

You can also use the ADSI Edit program to browse the Active Directory structure (**Start** > **Run** > **adsiedit.msc**). In ADSI Edit, right click any object, such as an organizational unit (OU), group, or user, and choose **Properties** to view the distinguished name. You can then copy the string of DC values as the base.

To verify that you have the correct base:

1. Click the Test Connection button in the directory properties to verify connectivity. Resolve any problems, and save the directory properties.

2. Commit changes to the device.
3. Create an access rule, select the **Users** tab, and try to add known user and group names from the directory. You should see auto-complete suggestions as you type for matching users and groups in the realm that contains the directory. If these suggestions appear in a drop-down list, then the system was able to query the directory successfully. If you see no suggestions, and you are certain the string you typed should appear in a user or group name, you need to correct the corresponding search base.

Configuring AD Identity Realms

An identity realm is a directory server plus other attributes required to provide authentication services. The directory server contains information about the users and user groups who are allowed access to your network.

For Active Directory, a realm is equivalent to an Active Directory domain. Create separate realms for each AD domain you need to support.

Realms are used in the following policies:

- Identity—The realm provides user identity and group membership information, which you can then use in access control rules. The system downloads updated information about all users and groups every day in the last hour of the day (UTC). The directory server must be reachable from the management interface.
- Remote access VPN—The realm provides authentication services, which determine whether a connection is allowed. The directory server must be reachable from the RA VPN outside interface.
- Access Control, SSL Decryption—You can select the realm in the user criteria for the rule to apply the rule to all users within the realm.

Work with your directory administrator to get the values required to configure the directory server properties.

Note If the directory server is not on an attached network or available through the default route, create a static route for the server. Select **Device** > **Routing** > **View Configuration** to create static routes.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create identity realm objects while editing a realm property by clicking the **Create New Identity Realm** link shown in the object list.

Before you begin

Ensure that time settings are consistent among the directory servers, the Firewall Threat Defense device, and clients. A time shift among these devices can prevent successful user authentication. "Consistent" means that you can use different time zones, but the time should be the same relative to those zones; for example, 10 AM PST = 1 PM EST.

Procedure

Step 1 Select **Objects**, then select **Identity Sources** from the table of contents.

Step 2 Do one of the following:

- To create an AD realm, click + > **AD**.
- To edit a realm, click the edit icon ( ) for the realm.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 3 Configure the basic realm properties.

- **Name**—A name for the directory realm.
- **Type**—The type of directory server. Active Directory is the only supported type, and you cannot change this field.
- **Directory Username**, **Directory Password**—The distinguished username and password for a user with appropriate rights to the user information you want to retrieve. For Active Directory, the user does not need elevated privileges. You can specify any user in the domain. The username must be fully qualified; for example, Administrator@example.com (not simply Administrator).

  Note
  The system generates ldap-login-dn and ldap-login-password from this information. For example, Administrator@example.com is translated as cn=administrator,cn=users,dc=example,dc=com. Note that cn=users is always part of this translation, so you must configure the user you specify here under the common name "users" folder.

- **Base DN**—The directory tree for searching or querying user and group information, that is, the common parent for users and groups. For example, cn=users,dc=example,dc=com. For information on finding the base DN, see Determining the Directory Base DN, on page 160.
- **AD Primary Domain**— The fully qualified Active Directory domain name that the device should join. For example, example.com.

Step 4 Configure the directory server properties.

- **Hostname/IP Address**—The hostname or IP address of the directory server. If you use an encrypted connection to the server, you must enter the fully-qualified domain name, not the IP address.
- **Port**—The port number used for communications with the server. The default is 389. Use port 636 if you select LDAPS as the encryption method.
- **Encryption**—To use an encrypted connection for downloading user and group information, select the desired method, **STARTTLS** or **LDAPS**. The default is **None**, which means that user and group information is downloaded in clear text.
  - **STARTTLS** negotiates the encryption method, and uses the strongest method supported by the directory server. Use port 389. This option is not supported if you use the realm for remote access VPN.
  - **LDAPS** requires LDAP over SSL. Use port 636.
- **Trusted CA Certificate**—If you select an encryption method, upload a Certificate Authority (CA) certificate to enable a trusted connection between the system and the directory server. 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 ad.example.com in the certificate, the connection fails.

Step 5 If there are multiple servers for the realm, click **Add Another Configuration** and enter the properties for each additional server.

You can add up to 10 AD servers to the realm. These servers need to be duplicates of each other and support the same AD domain.

You can collapse and expand each server entry for your convenience. The sections are labeled with the hostname/IP address and port.

Step 6 Click the **Test** button to verify the system can contact the server.

The system uses separate processes and interfaces to access the server, so you might get errors indicating that the connection works for one type of use but not another, for example, available for Identity policies but not for remote access VPN. If the server cannot be reached, verify that you have the right IP address and host name, that the DNS server has an entry for the hostname, and so forth. You might need to configure a static route for the server. For more information, see Troubleshooting Directory Server Connections, on page 164.

Step 7 Click **OK**.

Configuring an AD Realm Sequence

You can use an AD realm sequence in a passive identity rule so that the system can try to match users from more than one AD server. In a realm sequence, you configure an ordered list of AD realms where each AD server manages a different realm or domain, for example, engineering.example.com and marketing.example.com.

Realm sequences are useful only if you support more than one AD domain, and users from different domains might send traffic through the Firewall Threat Defense device. The realms are used in order to find the identity for a passively-authenticated user session. The order of the realms is used to resolve identity conflicts, in the rare cases where a conflict might arise.

Procedure

Step 1 Select **Objects**, then select **Identity Sources** from the table of contents.

Step 2 Do one of the following:

- To create an AD realm sequence, click **+** > **AD Realm Sequence**.
- To edit an AD realm sequence, click the edit icon () for the object.

To delete an unreferenced object, click the trash can icon () for the object.

Step 3 Configure the realm sequence properties:

- **Name**—The name of the object.
- **Description**—An optional description of the object.
- **AD Realms**—Click **+** to add AD realm objects to the sequence. After you add the realms, click and drag and drop the realms into the desired ordered sequence.

Step 4 Click **OK**.

You can now select the AD realm sequence in a passive identity rule.

Troubleshooting Directory Server Connections

The system uses different processes to communicate with your directory server depending on the feature. Thus, a connection for identity policies might work, whereas one for remote access VPN fails.

These processes use different interfaces to communicate with the directory server. You must ensure connectivity from these interfaces.

- Management interface, for: identity policies.
- Data interface, for: remote access VPN (outside interface).

When you configure the identity realm, use the **Test** button to verify that the connection can work. Failure messages should indicate the feature that is having connection problems. The following are the general issues you might encounter, based on authentication attributes and routing/interface configuration.

Directory user authentication issues.

If the problem is that the system could not log into the directory server because of the username or password, ensure that the name and password are correct and valid on the directory server. For Active Directory, the user does not need elevated privileges. You can specify any user in the domain. The username must be fully qualified; for example, Administrator@example.com (not simply Administrator).

Also, the system generates ldap-login-dn and ldap-login-password from the username and password information. For example, Administrator@example.com is translated as cn=administrator,cn=users,dc=example,dc=com. Note that cn=users is always part of this translation, so you must configure the user you specify here under the common name “users” folder.

The directory server is accessible through a data interface.

If the directory server is on a network that is either directly connected to a data interface (such as a GigabitEthernet interface), or routeable from a directly-connected network, you must ensure that there is a route between the virtual management interface and the directory server.

- Using **data-interfaces** as the management gateway should make routing successful.
- If you have an explicit gateway on the management interface, that gateway router needs to have a route to the directory server.
- You do not need to configure an IP address on the **diagnostic** interface, which is the physical interface used by the virtual management interface. However, if you do configure an address, do not also configure a static route (such as a default route) that would redirect traffic to the directory server to the diagnostic interface.
- If there is a router between the directly-connected network and the network that hosts the directory server, configure a static route for the directory server (**Device** > **Routing**).
- Verify that the data interface has the correct IP address and subnet mask.

The directory server is accessible through the Management physical interface.

If the directory server is on the network that is either directly connected to the Management physical interface (such as Management0/0) or routeable from that network, you must do the following:

- Configure an IPv4 address for the Management interface (with the logical name **diagnostic**) on **Device** > **Interfaces**. The IP address must be on the same subnet as the virtual management address (**Device** > **System Settings** > **Management Interface**).
- If there is a router between the directory server and the Management interface, configure a route for the directory server on **Device** > **Routing** for the **diagnostic** interface.
- Verify that the diagnostic and management interfaces have the correct IP address and subnet mask.

The directory server is on an external network.

If the directory server is on a network on the other side of the outside (uplink) interface, you might need to configure a site-to-site VPN connection. For the detailed procedure, see How to Use a Directory Server on an Outside Network with Remote Access VPN, on page 720.

RADIUS Servers and Groups

You can use RADIUS servers to authenticate and authorize remote access VPN connections, and the Firewall Device Manager and the Firewall Threat Defense CLI administration users. For example, if you also use Cisco Identity Services Engine (ISE) and its RADIUS server, you can use that server with the Firewall Device Manager.

When you configure a feature to use RADIUS servers, you select a RADIUS group instead of individual servers. A RADIUS group is a collection of RADIUS servers that are copies of each other. If a group has more than one server, they form a chain of backup servers to provide redundancy in case one server becomes unavailable. But even if you have only one server, you must create a one-member group to configure RADIUS support for a feature.

The following topics explain how to configure RADIUS servers and groups, so that they are available for use in the supported features.

Configure RADIUS Servers

RADIUS servers provide AAA (authentication, authorization, and accounting) services. If you use RADIUS servers to authenticate and authorize users, you can use those servers with the Firewall Device Manager.

After creating objects for each of your RADIUS servers, create RADIUS server groups to contain each group of duplicate servers.

Before you begin

If you want to configure a redirect ACL for RA VPN, you must use Smart CLI to create the extended ACL before creating or editing the server object. You cannot create the ACL while editing the object.

Procedure

Step 1 Select **Objects**, then select **Identity Sources** from the table of contents.

Step 2 Do one of the following:

- To create an object, click **+** > **RADIUS Server**.
- To edit an object, click the edit icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 3 Configure the following properties:

- **Name**—The name of the object. This does not have to match anything configured on the server.
- **Server Name or IP Address**—The fully-qualified host name (FQDN) or IP address of the server. For example, radius.example.com or 10.100.10.10.
- **Authentication Port**—The port on which RADIUS authentication and authorization are performed. The default is 1812.
- **Timeout**—The length of time, 1-300 seconds, that the system waits for a response from the server before sending the request to the next server. The default is 10 seconds. If you are using this server as a secondary authentication source for remote access VPN, for example, to prompt for an authentication token, increase this timeout to 60 seconds at least. This provides time for the user to obtain and enter the token.
- **Server Secret Key**—(Optional.) The shared secret that is used to encrypt data between the device and the RADIUS server. The key is a case-sensitive, alphanumeric string of up to 64 characters, with no spaces. The key must start with an alphanumeric character or an underscore, and it can contain the special characters: $ - _ . + @. The string must match the one configured on the RADIUS server. If you do not configure a secret key, the connection is not encrypted.
- **Require Message-Authenticator for all RADIUS Responses**—The Message-Authenticator attribute is used to protect against Blast-RADIUS attacks. If you have upgraded your RADIUS server so it supports the message authenticator, you can enable this option to help protect against these attacks. When enabled, all requests and responses must have the message authenticator, or authentication will fail.

  Requires Firewall Threat Defense Version 7.0.7+. For the Firepower 4100/9300, this feature may require an FXOS upgrade; for minimum builds, see *Cisco Secure Firewall Threat Defense Compatibility Guide*.

Step 4 (Optional.) If you are using the server for remote access VPN Change of Authorization configuration, you can click the **RA VPN Only** link and configure the following options.

- **Redirect ACL**—Select the extended ACL to use for the RA VPN redirect ACL. Create extended ACLs using the Smart CLI **Extended Access List** object on the **Device** > **Advanced Configuration** > **Smart CLI** > **Objects** page.

  The purpose of the redirect ACL is to send initial traffic to Cisco Identity Services Engine (ISE) so that ISE can assess the client posture. The ACL should send HTTPS traffic to ISE, but not traffic that is already destined for ISE, or traffic that is directed to a DNS server for name resolution. For an example, see Configure Change of Authorization on the Firewall Threat Defense Device, on page 703.
- **Interface Used to Connect to RADIUS Server**—Which interface to use when communicating with the server. If you select **Resolve via Route Lookup**, the system always uses the routing table to determine

the interface to use. If you select **Manually Choose Interface**, the system will always use the interface you select.

If you are configuring Change of Authorization, you must select a specific interface so that the system can correctly enable the CoA listener on the interface.

If the server is on the same network as the management address, which means you will select the Diagnostic interface, you must also configure an IP address on the Diagnostic interface. Having a management IP address is not sufficient. Go to **Device** > **Interfaces**, and configure an IP address on the diagnostic interface that is on the same subnet as the management IP address.

If you also use this server for the Firewall Device Manager administrative access, this interface is ignored. Administrative access attempts are always authenticated through the management IP address.

Step 5 (Optional, when editing the object only.) Click **Test** to check whether the system can connect to the server.

You are prompted for a username and password. The test confirms whether the server can be contacted, and if it can, that the username can be authenticated.

Step 6 Click **OK**.

Configure RADIUS Server Groups

A RADIUS server group contains one or more RADIUS server objects. The servers within a group must be copies of each other. These servers form a chain of backup servers, so that if the first server is unavailable, the system can try the next server in the list.

When you configure RADIUS support in a feature, you must select a server group. Thus, even if you have just one RADIUS server, you must create a server group to contain it.

Procedure

Step 1 Select **Objects**, then select **Identity Sources** from the table of contents.

Step 2 Do one of the following:

- To create an object, click + > **RADIUS Server Group**.
- To edit an object, click the edit icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 3 Configure the following properties:

- **Name**—The name of the object. This does not have to match anything configured on the servers.
- **Dead Time**—Failed servers are reactivated only after all servers have failed. The dead time is how long to wait, from 0 - 1440 minutes, after the last server fails before reactivating all servers. Deadtime applies only if you configure fallback to the local database; authentication is attempted locally until the deadtime elapses. The default is 10 minutes.
- **Maximum Failed Attempts**—The number of failed AAA transactions (that is, requests that do not get a response) sent to a RADIUS server in the group before trying the next server. You can specify 1-5, and

the default is 3. When the maximum number of failed attempts is exceeded, the system marks the server as Failed.

For a given feature, if you configured a fallback method using the local database, and all the servers in the group fail to respond, then the group is considered to be unresponsive, and the fallback method is tried. The server group remains marked as unresponsive for the duration of the dead time, so that additional AAA requests within that period do not attempt to contact the server group, and the fallback method is used immediately.

- **Dynamic Authorization (for RA VPN only)**, **Port**—If you enable RADIUS dynamic authorization or change of authorization (CoA) services for this RADIUS server group, the group will be registered for CoA notification and listen on the specified port for CoA policy updates from Cisco Identity Services Engine (ISE). The default listening port is 1700, or you can specify a different port in the range 1024 to 65535. Enable dynamic authorization only if you are using this server group in a remote access VPN in conjunction with ISE
- **Realm that Supports the RADIUS Server**—If the RADIUS server is configured to use an AD server for authenticating users, select the AD realm that specifies the AD server used in conjunction with this RADIUS server. If the realm does not already exist, click **Create New Identity Realm** at the bottom of the list and configure it now.
- **RADIUS Server list**—Select up to 16 RADIUS server objects that define the servers for the group. Add these objects in priority order. The first server in the list is used until it becomes unresponsive. After adding the objects, you can drag and drop to rearrange them. If the object you need does not yet exist, click **Create New RADIUS Server** and add it now.

  You can also click the **Test** link to verify the system can connect to the server. You are prompted for a username and password. The test confirms whether the server can be contacted, and if it can, that the username can be authenticated.

Step 4 (Optional.) Click the **Test All Servers** button to check connectivity to each server in the group.

You are prompted for a username and password. The system checks whether each server can be contacted, and whether the username can be authenticated on each server.

Step 5 Click **OK**.

Troubleshoot RADIUS Servers and Groups

Following are some things you can check if external authorization does not work.

- Use the **Test** buttons in the RADIUS server and server group objects to verify that the servers can be contacted from the device. Ensure that you save the objects before testing. If the test fails:
  - Please understand that the test ignores the interface configured for the server, and always uses the management interface. The test is expected to fail if the RADIUS authentication proxy is not configured to respond to requests from the management IP address.
  - Verify that you are entering a correct username/password combination during the test. You should get a Bad Credentials message if they are incorrect.
  - Verify the secret key, port, and IP address for the server. If you are using a hostname, verify that DNS is configured for the management interface. Consider the possibility that the secret key was changed on the RADIUS server but not in the device configuration.

  - If the test continues to fail, you might need to configure a static route to the RADIUS servers. Try pinging the server from the CLI Console or an SSH session to see if it can be reached.

- If external authentication has been working, but has stopped working, consider the possibility that all servers are in the dead time. If you configure fallback to local authentication, when all the RADIUS servers within a group have failed, the dead time is the number of minutes the system waits before trying the first server again. During the dead time, local authentication is used, so a given user's username and password would be the local username/password. The default is 10 minutes, but you can configure as long as 1440 minutes.
- If HTTPS external authentication works for some users but not others, evaluate the cisco-av-pair attribute defined in the RADIUS server for each user account. This attribute might be configured incorrectly. A missing or incorrect attribute will block all HTTS access for that user account.
- If SSH external authentication works for some users but not for others, evaluate the Service-Type attribute defined in the RADIUS server for each user account. This attribute might be configured incorrectly. A missing or incorrect attribute will block all SSH access for that user account.

Identity Services Engine (ISE)

You can integrate your Cisco Identity Services Engine (ISE) or ISE Passive Identity Connector (ISE-PIC) deployment with the Firewall Threat Defense device to use ISE/ISE-PIC for passive authentication.

ISE/ISE-PIC is an authoritative identity source, and provides user awareness data for users who authenticate using Active Directory (AD), LDAP, RADIUS, or RSA. However, for Firewall Threat Defense, you can use ISE for user identity awareness in conjunction with AD only. You can use the user identity in access control and SSL decryption policies as matching criteria, in addition to seeing user information in the various monitoring dashboards and events.

For more information on Cisco ISE/ISE-PIC, see the *Cisco Identity Services Engine Administrator Guide* (https://www.cisco.com/c/en/us/support/security/identity-services-engine/tsd-products-support-series-home.html) and the *Identity Services Engine Passive Identity Connector (ISE-PIC) Installation and Administrator Guide* (https://www.cisco.com/c/en/us/support/security/ise-passive-identity-connector/tsd-products-support-series-home.html).

Guidelines and Limitations for ISE

- The firewall system does not support 802.1x device authentication alongside Active Directory authentication because the system does not associate device authentication with users. If you use 802.1x active logins, configure ISE to report only 802.1x active logins (both device and user). That way, a device login is reported only once to the system.
- ISE/ISE-PIC does not report the activity of ISE Guest Services users.
- Synchronize the time on the ISE/ISE-PIC server and the device. Otherwise, the system might perform user timeouts at unexpected intervals.
- If you configure ISE/ISE-PIC to monitor a large number of user groups, the system might drop user mappings based on groups due to memory limitations. As a result, rules with realm or user conditions might not perform as expected.

- For the specific versions of ISE/ISE-PIC that are compatible with this version of the system, see the *Cisco Secure Firewall Compatibility Guide*, https://www.cisco.com/c/en/us/support/security/firepower-ngfw/products-device-support-tables-list.html.
- Use the IPv4 address of the ISE server, unless you confirm that your version of ISE supports IPv6.

Configure Identity Services Engine

To use the Cisco Identity Services Engine (ISE) or Cisco Identity Services Engine Passive Identity Connector (ISE PIC) as a passive identity source, you must configure the connection to the ISE Platform Exchange Grid (pxGrid) server.

Before you begin

- Export the pxGrid and MNT server certificates from ISE. For example, in ISE PIC 2.2, you find these on the **Certificates** > **Certificate Management** > **System Certificates** page. The MNT (Monitoring and Troubleshooting node) is shown as Admin in the Used By column in the certificates list. You can either upload them as trusted CA certificates on the **Objects** > **Certificates** page, or upload them during the following procedure. These nodes might be using the same certificate.
- You must also configure an AD identity realm. The system obtains the list of users from AD, and from ISE it gets information on the user-to-IP address mappings.
- If you will use security group tags (SGT) for access control, with or without static security group tag mappings, and listen to the SXP topic, you also need to configure SXP and these mappings in ISE. See Configure security groups and SXP publishing in ISE, on page 504.

Procedure

Step 1 Select **Objects**, then select **Identity Sources** from the table of contents.

Step 2 Do one of the following:

- To create an object, click + > **Identity Services Engine**. You can create at most one ISE object.
- To edit an object, click the edit icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 3 Configure the following properties:

- **Name**—The name of the object.
- **Status**—Click the toggle to enable or disable the object. When disabled, you cannot use ISE as an identity source in your identity rules.
- **Description**—An optional description of the object.
- **Primary Node Hostname/IP Address**—The hostname or IP address for the primary pxGrid ISE server. Do not specify an IPv6 address unless you verify that your version of ISE supports IPv6.

- **Secondary Node Hostname/IP Address**—If you set up a secondary ISE server for high availability, click **Add Secondary Node Hostname/IP Address** and enter the hostname or IP address of the secondary pxGrid ISE server.
- **pxGrid Server CA Certificate**—The trusted Certificate Authority certificate 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 Certificate**—The trusted Certificate Authority certificate for the ISE certificate when performing bulk downloads. This can be the same as the pxGrid server certificate if your MNT (Monitoring and Troubleshooting) server is not separate. If your deployment includes a primary and a secondary MNT node, the certificates for both nodes must be signed by the same certificate authority.
- **Server Certificate**—The internal identity certificate that the Firewall Threat Defense device must provide to ISE when connecting to ISE or performing bulk downloads.
- **Subscribe To**—Select which ISE pxGrid topics should be subscribed to. Subscribing to a topic means you will download data related to the topic.
  - **Session Directory Topic**—Whether to obtain information about user sessions, including SGT mappings for user sessions. This option is enabled by default. You should select this option if you want to obtain passive user identity for use in security policies and for visibility in the monitoring dashboards.
  - **SXP Topic**—Whether to obtain static SGT-to-IP address mappings. Select this topic if you want to write access control rules based on security group tags (SGT).
- **ISE Network Filters**—An optional filter you can set to restrict the data that ISE reports to the system. If you provide a network filter, ISE reports data from the networks within the filter only. Click **+**, select the network objects that identify the networks, and click **OK**. Click **Create New Network** if you need to create the objects. Configure IPv4 network objects only.

Step 4 Click the **Test** button to verify that the system can connect to your ISE server.

If the test fails, click the **See Logs** link to read the detailed error messages. For example, the following message indicates that the system could not connect to the server at the required port. The problem might be no route to the host, that the ISE server is not using the expected port, or that you have access control rules that prevent the connection.

```
Captured Jabberwerx log:2018-05-11T16:10:30 [   ERROR]: connection timed out while
trying to test connection to host=10.88.127.142:ip=10.88.127.142:port=5222
```

Step 5 Click **OK** to save the object.

What to do next

After you configure ISE, enable the identity policy, configure passive authentication rules, and deploy the configuration. Then, you must go into ISE/ISE PIC and accept the device as a subscriber. If you configure ISE/ISE PIC to automatically accept subscribers, you do not need to manually accept the subscription.

Troubleshoot the ISE/ISE-PIC Identity Source

ISE/ISE-PIC Connections

If you experience issues with the ISE or ISE-PIC connection, check the following:

- The pxGrid Identity Mapping feature in ISE must be enabled before you can successfully integrate ISE with the Firewall Threat Defense device.
- Before a connection between the ISE server and the Firewall Threat Defense device succeeds, you must manually approve the clients in ISE.

  Alternatively, you can 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 Firewall Threat Defense device (server) certificate must include the **clientAuth** extended key usage value, or it must not include any extended key usage values. If the clientAuth extended key usage is set, then there must also either be no key usage set, or the Digital Signature key usage value must be set. The self-signed identity certificates you can create using the Firewall Device Manager meet these requirements.
- The time on your ISE server must be synchronized with the time on the Firewall Threat Defense device. If the appliances are not synchronized, the system might perform user timeouts at unexpected intervals.

ISE/ISE-PIC User Data

If you experience issues with user data reported by ISE or ISE-PIC, 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 dashboards 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 system does not receive user data for ISE Guest Services users.

SAML Servers

You can configure Security Assertion Markup Language 2.0 (SAML 2.0) servers to use as single sign-on (SSO) authentication sources for remote access VPN connections. SAML is an open standard for exchanging authentication and authorization data between parties, specifically an Identity Provider (IdP) and Service Provider (SP).

Configure SAML Servers

You can configure Security Assertion Markup Language 2.0 (SAML 2.0) servers to use as single sign-on (SSO) authentication sources for remote access VPN connections. For example, the Duo Access Gateway (DAG) is a SAML server.

When you use a SAML server as the authentication method, the SAML server acts as the Identity Provider (IdP), whereas the Firewall Threat Defense device acts as the Service Provider (SP).

For RA VPN, you can use a SAML server as the primary authentication source, but you cannot configure a secondary authentication source, nor can you configure a fallback source.

Before you begin

Obtain the following information from the SAML server identity provider.

- Entity ID URL, which provides the SAML server metadata.
- Sign-in URL.
- Sign-out URL.
- Identity provider certificate.

The system does not support multiple values returned for a SAML attribute. For example, if you use SAML with RA VPN, you cannot provide multiple values for the Cisco Group Policy attribute: you must provide one value only. If you send multiple values, a default value is applied rather than picking one of the provided values. Ensure that the SAML server is configured to provide a single value for each attribute returned.

Procedure

Step 1 Do one of the following to get to the SAML Servers page:

- Select **Objects**, then select **Identity Sources** from the table of contents.
- Select **Device** > **Remote Access VPN** > **SAML Servers**.

Step 2 Do one of the following:

- To create an object, click **+** > **SAML Server**.
- To edit an object, click the edit icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 3 Configure the following properties:

- **Name**—The name of the object.
- **Description**—An optional description of the object.
- **Identity Provider (IDP) Entity ID URL**—The URL for a page that serves a metadata XML that describes how the SAML Issuer is going to respond to requests. Some SAML server products call this the Entity ID, others call it the Metadata URL. The URL must be from 4-128 characters, including the protocol, https://. For example, https://191.168.2.21/dag/saml2/idp/metadata.php.
- **Sign-In URL**—The URL for signing into the identity provider SAML server. The URL must be between 4-500 characters, including protocol. Both http:// and https:// are allowed. For example, https://191.168.2.21/dag/saml2/idp/SSOService.php.

- **Sign-Out URL**—The URL for signing out of the identity provider SAML server. The URL must be between 4-500 characters, including protocol. Both http:// and https:// are allowed. For example, https://191.168.2.21/dag/saml2/idp/SingleLogoutService.php.
- **Service Provider Certificate**—The internal certificate to use for the Firewall Threat Defense device. Ideally, you have already uploaded a certificate signed by a recognized third party, and you can select it now. You can also use the built-in DefaultInternalCertificate, or click **Create New Internal Certificate** and upload a signed certificate now. The SAML server identity provider will have to trust this certificate, so you might need to upload it to the SAML server. Consult the SAML server documentation for information on how to upload certificates or otherwise enable a trust relationship with a service provider.
- **Identity Provider Certificate**—The trusted CA certificate for the SAML server identity provider. Download this certificate from the SAML server. If you have not already uploaded it, click **Create New Trusted CA Certificate** and upload it now.
- **Request Signature**—The encryption algorithm to use when signing login request. Select None to disable encryption. Otherwise, choose one of the following, which are ordered from weakest to strongest: SHA1, SHA256, SHA384, SHA512.
- **Request Timeout**—SAML assertions have a valid period: the user must complete the single sign-on request within the valid period. You can set a timeout, in seconds, to change this period. If you set a timeout that is longer than the assertion's NotOnOrAfter condition, your timeout is ignored and NotOnOrAfter is honored. The range is 1-7200 seconds. The default is 300 seconds.
- **This SAML identity provider (IDP) is on an internal network**—Whether the SAML server is operating on an internal network, rather than external to the protected networks.
- **Request IDP re-authentication at login**—Select this option for the user to re-authenticate at each login, rather than have the SAML server re-use a previous authentication session. This option is enabled by default.

Step 4 Click **OK**.

What to do next

If you enabled **Request Signature** to encrypt communications, you need to upload the device manager information to the SAML server. From the identity sources list, click the **Download** ( ) button for the server, and save the XML file. Then, log into the SAML server and upload the information. See your SAML provider documentation for detailed information.

If you use the server for device manager login, and it does not work, verify the SAML server configuration.

- Log In into the SAML IdP and verify that the device manager SAML Response consumer is configured correctly. The value should be: https://<FDM_URL>/api/fdm/latest/fdm/token
- If signing is enabled in the SAML server object, make sure that the device manager public certificate is uploaded into the SAML application, and that encryption is enabled. Uploading the device manager XML file should add the certificate to the SAML server. You can also retrieve the device manager certificate through the FDM API: https://<FDM_URL>/saml/metadata

Local Users

The local user database (LocalIdentitySource) includes users that you have defined in the Firewall Threat Defense.

You can use locally-defined users for the following purposes:

- Remote Access VPN, as a primary or fallback identity source.
- Management access, as a primary or secondary source for the Firewall Threat Defense users.

  The **admin** user is a system-defined local user. However, the admin user cannot log into a remote access VPN. You cannot create additional local administrative users.

  If you define external authentication for management access, external users who log into the device appear on the local users list.
- Identity policy, indirectly, as a passive identity source to collect user identity from remote access VPN logins.

The following topic explains how to configure local users.

Configure Local Users

You can create user accounts directly on the device for use with remote access VPN. You can use the local user accounts instead of, or in addition to, an external authentication source.

If you use the local user database as a fallback authentication method for remote access VPN, ensure that you configure the same usernames/passwords in the local database as the names in the external database. Otherwise, the fallback mechanism will be ineffective.

The users defined here cannot log into the device CLI.

Procedure

Step 1 Select **Objects** > **Users**.

The list shows the usernames and service types, which can be:

- MGMT—For administrative users who can log into the Firewall Device Manager. The admin user is always defined, and you cannot delete it. You also cannot configure additional MGMT users. However, if you define external authentication for management access, external users who log into the device appear on the local users list as MGMT users.
- RA VPN—For users who can log into a remote access VPN configured on the device. You must also select the local database for the primary or secondary (fallback) source.

Step 2 Do one of the following:

- To add a user, click +.
- To edit a user, click the edit icon ( ) for the user.

If you no longer need a particular user account, click the delete icon ( ) for the user.

Step 3 Configure the user properties:

The name and password can contain any printable ASCII alphanumeric or special character except spaces and question marks. Printable characters are ASCII codes 33-126.

- **Name**—The username for logging into the remote access VPN. The name can be 4-64 characters, and it cannot contain spaces. For example, johndoe.
- **Password**, **Confirm Password**—Enter the password for the account. The password must be 8-16 characters long. It cannot contain consecutive letters that are the same. It must also contain at least one of each of the following: number, upper and lower case characters, and a special character.

Note
Users cannot change their passwords. Notify them of their passwords, and when they need to change them, you must edit the user account. Also, do not update the password for external MGMT users: the passwords are controlled by the external AAA server.

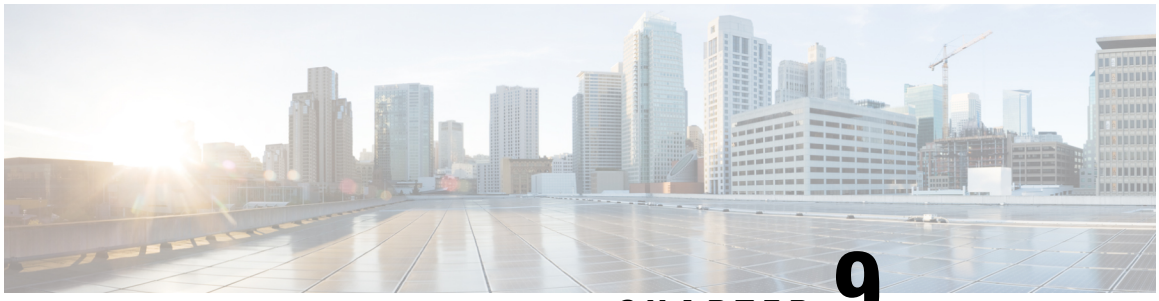
Step 4 Click **OK**.



PART III

The Basics

- Logical Devices on the Firepower 4100/9300, on page 179
- High Availability (Failover), on page 191
- Interfaces, on page 235



CHAPTER 9

Logical Devices on the Firepower 4100/9300

The Firepower 4100/9300 is a flexible security platform on which you can install one or more *logical devices*.

You must configure chassis interfaces, add a logical device, and assign interfaces to the device on the Firepower 4100/9300 chassis using the Firewall Chassis Manager or the FXOS CLI. You cannot perform these tasks in the Firewall Device Manager.

This chapter describes basic interface configuration and how to add a standalone or High Availability logical device using the Firewall Chassis Manager. 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 179
- Requirements and Prerequisites for Firepower 9300 Hardware and Software Combinations, on page 181
- Guidelines and Limitations for Logical Devices, on page 182
- Configure Interfaces, on page 182
- Configure a Logical Device, on page 184
- History for Logical Devices, on page 190

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 Firewall Chassis Manager. 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, or if the logical device is offline.

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.
- Data-sharing—Use for regular data. Only supported with container instances, these data interfaces can be shared by one or more logical devices/container instances (Firewall Threat Defense-using-Firewall Management Center only).
- 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. For information about the separate chassis management interface, see Chassis Management Interface, on page 179.

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 Firewall Threat Defense-using-Firewall Management Center devices.

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. The Firewall Device Manager and Security Cloud Control does not support clustering.

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.

Requirements and Prerequisites for Firepower 9300 Hardware and Software Combinations

The Firepower 9300 includes 3 security module slots and multiple types of security modules. See the following requirements:

- Security Module Types—You can install modules of different types in the Firepower 9300. For example, you can install the SM-48 as module 1, SM-40 as module 2, and SM-56 as module 3.
- Native and Container instances—When you install a container instance on a security module, that module can only support other container instances. A native instance uses all of the resources for a module, so you can only install a single native instance on a module. You can use native instances on some modules, and container instances on the other module. For example, you can install a native instance on module 1 and module 2, but container instances on module 3.
- High Availability—High Availability is only supported between same-type modules on the Firepower 9300. However, the two chassis can include mixed modules. For example, each chassis has an SM-40, SM-48, and SM-56. You can create High Availability pairs between the SM-40 modules, between the SM-48 modules, and between the SM-56 modules.
- ASA and Firewall Threat Defense application types—You can install different application types on separate modules in the chassis. For example, you can install ASA on module 1 and module 2, and Firewall Threat Defense on module 3.
- ASA or Firewall Threat Defense versions—You can run different versions of an application instance type on separate modules, or as separate container instances on the same module. For example, you can install the Firewall Threat Defense 6.3 on module 1, Firewall Threat Defense 6.4 on module 2, and Firewall Threat Defense 6.5 on module 3.

Guidelines and Limitations for Logical Devices

See the following sections for guidelines and limitations.

Guidelines and Limitations for Interfaces

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

High Availability

- Configure high availability within the application configuration.
- You can use any data interfaces as the failover and state links.
- The two units in a High Availability Failover configuration must:
  - 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 more information, see System Requirements for High Availability, on page 199.

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.
- If you replace an SFP on a port with a different SFP module, the speed, duplex, and auto-negotiation of the interface is not updated automatically. You must manually re-configure the interface.

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.

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.

You can configure each physical Data 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.
- On—The EtherChannel is always on, and LACP is not used. An "on" EtherChannel can only establish a connection with another "on" EtherChannel.



Note It may take up to three minutes for an EtherChannel to come up to an operational state if you change its mode from On to Active or from Active to On.

The Firepower 4100/9300 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. "On" mode cannot use standby interfaces in the channel group when an interface goes down, and the connectivity and configurations are not checked.

When the Firepower 4100/9300 creates an EtherChannel, the EtherChannel stays in a **Suspended** state for Active LACP mode or a **Down** state for On LACP mode 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** or **Down** state.

Configure a Logical Device

Add a standalone logical device or a High Availability pair on the Firepower 4100/9300.

Add a Standalone Firewall Threat Defense for the Firewall Device Manager

You can use the Firewall Device Manager with a native instance. Container instances are not supported. Standalone logical devices work either alone or in a High Availability pair.

Before you begin

- Download the application image you want to use for the logical device from Cisco.com, and then that image to the Firepower 4100/9300.
- 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.
- You must also configure at least one Data type interface.
- Gather the following information:
  - Interface IDs for this device
  - Management interface IP address and network mask
  - Gateway IP address
  - DNS server IP address
  - Firewall Threat Defense hostname and domain name

Procedure

See the Firewall Device Manager configuration guide to start configuring your security policy.

Add a High Availability Pair

Firewall Threat Defense 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 System Requirements for High Availability, on page 199.

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 (Failover), on page 191.

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 Firewall Threat Defense Logical Device

You can allocate or unallocate an interface on the Firewall Threat Defense logical device. You can then sync the interface configuration in the Firewall Management Center or the Firewall Device Manager.

Adding a new interface, or deleting an unused interface has minimal impact on the Firewall 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 Firewall 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 Firewall Management Center or the Firewall Device Manager.

For the Firewall Device Manager: You can migrate the configuration from one interface to another interface before you delete the old interface.

Before you begin

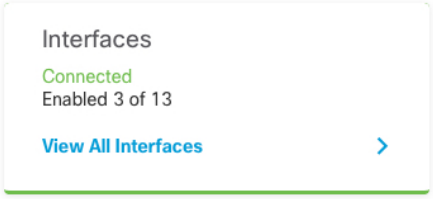
- Configure your interfaces, and add any EtherChannels according to Configure a Physical Interface, on page 183 and Add an EtherChannel (Port Channel), on page 183.
- If you want to add an already-allocated interface to an EtherChannel (for example, all interfaces are allocated by default to a cluster), 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.
- For High Availability, make sure you add or remove the interface on all units before you sync the configuration in the Firewall Management Center or the Firewall Device Manager. We recommend that you make the interface changes on the standby unit first, and then on the 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 Firewall Chassis Manager. Once unallocated, add the new interface and then use sync interfaces from the Firewall Management Center.

Procedure

Step 1 Sync and migrate the interfaces in the Firewall Device Manager.

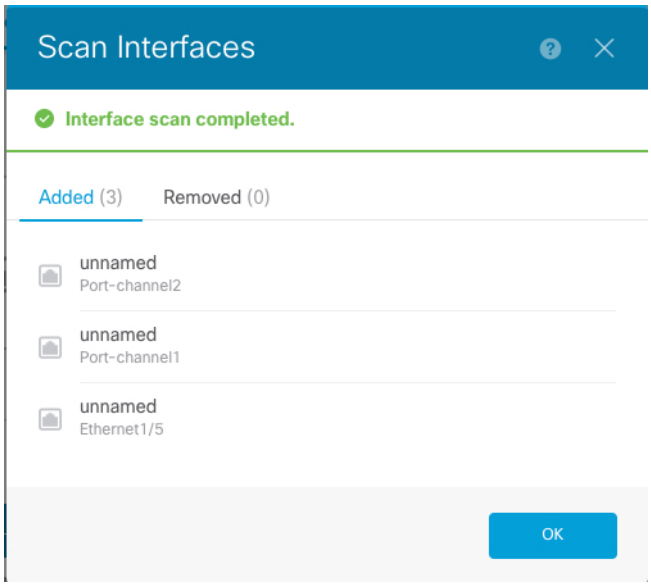
a) Log into the Firewall Device Manager.

b) Click **Device**, then click the **View All Interfaces** link in the **Interfaces** summary.



c) Click the **Scan Interfaces icon**.

d) Wait for the interfaces to scan, and then click **OK**.



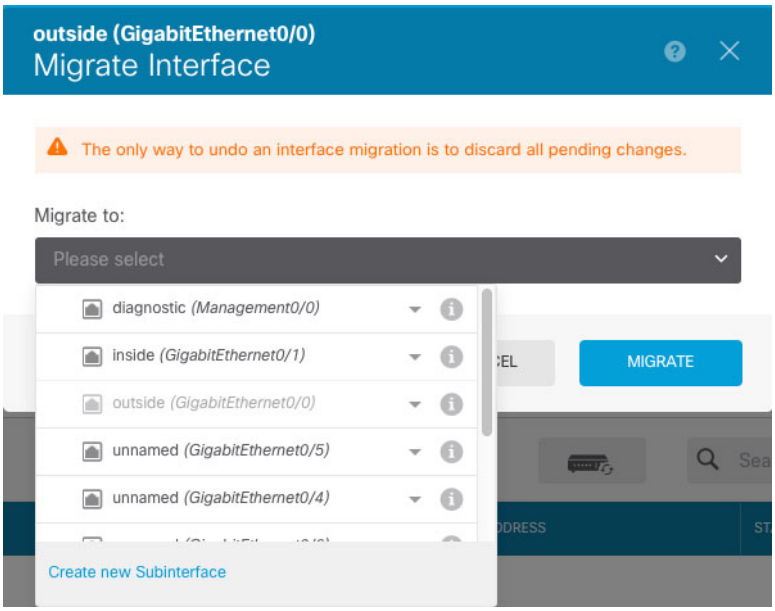
e) Configure the new interfaces with names, IP addresses, and so on.

If you want to use the existing IP address and name of an interface that you want to delete, then you need to reconfigure the old interface with a dummy name and IP address so that you can use those settings on the new interface.

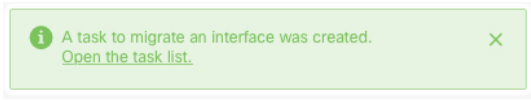
f) To replace an old interface with a new interface, click the Replace icon for the old interface.

This process replaces the old interface with the new interface in all configuration settings that refer to the interface.

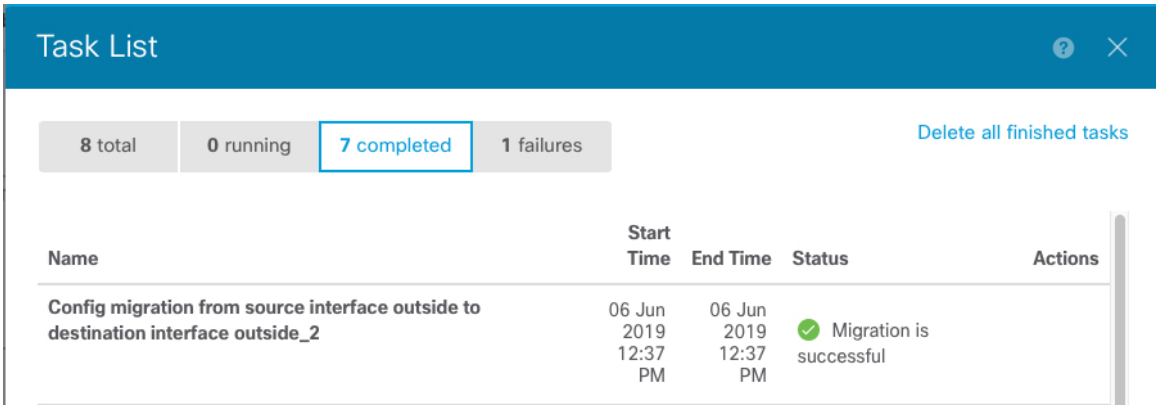
g) Choose the new interface from the **Replacement Interface** drop-down list.



h) A message appears on the **Interfaces** page. Click the link in the message.

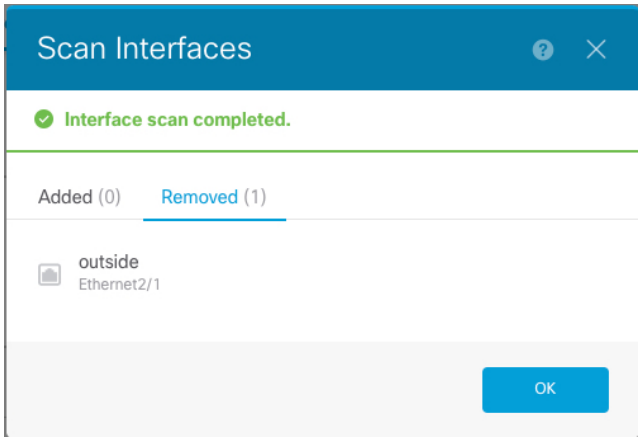


i) Check the **Task List** to ensure that the migration was successful.



Step 2 Sync the interfaces again in the Firewall Management Center or the Firewall Device Manager.

Figure 1: Firewall Device Manager Scan Interfaces



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 using a console connection or a Telnet connection.

connect module *slot_number* {**console** | **telnet**}

To connect to the security engine of a device that does not support multiple security modules, always use **1** as the *slot_number*.

The benefits of using a Telnet connection is that you can have multiple sessions to the module at the same time, and the connection speed is faster.

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 *name*

To view the instance names, enter the command without a name.

Example:

```
Firepower-module1> connect ftd ftd1
Connecting to ftd(ftd-native) console... enter exit to return to bootCLI
[...]
>
```

Step 3 Exit the application console to the FXOS module CLI.

- Firewall Threat Defense—Enter **exit**

Step 4 Return to the supervisor level of the FXOS CLI.

Exit the console:

a) Enter **~**

You exit to the Telnet application.

b) To exit the Telnet application, enter:

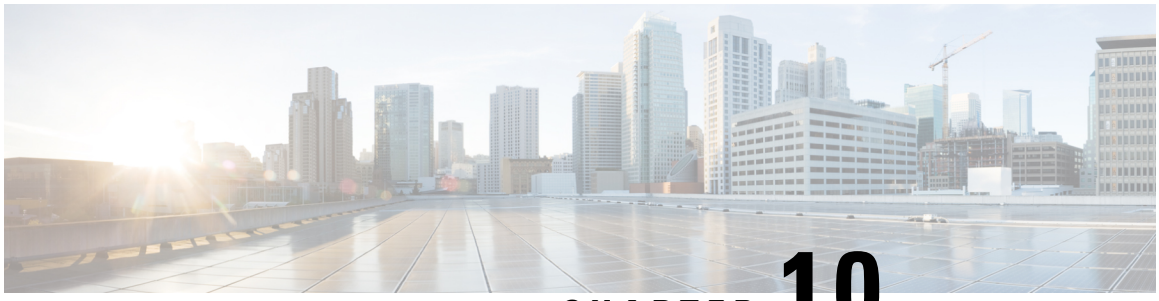
telnet>**quit**

Exit the Telnet session:

a) Enter **Ctrl-], .**

History for Logical Devices

Feature	Version	Details
Support for the Firewall Device Manager on the Firepower 4100/9300.	6.5.0	You can now use the Firewall Device Manager with Firewall Threat Defense logical devices on the Firepower 4100/9300. Firewall Device Manager does not support Multi-Instance capability; only native instances are supported. **Note** Requires FXOS 2.7.1.



CHAPTER 10

High Availability (Failover)

The following topics describe how to configure and manage active/standby failover to accomplish high availability of the Firewall Threat Defense system.

- About High Availability (Failover), on page 191
- System Requirements for High Availability, on page 199
- Guidelines for High Availability, on page 201
- Configuring High Availability, on page 203
- Managing High Availability, on page 215
- Monitoring High Availability, on page 225
- Troubleshooting High Availability (Failover), on page 228

About High Availability (Failover)

A high availability or failover setup joins two devices so that if the primary device fails, the secondary device can take over. This helps you keep your network operational in case of device failure.

Configuring high availability requires two identical Firewall Threat Defense devices connected to each other through a dedicated failover link and, optionally, a state link. The two units constantly communicate over the failover link to determine the operating status of each unit and to synchronize deployed configuration changes. The system uses the state link to pass connection state information to the standby device, so that if a failover occurs, user connections are preserved.

The units form an active/standby pair, 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.

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, the active unit fails over to the standby unit, which then becomes active.

Active/standby failover

Active/standby failover is a high availability configuration that

- lets you use a standby Firewall Threat Defense device to take over the functionality of a failed unit,
- enables the standby unit to become the active unit when the active unit fails, and
- allows the failed unit to come back online as the standby unit if the problem is temporary.

The failed unit must be replaced if the problem is not temporary. Set the standby unit to primary before replacing the failed unit to retain the configuration of the secondary unit.

For multiple context mode, the ASA can fail over the entire unit (including all contexts) but cannot fail over individual contexts separately.

Primary/secondary roles and active/standby status

Primary/secondary roles and active/standby status are failover configuration concepts that

- define unit hierarchy, with primary units taking precedence during simultaneous startup
- determine IP address and MAC address usage patterns during failover operations, and
- establish traffic flow responsibilities between paired units in Active/Standby failover configurations.

Primary and secondary unit differences

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 depending 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

Active unit determination at startup is a high availability process that

- sets a unit to standby if it boots and detects a peer already running as active,
- sets a unit to active if it boots and does not detect a peer, and
- sets the primary unit as active and secondary unit as standby if both units boot simultaneously.

Failover events

A failover event is a system condition that

- triggers an active unit to transfer control to a standby unit in Active/Standby failover configurations
- occurs on a unit basis, and
- follows specific failover policies that determine whether failover occurs and what actions each unit takes.

Failover event policies and actions

This 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 6: 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.

Failover and Stateful Failover Links

The failover link is a dedicated connection between the two units. The stateful failover link is also a dedicated connection, but you can either use the one failover link as a combined failover/state link, or you can create a separate, dedicated state link. If you use just the failover link, the stateful information also goes over that link: you do not lose stateful failover capability.

By default, the communications on the failover and stateful failover links are plain text (unencrypted). You can encrypt the communications for enhanced security by configuring an IPsec encryption key.

The following topics explain these interfaces in more detail, and include recommendations on how to wire the devices for the best results.

Failover Link

The two units in a failover pair constantly communicate over a failover link to determine the operating status of each unit and to synchronize configuration changes.

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.
- System database updates, including VDB and rules, but not including the geolocation and Security Intelligence databases. Each system separately downloads geolocation and Security Intelligence updates. If you create an update schedule, these should remain synchronized. However, if you do a manual geolocation or Security Intelligence update on the active device, you should also do one on the standby device.



Note Eventing, reporting, and audit log data are not synchronized. Event viewer and the dashboards show data related to the given unit only. In addition, deployment history, task history, and other audit log events are not synchronized.

Stateful Failover Link

The system uses the state link to pass connection state information to the standby device. This information helps the standby unit maintain existing connections when a failover occurs.

Using a single link for both the failover and stateful failover links 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.

Interfaces for the Failover and State Links

You can use an unused, but enabled, data interface (physical or EtherChannel) as the failover link; however, you cannot specify an interface that is currently configured with a name. 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). You cannot use a management interface, subinterface, VLAN interface, or switch port for failover.

The Firewall Threat Defense device does not support sharing interfaces between user data and the failover link.

See the following guidelines for sizing the failover and state 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.

When using an EtherChannel interface as the failover or state link, you must confirm that the same EtherChannel with the same ID and member interfaces exists on both devices before establishing high availability. If there is an EtherChannel mismatch, you need to disable HA and corerct the configuration on the secondary unit before. 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 and Stateful Failover Interfaces

You can use any unused data physical interfaces as the failover link and optional dedicated state link. However, you cannot select an interface that is currently configured with a name, or one that has subinterfaces. The failover and stateful failover link interfaces are not configured as normal networking interfaces. They exist for failover communication only, and you cannot use them for through traffic or management access.

Because the configuration is synchronized between the devices, you must select the same port number for each end of a link. For example, GigabitEthernet1/3 on both devices for the failover link.

Connect the failover link, and the dedicated state link if used, 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 Firewall Threat Defense device. A dedicated state link has the same requirement, but must be on a different network segment than the failover link.



Note The advantage of using a switch is that if one of the unit's interfaces goes down, it is easy to troubleshoot which interface failed. If you are using a direct cable connection, if one interface fails, the link is brought down on both peers, which makes it difficult to determine which device is at fault.

- Using an Ethernet cable to connect the units directly, without the need for an external switch. The Firewall Threat Defense supports Auto-MDI/MDIX on its copper Ethernet ports, so you can either use a crossover cable or a straight-through cable. If you use a straight-through cable, the interface automatically detects the cable and swaps one of the transmit/receive pairs to MDIX.

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 Firewall 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 Firewall Threat Defense devices, then when a switch or inter-switch-link is down, both Firewall Threat Defense devices become active. Therefore, the two connection methods shown in the following figures are **not** recommended.

Figure 2: Connecting with a Single Switch—Not Recommended

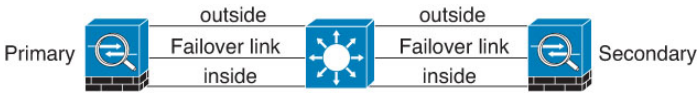
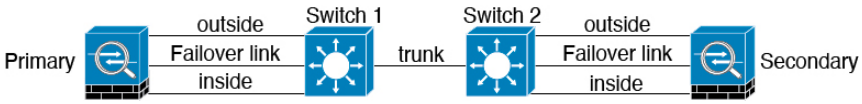


Figure 3: Connecting with a Double-Switch—Not Recommended



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 4: Connecting with a Different Switch

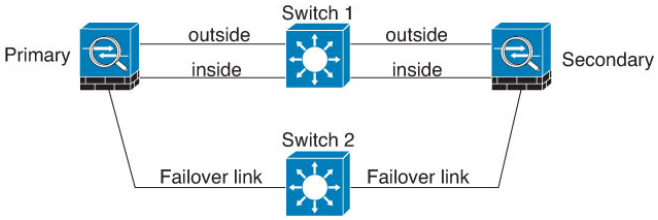
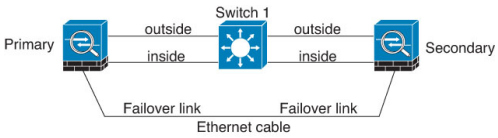


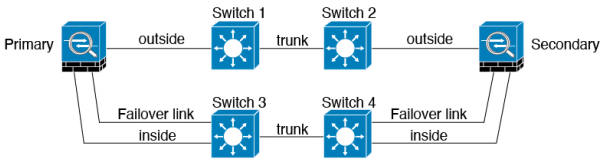
Figure 5: Connecting with a Cable



Scenario 3—Recommended

If the Firewall 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 6: Connecting with a Secure Switch



How Stateful Failover Affects User Connections

The active unit shares connection state information with the standby unit. This means that the standby unit can maintain certain types of connections without impacting the user.

However, there are some types of connections that do not support stateful failover. For these connections, the user will need to reestablish the connection if there is a failover. Often times, this happens automatically based on the behavior of the protocol used in the connection.

The following topics explain which features are supported or not supported for stateful failover.

Supported features

Supported features are high availability capabilities that

- enable state information to be passed to the standby Firewall Threat Defense device during Stateful Failover
- maintain connection states and critical information for seamless traffic flow, and
- ensure minimal disruption during failover events.

State information types

For Stateful Failover, these state information types are passed to the standby 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.
- Passive user identity decisions from the identity policy, but not those gathered through active authentication through captive portal.
- Security Intelligence decisions.
- RA VPN—Remote access VPN end users do not have to reauthenticate or reconnect the VPN session after a failover. However, applications operating over the VPN connection could lose packets during the failover process and not recover from the packet loss.
- From all the connections, only established ones will be replicated on the Standby device.

Unsupported features

Unsupported features are device capabilities that

- are not passed to the standby Firewall Threat Defense in Stateful Failover
- do not maintain state information during failover events, and
- may require re-establishment after a failover occurs.

State information limitations

For Stateful Failover, the following state information is not passed to the standby Firewall Threat Defense:

- 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.

Configuration Changes and Actions Allowed on a Standby Unit

When operating in high-availability mode, you make configuration changes to the active unit only. When you deploy the configuration, the new changes are also transmitted to the standby unit.

However, some properties are unique to the standby unit. You can change the following on a standby unit:

- Management IP address and gateway.
- (CLI only.) The password for the admin user account and other local user accounts. You can make this change in the CLI only, you cannot make it in the Firewall Device Manager. Any local user will have to change their password on both units separately.

In addition, the following actions are available on a standby device.

- High availability actions, such as suspend, resume, reset, and break HA, and switch modes between active and standby.
- Dashboard and eventing data are unique per device, and are not synchronized. This includes custom views in Event Viewer.
- Audit log information is unique per device.
- Smart Licensing registration. However, you must enable or disable the optional licenses on the active unit, and the action is synchronized with the standby unit, which requests or releases the appropriate license.
- Backup, but not restore. You must break HA on the unit to restore a backup. If the backup includes the HA configuration, the unit will rejoin the HA group.
- Software upgrade installation.
- Generating troubleshooting logs.
- Manually updating the Geolocation or Security Intelligence databases. These databases are not synchronized between the units. If you create an update schedule, the units can independently maintain consistency.
- You can view active the Firewall Device Manager user sessions, and delete sessions, from the **Monitoring** > **Sessions** page.

System Requirements for High Availability

The following topics explain the requirements you must meet before incorporating two devices in a high availability configuration.

Hardware Requirements for HA

To link two devices together in a high availability configuration, you must meet the following hardware requirements.

- The devices must be the exact same hardware model.

For the Firepower 9300, High Availability is only supported between same-type modules; but the two chassis can include mixed modules. For example, each chassis has an SM-36 and SM-44. You can create High Availability pairs between the SM-36 modules and between the SM-44 modules.

- The devices must have the same number and type of interfaces.

  For the Firepower 4100/9300 chassis, all interfaces must be preconfigured in FXOS identically before you enable HA. If you change the interfaces after you enable HA, make the interface changes in FXOS on the standby unit, and then make the same changes on the active unit.

- The devices must have the same modules installed. For example, if one has an optional network interface module, then you must install the same module in the other device.
- Intra-chassis High Availability for the Firepower 9300 is not supported. You cannot configure HA between separate logical devices on the same chassis.

Software Requirements for HA

To link two devices together in a high availability configuration, you must meet the following software requirements.

- The devices must run the exact same software version, which means the same major (first), minor (second), and maintenance (third) numbers. You can find the version in the Firewall Device Manager on the **Devices** page, or you can use the **show version** command in the CLI. Devices with different versions are allowed to join, but the configuration is not imported into the standby unit and failover is not functional until you upgrade the units to the same software version.
- Both devices must be in local manager mode, that is, configured using the Firewall Device Manager. If you can log into the Firewall Device Manager on both systems, they are in local manager mode. You can also use the **show managers** command in the CLI to verify.
- You must complete the initial setup wizard for each device.
- Each device must have its own management IP address. The configuration for the management interface is not synchronized between the devices.
- The devices must have the same NTP configuration.
- You cannot configure any interface to obtain its address using DHCP. That is, all interfaces must have static IP addresses.
- For Cloud Services, either both devices must be enrolled in the same region, or neither device can be enrolled. You cannot have mixed Cloud Services enrollment.
- You must deploy any pending changes before you configure high availability.

License Requirements for HA

Before configuring high availability, the units must be in the same state: either both registered with the Base license, or both in evaluation mode. If the devices are registered, they can be registered to different Cisco Smart Software Manager accounts, but the accounts must have the same state for the export-controlled functionality setting, either both enabled or both disabled. However, it does not matter if you have enabled different optional licenses on the units. If you register both units, you must select the same Cisco Cloud Services region for the devices.

If the devices are registered, they must use the same mode, either Smart License or Permanent License Reservation (PLR).

During operation, the units in the high availability pair must have the same licenses. Any license changes you make on the active unit are repeated on the standby unit during deployment.

High availability configurations require two Smart License entitlements; one for each device in the pair. You must ensure there are sufficient licenses in your account to apply to each device. It is possible to be in compliance on one device, but out of compliance on the other, if there are insufficient licenses.

For example, if the active device has the Base license and the Threat, and the standby device has only the Base license, the standby unit communicates with the Cisco Smart Software Manager to obtain an available Threat from your account. If your Smart Licenses account does not include enough purchased entitlements, your account becomes Out-of-Compliance (and the standby device is Out-of-Compliance even though the active device is compliant) until you purchase the correct number of licenses.

Watch Out For:

- If you register the devices to accounts that have different settings for export controlled features, or try to create an HA pair with one unit registered and the other in evaluation mode, the HA join might fail.
- If you configure an IPsec encryption key with inconsistent settings for export controlled features, both devices will become active after you activate HA. This will impact routing on the supported network segments, and you will have to manually break HA on the secondary unit to recover.
- Do not change licensing in the middle of creating the HA group. Both units must have the same configuration at the time of HA join or you will see the following error: “FDM validation failure - Cloud Service enrollment status mismatch between Primary and Secondary Node. Check app-sync-history CLI for details.”

Guidelines for High Availability

Model Support

- Firepower 9300—You can configure HA on the Firepower 9300. However, you cannot configure HA between separate logical devices on the same Firepower 9300 chassis.
- Firepower 1010:
  - You should not use the switch port functionality when using high availability. Because the switch ports operate in hardware, they continue to pass traffic on both the active *and* the standby units. High availability is designed to prevent traffic from passing through the standby unit, but this feature does not extend to switch ports. In a normal high availability network setup, active switch ports on both units will lead to network loops. We suggest that you use external switches for any switching capability. Note that VLAN interfaces can be monitored by failover, while switch ports cannot. Theoretically, you can put a single switch port on a VLAN and successfully use high availability, but a simpler setup is to use physical firewall interfaces instead.
  - You can only use a firewall interface as the failover link.
  - When the chassis is in a high availability pair, the "Active" LED is amber for the standby unit.

- (Firepower 1000 series, Firepower 2100)—Deploying the devices in HA with hundreds of interfaces configured on them can result in increased delay in the failover time (seconds).
- Firewall Threat Defense Virtual—HA configuration is not supported for Firewall Threat Defense Virtual for the Microsoft Azure Cloud or the Amazon Web Services (AWS) Cloud.

Additional Guidelines

- 169.254.0.0/16 and fd00:0:0:*::/64 are internally used subnets and you cannot use them for the failover or state links.
- The configuration from the active unit is synchronized to the standby unit when you run a deployment job on the active unit. However, some changes do not show up in the pending changes even though they are not synchronized on the standby unit until you deploy changes. If you alter any of the following, the changes are hidden and you must run a deployment job before they are configured on the standby unit. If you need to apply the change immediately, you will need to make some other change that does appear in the pending changes. Hidden changes include edits to the following: schedules for rule, geodatabase, Security Intelligence, or VDB updates; schedules for backups; NTP; HTTP proxy for management connections; license entitlement; cloud services options; URL filtering options.
- You should do backups on both the primary and secondary units. To restore a backup, you must first break HA. Do not restore the same backup on both units, because they would then both go active. Instead, restore the backup on the unit you want to go active first, then restore the equivalent backup on the other unit.
- The **Test** button for the various identity sources works on the active unit only. If you need to test identity source connectivity for the standby device, you must first switch modes to make the standby peer the active peer.
- Creating or breaking the high availability configuration restarts the Snort inspection process on both devices when the configuration change is deployed. This can result in through traffic disruption until the process completely restarts.
- When you initially configure high availability, if the Security Intelligence and Geolocation database versions on the secondary are different than they are on the primary, jobs to update the databases are scheduled on the secondary unit. These jobs are run on the next deployment from the active unit. Even if the HA join fails, these jobs remain and will execute on the next deployment.
- 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 high availability pair can cause communication problems when a failover event occurs. This problem occurs because 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 network management system (NMS). You can instead use SNMPv3 with encryption so the IPsec tunnel is not required.
- Interfaces that you use for high-availability failover and stateful failover links do not have to be enabled. The interface status should show that the link is up, but the interfaces themselves might appear to be disabled. Additionally, the interface information is not updated with the IP addresses defined in the high-availability configuration.

Configuring High Availability

Use a high availability setup to ensure network connectivity even if a device fails. With active/standby high availability, two devices are linked, so that if the active device fails, the standby device takes over and users should see no more than a brief connectivity problem.

The following procedure explains the end-to-end process for setting up an active/standby high availability (HA) pair.

Procedure

Step 1 Prepare the Two Units for High Availability, on page 203.

Step 2 Configure the Primary Unit for High Availability, on page 205.

Step 3 Configure the Secondary Unit for High Availability, on page 208.

Step 4 Configure Failover Criteria for Health Monitoring, on page 209.

The criteria includes peer monitoring and interface monitoring. Although all failover criteria have default settings, you should at least examine them to verify that the default settings work for your network.

- Configure Peer Unit Health Monitoring Failover Criteria, on page 209.
- Configure Interface Health Monitoring Failover Criteria, on page 210.

  For information on interface testing, see How the System Tests Interface Health, on page 212.

Step 5 (Optional but recommended.) Configure Standby IP and MAC Addresses, on page 213.

Step 6 (Optional.) Verify the High Availability Configuration, on page 214.

Prepare the Two Units for High Availability

There are many things that you must prepare correctly before you can successfully configure high availability.

Procedure

Step 1 Ensure that the devices meet the requirements explained in Hardware Requirements for HA, on page 199.

Step 2 Determine whether you will use a single failover link, or separate failover and stateful failover links, and identify the ports you will use.

You must use the same port number on each device for each link. For example, GigabitEthernet 1/3 on both devices for the failover link. Know which ones you will use so that you do not accidentally use them for other purposes. For more information, see Failover and Stateful Failover Links, on page 193.

Step 3 Install the devices, connect them to the network, and complete the initial setup wizard on each device.

a) Review the recommended network designs in Avoiding Interrupted Failover and Data Links, on page 195.

b) Connect at least the outside interfaces, as explained in Connect the Interfaces, on page 19.

You can also connect the other interfaces, but you must ensure that you use the same port on each device to connect to a given subnet. Because the devices will share the same configuration, you must connect them to your networks in a parallel manner.

Note
The setup wizard does not let you change the IP addresses on the management and inside interface. Thus, if you connect either of these interfaces on the primary device to the network, do not also connect the interfaces on the secondary device, or you will get an IP address conflict. You can directly connect your workstation to one of these interfaces and get an address through DHCP, so that you can connect to the Firewall Device Manager and configure the device.

c) Complete the initial setup wizard on each device. Ensure that you specify static IP addresses for the outside interface. In addition, configure the same NTP servers. For more information, see Complete the Initial Configuration Using the Setup Wizard, on page 21.

Choose the same licensing and Cisco Success Network options for the units. For example, evaluation mode for each or register the devices.

d) On the secondary device, select **Device** > **System Settings** > **Management Interface** and configure a unique IP address, change the gateway if necessary, and disable or change the DHCP server settings to suit your needs.

e) On the secondary device, select **Device** > **Interface** and edit the inside interface. Either delete the IP address, or change it. Also, delete the DHCP server defined for the interface, because you cannot have two DHCP servers on the same network.

f) Deploy the configuration on the secondary device.

g) If necessary based on your network topology, log into the primary device and change the management address, gateway, and DHCP server settings, and the inside interface IP address and DHCP server settings. Deploy the configuration if you make any changes.

h) If you have not connected the inside interface, or management interface if you use a separate management network, you can now connect them to the switches.

Step 4 Verify that the devices have the exact same software version, which means the same major (first), minor (second), and maintenance (third) numbers. You can find the version in the Firewall Device Manager on the Devices page, or you can use the **show version** command in the CLI.

If they are not running the same software versions, obtain the preferred software version from Cisco.com and install it on each device. For details, see Upgrading Firewall Threat Defense, on page 791.

Step 5 Connect and configure the failover and stateful failover links.

a) Following your preferred network design (chosen from Avoiding Interrupted Failover and Data Links, on page 195), connect the failover interfaces for each device appropriately, either to a switch or directly to each other.

b) If you are using a separate state link, also connect the stateful failover interfaces for each device appropriately.
c) Log into each device in turn and go to **Device** > **Interface**. Edit each interface and verify there are no interface names or IP addresses configured.

   If the interfaces are configured with names, you might need to remove them from security zones and delete other configurations before you can delete the name. If deleting the name fails, examine the error messages to determine what other changes you need to make.

Step 6 On the primary device, connect the remaining data interfaces and configure the device.
a) Select **Device** > **Interface**, edit each interface used for through traffic and configure the primary static IP addresses.
b) Add the interfaces to security zones, and configure the basic policies needed to handle traffic on the connected networks. For example configurations, see the topics listed in Best Practices: Use Cases for Firewall Threat Defense, on page 39.
c) Deploy the configuration.

Step 7 Verify that you meet all the requirements explained in Software Requirements for HA, on page 200.

Step 8 Verify that you have consistent licensing (registered or in evaluation mode). For more information, see License Requirements for HA, on page 200.

Step 9 On the secondary device, connect the remaining data interfaces to the same networks as the equivalent interfaces on the primary device. Do not configure the interfaces.

Step 10 On each device, select **Device** > **System Settings** > **Cloud Services** and verify that you have the same settings.

You are now ready to configure high availability on the primary device.

Configure the Primary Unit for High Availability

To set up an active/standby high availability pair, you must first configure the primary device. The primary device is the unit that you intend should be active under normal circumstances. The secondary device remains in standby mode until the primary unit becomes unavailable.

Select which device you want to be primary, then log into the Firewall Device Manager on that device and follow this procedure.

Note Once you establish the high availability pair, you must break the pair in order to edit the configuration described in this procedure.

Before you begin

Ensure that the interfaces you will configure for the failover and stateful failover link are not named. If they currently are named, you must remove the interfaces from any policies that use them, including security zone objects, then edit the interfaces to delete the name. The interfaces must also be in routed mode, not passive mode. These interfaces must be dedicated for use in the HA configuration: you cannot use them for any other purposes.

If there are any pending changes, you must deploy them before you can configure HA.

Procedure

Step 1 Click **Device**.

Step 2 On the right side of the device summary, click **Configure** next to the **High Availability** group.

If you are configuring HA for the first time on the device, the group would look like the following.



Step 3 On the High Availability page, click the **Primary Device** box.

If the secondary device is already configured, and you copied the configuration to the clipboard, you can click the **Paste from Clipboard** button and paste in the configuration. This will update the fields with the appropriate values, which you can then verify.

Step 4 Configure the **Failover Link** properties.

The two units in a failover pair constantly communicate over a failover link to determine the operating status of each unit and to synchronize configuration changes. For more information, see Failover Link, on page 194.

- **Physical Interface**—Select the interface you connected to the secondary device for use as the failover link. This must be an unnamed interface.

  When using an EtherChannel interface as the failover or state link, you must confirm that the same EtherChannel with the same ID and member interfaces exists on both devices before establishing high availability. If there is an EtherChannel mismatch, you need to disable HA and corerct the configuration on the secondary unit before. 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.

- **Type**—Select whether you will use an IPv4 or IPv6 address for the interface. You can configure one type of address only.

- **Primary IP**—Enter the IP address for the interface on this device. For example, 192.168.10.1. For IPv6 addresses, you must include the prefix length in standard notation, for example, 2001:a0a:b00::a0a:b70/64.

- **Secondary IP**—Enter the IP address that should be configured on the other end of the link for the interface on the secondary device. The address must be on the same subnet as the primary address, and it must be different than the primary address. For example, 192.168.10.2 or 2001:a0a:b00::a0a:b71/64.

- **Netmask** (IPv4 only)—Enter the subnet mask for the primary/secondary IP address.

Step 5 Configure the Stateful Failover Link properties.

The system uses the state link to pass connection state information to the standby device. This information helps the standby unit maintain existing connections when a failover occurs. You can either use the same link as the failover link, or configure a separate link.

- **Use the Same Interface as the Failover Link**—Select this option if you want to use a single link for the failover and stateful failover communications. If you select this option, continue with the next step.

- **Physical Interface**—If you want to use a separate stateful failover link, select the interface you connected to the secondary device for use as the stateful failover link. This must be an unnamed interface. Then, configure the following properties:
  - **Type**—Select whether you will use an IPv4 or IPv6 address for the interface. You can configure one type of address only.
  - **Primary IP**—Enter the IP address for the interface on this device. The address must be on a different subnet than the one used for the failover link. For example, 192.168.11.1. For IPv6 addresses, you must include the prefix length in standard notation, for example, 2001:a0a:b00:a::a0a:b70/64.
  - **Secondary IP**—Enter the IP address that should be configured on the other end of the link for the interface on the secondary device. The address must be on the same subnet as the primary address, and it must be different than the primary address. For example, 192.168.11.2 or 2001:a0a:b00:a::a0a:b71/64.
  - **Netmask** (IPv4 only)—Enter the subnet mask for the primary/secondary IP address.

Step 6 (Optional.) Enter an **IPsec Encryption Key** string if you want to encrypt communication between the two units in the pair.

You must configure the exact same key on the secondary node, so make a note of the string you enter.

If you do not enter a key, all communication on the failover and stateful failover links is in plain text. If you are not using direct cable connections between the interfaces, this could be a security problem.

Note
If you configure HA failover encryption in evaluation mode, the systems use DES for the encryption. If you then register the devices using an export-compliant account, the devices will use AES after a reboot. Thus, if a system reboots for any reason, including after installing an upgrade, the peers will be unable to communicate and both units will become the active unit. We recommend that you do not configure encryption until after you register the devices. If you do configure this in evaluation mode, we recommend you remove the encryption before registering the devices.

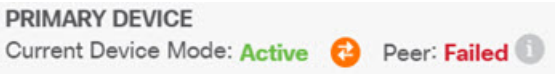
Step 7 Click **Activate HA**.

The system immediately deploys the configuration to the device. You do not need to start a deployment job. If you do not see a message saying that your configuration was saved and deployment is in progress, scroll to the top of the page to see the error messages.

The configuration is also copied to the clipboard. You can use the copy to quickly configure the secondary unit. For added security, the encryption key is not included in the clipboard copy.

After configuration completes, you get a message explaining the next steps you need to take. Click **Got It** after reading the information.

At this point, you should be on the High Availability page, and your device status should be “Negotiating.” The status should transition to Active even before you configure the peer, which should appear as Failed until you configure it.



You can now configure the secondary unit. See Configure the Secondary Unit for High Availability, on page 208.

Note
The selected interfaces are not configured directly. However, if you enter **show interface** in the CLI, you will see that the interfaces are using the specified IP addresses. The interfaces are named "failover-link" and if you configure a separate state link, "stateful-failover-link."

Configure the Secondary Unit for High Availability

After you configure the primary device for active/standby high availability, you must then configure the secondary device. Log into the Firewall Device Manager on that device and follow this procedure.

Note If you have not done so already, copy the high availability configuration from the primary device to the clipboard. It is much easier to configure the secondary device using copy/paste than to manually enter the data.

Procedure

Step 1 Click **Device**.

Step 2 On the right side of the device summary, click **Configure** next to the **High Availability** group.

If you are configuring HA for the first time on the device, the group would look like the following.



Step 3 On the High Availability page, click the **Secondary Device** box.

Step 4 Do one of the following:

- **Easy method**—Click the **Paste from Clipboard** button, paste in the configuration and click **OK**. This will update the fields with the appropriate values, which you can then verify.
- **Manual method**—Configure the failover and stateful failover links directly. Enter the exact same settings on the secondary device that you entered on the primary device.

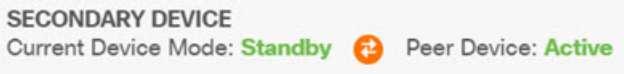
Step 5 If you configured an **IPSec Encryption Key** on the primary device, enter the exact same key for the secondary device.

Step 6 Click **Activate HA**.

The system immediately deploys the configuration to the device. You do not need to start a deployment job. If you do not see a message saying that your configuration was saved and deployment is in progress, scroll to the top of the page to see the error messages.

After configuration completes, you get a message saying that you have configured HA. Click **Got It** to dismiss the message.

At this point, you should be on the High Availability page, and your device status should indicate that this is the secondary device. If the join with the primary device was successful, the device will synchronize with the primary, and eventually the mode should be Standby and the peer should be Active.



Note
The selected interfaces are not configured directly. However, if you enter **show interface** in the CLI, you will see that the interfaces are using the specified IP addresses. The interfaces are named "failover-link" and if you configure a separate state link, "stateful-failover-link."

Configure Failover Criteria for Health Monitoring

The units in a high availability configuration monitor themselves for overall health and for interface health.

The failover criteria define the health monitoring metrics that determine whether a peer has failed. If the active peer is the unit that violates the criteria, then it triggers a failover to the standby unit. If the standby peer is the unit that violates the criteria, it is marked as failed and is not available for failover.

You can configure failover criteria on the active device only.

The following table shows the failover triggering events and associated failure detection timing.

Table 7: Failover Times Based on Failover Criteria

Failover Triggering Event	Minimum	Default	Maximum
The active unit loses power or stops normal operation.	800 milliseconds	15 seconds	45 seconds
An active unit interface physical link is down.	500 milliseconds	5 seconds	15 seconds
An active unit interface is up, but a connection problem causes interface testing.	5 seconds	25 seconds	75 seconds

The following topics explain how to customize the failover health monitoring criteria and also how the system tests interfaces.

Configure Peer Unit Health Monitoring Failover Criteria

Each peer in a high availability configuration 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 the peer is responsive. The action that the device takes depends on the response from the other unit:

- If the device receives a response on the failover link, then it does not fail over.
- If the 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 fail over. 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 device does not receive a response on any interface, then the standby unit switches to active mode and classifies the other unit as failed.

You can configure the poll and hold time for the hello messages.

Procedure

Step 1 On the active device, click **Device**.

Step 2 Click the **High Availability** link on the right side of the device summary.

The Failover Criteria are listed in the right column of the High Availability page.

Step 3 Define the **Peer Timing Configuration**.

These settings determine how quickly the active device can fail over to the standby device. With a faster poll time, the device can detect failure and trigger failover faster. However, faster detection can cause unnecessary switchovers when the network is temporarily congested. The default settings are appropriate for most situations.

If a unit does not hear hello packet on the failover interface for one polling period, additional testing occurs through the remaining interfaces. If there is still no response from the peer unit during the hold time, the unit is considered failed and, if the failed unit is the active unit, the standby unit takes over as the active unit.

- **Poll Time**—The amount of time between hello messages. Enter 1 - 15 seconds, or 200 - 999 milliseconds. The default is 1 second.
- **Hold Time**—The time during which a unit must receive a hello message on the failover link, after which the peer unit is declared failed. The hold time must be at least 3 times more than the poll time. Enter 1 - 45 seconds, or 800 - 999 milliseconds. The default is 15 seconds.

Step 4 Click **Save**.

Configure Interface Health Monitoring Failover Criteria

You can monitor up to 211 interfaces, depending on your device model. You should monitor important interfaces. For example, interfaces that ensure throughput between important networks. Monitor an interface only if you configure standby IP addresses for it, and if the interface should be always up.

When a unit does not receive hello messages on a monitored interface for 2 polling periods, it runs interface tests. If all interface tests fail for an interface, but this same interface on the other unit continues to successfully pass traffic, then the interface is considered to be failed. If the threshold for failed interfaces is met, then a failover occurs. If the other unit's interface also fails all the network tests, then both interfaces go into the “Unknown” state and do not count towards the failover limit.

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.

You can monitor interface HA status from the CLI or CLI Console using the **show monitor-interface** command. For more information, see Monitoring Status for HA-Monitored Interfaces, on page 227.



Note When an interface goes down, for failover it is still considered to be a unit issue. If the unit detects that an interface is down, failover occurs immediately (if you keep the default threshold of 1 interface), without waiting for the interface holdtime. The interface holdtime is only useful when the unit considers its status to be OK, although it is not receiving hello packets from the peer.

Before you begin

By default, all named physical interfaces are selected for HA monitoring. Thus, you should disable monitoring on unimportant physical interfaces. For subinterfaces or bridge groups, you must manually enable monitoring.

To disable interface monitoring completely and prevent failover due to interface failure, simply ensure that no interface is enabled for HA monitoring.

Procedure

Step 1 On the active device, click **Device**.

Step 2 Click the **High Availability** link on the right side of the device summary.

The Failover Criteria are listed in the right column of the High Availability page.

Step 3 Define the **Interface Failure Threshold**.

If the number of failed interfaces meets the threshold, the unit marks itself as failed. If the unit is the active unit, it fails over to the standby unit. If the unit is the standby unit, by marking itself as failed, the active unit will not consider the unit as available for failover.

When setting this criteria, consider how many interfaces you are monitoring. For example, if you enable monitoring on only 2 interfaces, then a threshold of 10 interfaces will never be reached. You configure monitoring for an interface by selecting the **Enable for HA Monitoring** option on the **Advanced Options** tab when editing interface properties.

By default, the unit marks itself as failed if one monitored interface fails.

You can set the interface failure threshold by selecting one of the following **Failover Criteria** options:

- **Number of failed interfaces exceeds**—Enter the raw number of interfaces. The default is 1. The maximum actually depends on the device model and can vary, but you cannot enter more than 211. If you use this criteria, you will get a deployment error if you enter a number larger than the device supports. Try a smaller number or use percentage instead.
- **Percentage of failed interfaces exceeds**—Enter a number from 1 - 100. For example, if you enter 50%, and you are monitoring 10 interfaces, then the device marks itself as failed if 5 interfaces fail.

Step 4 Define the **Interface Timing Configuration**.

These settings determine how quickly the active device can determine if an interface has failed. With a faster poll time, the device can detect interface failure faster. However, faster detection can mean that busy interfaces get marked as failed when in fact they are healthy, which can result in unnecessarily frequent failovers. The default settings are appropriate for most situations.

If an interface link is down, interface testing is not conducted and the standby unit could become active in just one interface polling period if the number of failed interfaces meets or exceeds the configured interface failover threshold.

- **Poll Time**—The frequency that hello packets are sent out on data interfaces. Enter 1 - 15 seconds, or 500 - 999 milliseconds. The default is 5 seconds.
- **Hold Time**—The hold time determines how long it takes from the time a hello packet is missed to when the interface is marked as failed. Enter 5 - 75 seconds. You cannot enter a hold time that is less than 5 times the poll time.

Step 5 Click **Save**.

Step 6 Enable HA monitoring for each interface you want to monitor.

a) Choose **Device** > **Interfaces**.

If an interface is being monitored, the Monitor for HA column indicates Enabled.

b) Click the edit icon ( ) for an interface whose monitoring status you want to change.

You cannot edit the failover or stateful failover interfaces. Interface monitoring does not apply to them.

c) Click the **Advanced Options** tab.

d) Select or deselect the **Enable for HA Monitoring** checkbox as preferred.

e) Click **OK**.

Step 7 (Optional, but recommended.) Configure standby IP addresses and MAC addresses for monitored interfaces. See Configure Standby IP and MAC Addresses, on page 213.

How the System Tests Interface Health

The system continuously tests interfaces that you are monitoring for high availability health. The address used for testing an interface is based on the address types you configure:

- If an interface has both 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).

The system performs the following tests on each unit:

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 unit considers it failed. If the status is Up, then the unit performs the Network Activity test.
2. Network Activity test—A received network activity test. The purpose of this test is to generate network traffic using LANTEST messages to determine which (if either) unit has failed. At the start of the test, each unit clears its received packet count for its interfaces. As soon as a unit receives any packets during the test (up to 5 seconds), then the interface is considered operational. If one unit receives traffic and the other unit does not, then the unit that received no traffic is considered failed. If neither unit received traffic, then the unit starts the ARP test.

3. ARP test—A reading of the unit ARP cache for the 2 most recently acquired entries. One at a time, the unit sends ARP requests to these devices, attempting to stimulate network traffic. After each request, the unit counts all received traffic for up to 5 seconds. If traffic is received, the interface is considered operational. If no traffic is received, an ARP request is sent to the next device. If at the end of the list no traffic has been received, the unit starts the ping test.
4. Broadcast Ping test—A ping test that consists of sending out a broadcast ping request. The unit then counts all received packets for up to 5 seconds. If any packets are received at any time during this interval, the interface is considered operational and testing stops. If no traffic is received, the testing starts over again with the ARP test.

Configure Standby IP and MAC Addresses

When you configure your interfaces, you can specify an active IP address and a standby IP address on the same network. 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.

1. When the primary unit fails over, the secondary unit assumes the IP addresses and MAC addresses of the primary unit and begins passing traffic.
2. The unit that is now in standby state takes over the standby 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.

If the secondary unit boots without detecting the primary unit, the secondary unit becomes the active unit and uses its own MAC addresses, because it does not know the primary unit MAC addresses. However, 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.

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. You can manually configure virtual MAC addresses.

If you do not configure virtual MAC addresses, you might need to clear the ARP tables on connected routers to restore traffic flow. The 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.

Procedure

Step 1 Choose **Device** > **Interfaces**.

You should at least configure standby IP and MAC addresses for the interfaces you are monitoring for HA. If an interface is being monitored, the Monitor for HA column indicates Enabled.

Step 2 Click the edit icon ( ) for the interface whose standby addresses you want to configure.

You cannot edit the failover or stateful failover interfaces. You set the IP addresses for these interfaces when you configure high availability.

Step 3 Configure the Standby IP addresses on the **IPv4 Address** and **IPv6 Address** tabs.

The standby address is used by this interface on the standby device. 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. Configure standby addresses for each IP version you are using.

Step 4 Click the **Advance Options** tab and configure the MAC Addresses.

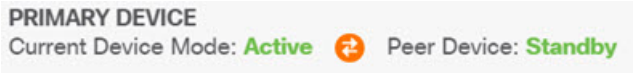
By default, the system uses the MAC address burned into the network interface card (NIC) for the interface. Thus, all subinterfaces on an interface use the same MAC address, so you might want to create unique addresses per subinterface. Manually configured active/standby MAC addresses are also recommended if you configure high availability. Defining the MAC addresses helps maintain consistency in the network in the event of failover.

- **MAC Address**—The Media Access Control address in H.H.H format, where H is a 16-bit hexadecimal digit. For example, you would enter the MAC address 00-0C-F1-42-4C-DE 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.)
- **Standby 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 5 Click **OK**.

Verify the High Availability Configuration

After completing the high availability configuration, verify that the device status indicates that both devices are operational and in active/standby mode.



You can verify that the high availability configuration is working by following this procedure.

Procedure

Step 1 Test that your active unit is passing traffic as expected by using FTP (for example) to send a file between hosts on different interfaces.

At least test connections from one workstation to systems that are connected to each of the configured interfaces.

Step 2 Switch modes so that the active unit is now the standby unit by doing one of the following:

- In the Firewall Device Manager, select **Switch Mode** from the gear menu on the **Device** > **High Availability** page.
- In the CLI of the active unit, enter **no failover active**.

Step 3 Repeat the connection testing to verify that you can make the same connections through the other unit in the high availability pair.

If the test is not successful, verify that you connected the unit's interfaces to the same networks as the equivalent interfaces on the other unit.

You can see the HA status from the High Availability page. You can also use the CLI or CLI Console of the unit, and enter the **show failover** command to check the failover status. Also, use the **show interface** command to verify the interface configuration for the interfaces used in any connection tests that failed.

If these actions do not identify the problem, there are other steps you can take. See Troubleshooting High Availability (Failover), on page 228.

Step 4 When you are finished, you can switch modes to return active status to the original unit that was active.

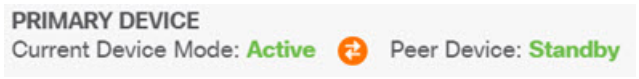
Managing High Availability

You can manage a high availability pair by clicking the **High Availability** link on the **Device** Summary page.



The High Availability page includes the following:

- **Role and Mode Status**—The left status area shows whether the device is the Primary or Secondary device in the group. The mode indicates whether this device is active or standby, or whether HA has been suspended or the device is waiting to join the peer device. It also shows the status of the peer device, which can be active, standby, suspended, or failed. For example, when you are logged into the primary device and it is also the active device, and the secondary device is healthy and ready for failover if needed, the status would look like the following. You can click the icon between the peers to get information on the configuration synchronization status between the devices.

- **Last Failure Reason**—If the High Availability (HA) configuration fails for some reason, such as the active device becoming unavailable and failing over to the standby device, the last reason for failure is shown below the status information for the role and mode status. This message is derived from the failover history.
- **Failover History** link—Click this link to see the detailed history of status of the devices in the pair. The system opens the CLI Console and executes the **show failover history details** command.
- **Deployment History** link—Click this link to go to the audit log with the events filtered to show deployment jobs only.
- **Gear button** —Click this button to perform actions on the devices.
  - **Suspend HA/Resume HA**—Suspending HA stops the devices from functioning as a high availability pair without removing the HA configuration. You can subsequently resume, that is, re-enable, HA on the devices. For details, see Suspending or Resuming High Availability, on page 216.
  - **Break HA**—Breaking HA removes the high availability configuration from both devices and returns them to standalone devices. For details, see Breaking High Availability, on page 217.

- **Switch Mode**—Switching mode lets you force an active device to become standby, or a standby device to become active, depending from which device you perform the action. For details, see Switching the Active and Standby Peers (Forcing Failover), on page 219.

- **High Availability Configuration**—This panel shows the configuration of the failover pair. Click the **Copy to Clipboard** button to load the information into the clipboard, from where you can paste it into the secondary device's configuration. You can also copy it into another file for your records. This information does not show whether you defined an IPsec encryption key.



Note The interface configuration for HA is not reflected on the Interfaces page (**Device** > **Interfaces**). You cannot edit the interfaces that you use in an HA configuration.

- **Failover Criteria**—This panel includes the settings that determine the health criteria used when evaluating whether the active unit has failed and the standby unit should become the active unit. Adjust these criteria so that you get the failover performance required in your network. For details, see Configure Failover Criteria for Health Monitoring, on page 209.

The following topics explain various management tasks related to a high availability configuration.

Suspending or Resuming 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.
- You want to prevent failover while installing a software upgrade on the standby device.

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.

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.

You can resume a unit only if it is in Suspended state. The unit will negotiate active/standby status with the peer unit.

Note If necessary, you can suspend HA from the CLI by entering the **configure high-availability suspend** command. To resume HA, enter **configure high-availability resume**.

Before you begin

If you suspend high availability through the Firewall Device Manager, it stays suspended until you resume it, even if you reload the unit. However, if you suspend it through the CLI, it is a temporary state, and upon reload, the unit resumes the high-availability configuration automatically and negotiates the active/standby state with the peer.

If you are suspending high availability on the standby unit, please check whether the active unit is currently running a deployment job. If you switch modes while a deployment job is in progress, the job will fail and you will lose your configuration changes.

Procedure

Step 1 Click **Device**.

Step 2 Click the **High Availability** link on the right side of the device summary.

Step 3 Choose the appropriate command from the gear icon ( ).

- **Suspend HA**—You are prompted to confirm the action. Read the message and click **OK**. The HA status should show that the device is in Suspended mode.
- **Resume HA**—You are prompted to confirm the action. Read the message and click **OK**. The HA status should return to normal, either active or standby, after the unit negotiates with the peer.

Breaking High Availability

If you no longer want the two devices to operate as a high availability pair, you can break the HA configuration. When you break HA, each device becomes a standalone device. Their configurations are changed as follows:

- The active device retains the full configuration as it is prior to the break, with the HA configuration removed.
- The standby device has all interface configuration removed in addition to the HA configuration. All physical interfaces are disabled, although subinterfaces are not disabled. The management interface remains active, so you can log into the device and reconfigure it.

Note Alternatively, you can use the BreakHAStatus API resource (from the API Explorer) and use the **interfaceOption** attribute to direct the system to reconfigure the standby device's interfaces using the standby IP addresses. You must use the API if you want this result; Firewall Device Manager always disables the interfaces. Note that the system reconfigures the IP addresses but otherwise does not reconfigure all interface options, so traffic might not behave as expected until you deploy changes after the break.

How the break actually affects the units depends on the state of each unit when you perform the break.

- If the units are in a healthy active/standby state, break HA from the active unit. This will remove the HA configuration from both devices in the HA pair. If you want to break HA on the standby unit only, you must log into it and first suspend HA, then you can break HA.
- If the standby unit is in a suspended or failed state, breaking HA from the active unit removes the HA configuration from the active unit only. You must log into the standby unit and also break HA on that unit.
- If the peers are still negotiating HA or are synchronizing their configuration, you cannot break HA. Wait for the negotiation or synchronization to complete or time out. If you believe the systems are stuck in this state, you can suspend HA and then break HA.

Note When using the Firewall Device Manager, you cannot break HA from the CLI using the **configure high-availability disable** command.

Before you begin

For ideal results, bring the devices into a healthy active/standby state, and perform this action from the active device.

Procedure

Step 1 Click **Device**.

Step 2 Click the **High Availability** link on the right side of the device summary.

Step 3 From the gear icon ( ), choose **Break HA**.

Step 4 Read the confirmation message, decide whether to select the option to disable interfaces, and click **OK**.

You must select the option to disable interfaces if you are breaking HA from the standby unit.

The system immediately deploys your changes on both this device and the peer device (if possible). It might take a few minutes for deployment to complete on each device and for each device to become independent.

Switching the Active and Standby Peers (Forcing Failover)

You can switch the active/standby modes for a functioning high-availability pair, that is, one peer is active, the other is standby. For example, if you are installing a software upgrade, you can switch the active unit to standby so that the upgrade does not impact user traffic.

You can switch modes from either the active or standby unit, but the peer unit must be functioning from the other unit's point of view. You cannot switch modes if any unit is suspended (you must resume HA first) or failed.



Note If necessary, you can switch between active and standby modes from the CLI. From the standby unit, enter the **failover active** command. From the active unit, enter the **no failover active** command.

Before you begin

Before switching modes, verify that the active unit is not in the middle of a deployment job. Wait for the deployment to complete before switching modes.

If the active unit has pending undeployed changes, deploy them before switching modes. Otherwise, you will lose your changes if you run a deployment job from the new active unit.

Procedure

Step 1 Click **Device**.

Step 2 Click the **High Availability** link on the right side of the device summary.

Step 3 From the gear icon (⚙), choose **Switch Mode**.

Step 4 Read the confirmation message and click **OK**.

The system forces failover so that the active unit becomes standby, and the standby unit becomes the new active unit.

Preserving Undeployed Configuration Changes After Failover

When you make configuration changes to the units in a high availability pair, you edit the configuration on the active unit. You then deploy your changes, and both the active and standby unit are updated with the new configuration. It does not matter if the active unit is the primary or secondary device.

However, undeployed changes are not synchronized between the units. Any undeployed changes are available on the unit where you made those changes only.

Thus, if a failover happens when you have undeployed changes, those changes are not available on the new active unit. The changes do, however, remain in place on the unit that is now standby.

To retrieve your undeployed changes, you must switch modes to force a failover and return the other unit to active status. When you log into the newly-active unit, your undeployed changes are available, and you can deploy them. Use the **Switch Modes** command from the **High Availability** settings gear menu (⚙).

Please keep the following in mind:

- If you deploy changes from the active unit while there are undeployed changes on the standby unit, the undeployed changes on the standby unit will be erased. You will not be able to retrieve them.
- When a standby unit joins a high availability pair, any undeployed changes on the standby unit will be erased. The configuration is synchronized whenever a unit joins, or rejoins, the pair.
- If the unit that contains the undeployed changes failed catastrophically, and you had to replace or reimage the unit, your undeployed changes are permanently lost.

Changing Licenses and Registration in High Availability Mode

The units in a high availability pair must have the same licenses and registration status. To make changes:

- You enable or disable optional licenses on the active unit. Then, you deploy the configuration, and the standby unit requests (or frees) the necessary licenses. When enabling licenses, you must ensure that your Cisco Smart Software Manager account has sufficient licenses available, or you could have one unit compliant while the other unit is non-compliant.
- You register, or unregister, the units separately. To function correctly, the units must both be in evaluation mode, or both be registered. You can register the units to different Cisco Smart Software Manager accounts, but the accounts must have the same state for the export-controlled functionality setting, either both enabled or both disabled. You cannot deploy configuration changes if the units have inconsistent registration status.

Editing the HA IPsec Encryption Key or HA Configuration

You can change any of the failover criteria by logging into the active unit, making your changes, and deploying them.

However, if you need to change the IPsec encryption key used on your failover links, or change the interfaces or IP addresses for either the failover or stateful failover links, you must first break the HA configuration. You can then reconfigure the primary and secondary units with the new encryption key or failover/stateful failover link settings.

Marking a Failed Unit Healthy

A unit in a high availability configuration can be marked as failed due to regular health monitoring. If the unit is healthy, it should return to normal status when it meets health monitoring requirements again. If you see a healthy device failing frequently, you might want to increase the peer timeouts, stop monitoring specific interfaces that are less important, or change interface monitoring timeouts.

You can force a failed unit to be seen as healthy by entering the **failover reset** command from the CLI. We recommend that you enter the command from the active unit, which will reset the status of the standby unit. You can display the failover status of the unit with the **show failover** or **show failover state** commands.

Restoring a failed unit to an unfailed state does not automatically make it active. Restored units remain in the standby state until made active by failover (forced or natural).

Resetting the device status does not resolve the problems that led to the device being marked failed. If you do not address the problems, or relax your monitoring timeouts, the device can be marked as failed again.

Upgrading High Availability Firewall Threat Defense

Use this procedure to upgrade high availability devices. Upgrade them one at a time. To minimize disruption, always upgrade the standby. That is, upgrade the current standby, switch roles, then upgrade the new standby. If you need to update FXOS, do that on both chassis before you upgrade Firewall Threat Defense on either. Again, always upgrade the standby.

Caution Do not make or deploy configuration changes on one unit while the other is upgrading, or to a mixed version pair. Even if the system appears inactive, do not manually reboot or shut down during upgrade; you could place the system in an unusable state and require a reimage. You can manually cancel failed or in-progress major and maintenance upgrades, and retry failed upgrades. If you continue to have issues, contact Cisco TAC.

For details on these and other issues you may encounter during upgrade, see Troubleshooting High Availability Threat Defense Upgrades, on page 222.

Before you begin

Complete upgrade planning. Make sure your deployment is healthy and successfully communicating.

Note Upgrade planning starts with reading the Secure Firewall Threat Defense release notes. It then includes taking backups, obtaining upgrade packages, and performing associated upgrades (such as FXOS for the Firepower 4100/9300). It also includes checks for necessary configuration changes, readiness checks, disk space checks, and checks for both running and scheduled tasks. For details, see the Secure Firewall Threat Defense upgrade guides for Firewall Device Manager for your version.

Procedure

Step 1 Log into the standby unit.

Step 2 Select **Device**, then click **View Configuration** in the Updates panel.
The System Upgrade panel shows the currently running software version and any upgrade package that you have already uploaded.

Step 3 Upload the upgrade package.

You can upload one package only. If you upload a new package, it replaces the old one. Make sure you have the correct package for your target version and device model. Click **Browse** or **Replace File** to begin the upload.

When the upload completes, the system displays a confirmation dialog box. Before you click **OK**, optionally select **Run Upgrade Immediately** to choose rollback options and upgrade now. If you upgrade now, it is especially important to have completed as much of the pre-upgrade checklist as possible (see the next step).

Step 4 Perform final pre-upgrade checks, including the readiness check.

Revisit the pre-upgrade checklist. Make sure you have completed all relevant tasks, especially the final checks. If you do not run the readiness check manually, it runs when you initiate the upgrade. If the readiness check fails, the upgrade is canceled. For more information, see Running an Upgrade Readiness Check, on page 793.

Step 5 Click **Upgrade Now** to start the upgrade.

a) Choose rollback options.

You can **Automatically cancel on upgrade failure and roll back to the previous version**. With this option enabled, the device automatically returns to its pre-upgrade state upon major or maintenance upgrade failure. Disable this option if you want to be able to manually cancel or retry a failed upgrade.

b) Click **Continue** to upgrade and reboot the device.

You are automatically logged off and taken to a status page where you can monitor the upgrade until the device reboots. The page also includes an option to cancel the in-progress installation. If you disabled automatic rollback and the upgrade fails, you can manually cancel or retry the upgrade.

Traffic is dropped while you upgrade. For the ISA 3000 only, if you configured hardware bypass for power failure, traffic is dropped during the upgrade but is passed without inspection while the device completes its post-upgrade reboot.

Step 6 Log back in when you can and verify upgrade success.

The Device Summary page shows the currently running software version and high availability status. Do not proceed until you have verified success *and* high availability has resumed. If high availability remains suspended after successful upgrade, see Troubleshooting High Availability Threat Defense Upgrades, on page 222.

Step 7 Upgrade the second unit.

a) Switch roles, making this device active: Select **Device** > **High Availability**, then select **Switch Mode** from the gear menu ( ). Wait for the unit's status to change to active and confirm that traffic is flowing normally. Log out.

b) Upgrade: Repeat the previous steps to log into the new standby, upload the package, upgrade the device, monitor progress, and verify success.

Step 8 Examine device roles.

If you have preferred roles for specific devices, make those changes now.

Step 9 Log into the active unit.

Step 10 Complete post-upgrade tasks.

a) Update system databases. If you do not have automatic updates configured for intrusion rules, VDB, and GeoDB, update them now.

b) Complete any other required post-upgrade configuration changes.

c) Deploy.

Troubleshooting High Availability Threat Defense Upgrades

General Upgrade Troubleshooting

These issues can occur when you are upgrading any device, whether standalone or in a high availability pair.

Upgrade package errors.

To find the correct upgrade package, select or search for your model on the Cisco Support & Download site, then browse to the software download page for the appropriate version. Available upgrade packages are listed along with installation packages, hotfixes, and other applicable downloads. Upgrade package file names reflect the platform, package type (upgrade, patch, hotfix), software version, and build.

Upgrade packages from Version 6.2.1+ are signed, and terminate in .sh.REL.tar. If you are upgrading from an older version, download the package that terminates in .sh instead. The Cisco Support & Download site indicates the correct package for your version. Do not untar signed upgrade packages. Do not rename upgrade packages or transfer them by email.

Cannot reach the device at all during upgrade.

Devices stop passing traffic during the upgrade or if the upgrade fails. Before you upgrade, make sure traffic from your location does not have to traverse the device itself to access the device's management interface.

Device appears inactive or is unresponsive during upgrade.

You can manually cancel in-progress major and maintenance upgrades; see Canceling or Retrying Firewall Threat Defense Upgrades, on page 794. If the device is unresponsive, or if you cannot cancel the upgrade, contact Cisco TAC.



Caution Even if the system appears inactive, do *not* manually reboot or shut down during upgrade. You could place the system in an unusable state and require a reimage.

Upgrade is successful but the system does not function to your expectations.

First, make sure that cached information gets refreshed. Do not simply refresh the browser window to log back in. Instead, delete any "extra" path from the URL and reconnect to the home page; for example, http://threat-defense.example.com/.

If you continue to have issues and need to return to an earlier major or maintenance release, you may be able to revert; see Reverting Firewall Threat Defense, on page 794. If you cannot revert, you must reimage.

Upgrade fails.

When you initiate a major or maintenance upgrade, you use the **Automatically cancel on upgrade failure...** (auto-cancel) option to choose what happens if upgrade fails, as follows:

- Auto-cancel enabled (default): If upgrade fails, the upgrade cancels and the device automatically reverts to its pre-upgrade state. Correct any issues and try again later.
- Auto-cancel disabled: If upgrade fails, the device remains as it is. Correct the issues and retry immediately, or manually cancel the upgrade and try again later.

For more information, see Canceling or Retrying Firewall Threat Defense Upgrades, on page 794. If you cannot retry or cancel, or if you continue to have issues, contact Cisco TAC.

High Availability Upgrade Troubleshooting

These issues are specific to high availability upgrades.

Upgrade will not begin without deploying uncommitted changes.

If you get an error message stating that you must deploy all uncommitted changes even though there are none, log into the active unit (remember, you should be upgrading the standby), create some minor change, and deploy. Then, undo the change, redeploy, and try the upgrade again on the standby.

If this does not work, and the units are running different software versions against recommendations, switch roles to make the standby unit active, then suspend high availability. Deploy from the active/suspended unit, resume high availability, then switch roles to make the active unit the standby again. Upgrade should then work.

Deployment from active unit fails during standby upgrade, or causes an application synchronization error.

This can happen if you deploy from the active unit while the standby is upgrading, which is not supported. Proceed with the upgrade despite the error. After you upgrade both units, make any required configuration changes and deploy from the active unit. The error should resolve.

To avoid these issues, do not make or deploy configuration changes on one unit while the other is upgrading, or to a mixed version pair.

Configuration changes made during upgrade are lost.

If you absolutely must make and deploy changes to a mixed version pair, you must make the changes to both units or they will be lost after you upgrade the down-level active unit.

High availability is suspended after upgrade.

After the post-upgrade reboot, high availability is briefly suspended while the system performs some final automated tasks, such as updating libraries and restarting Snort. You are most likely to notice this if you log into the CLI *very* shortly after upgrade. If high availability does not resume on its own after the upgrade fully completes and Firewall Device Manager is available, do it manually:

1. Log into both the active device and the standby device and check the task lists. Wait until all tasks are finished running on both devices. If you resume high availability too soon, you may have a future issue where failover causes an outage.
2. Select **Device** > **High Availability**, then select **Resume HA** from the gear menu ( ).

Failover does not occur with a mixed version pair.

Although an advantage of high availability is that you can upgrade your deployment without traffic or inspection interruptions, failover is disabled during the entire upgrade process. That is, not only is failover necessarily disabled when one device is offline (because there is nothing to fail over to—you are essentially already failed over), but failover is also disabled with mixed version pairs. During upgrade is the only time where mixed version pairs are (temporarily) allowed. Schedule upgrades during maintenance windows when they will have the least impact if something goes wrong, and make sure you have enough time to upgrade both devices in that window.

Upgrade failed on only one device, or one device was reverted. The pair is now running mixed versions.

Mixed version pairs are not supported for general operations. Either upgrade the down-version device or revert the higher version device. For patches, because revert is not supported, if you cannot successfully upgrade the down-version device you must break high availability, reimage one or both devices, then re-establish high availability.

Replacing a Unit in a High Availability Pair

If necessary, you can replace a unit in a high availability group without disrupting network traffic.

Procedure

Step 1 If the unit you are replacing is functional, ensure that you fail over to the peer unit, then use the **shutdown** command from the device CLI to bring down the device gracefully. If the unit is not functional, confirm that the peer is operating in Active mode.

If you have Administrator privileges, you can also enter the **shutdown** command through the Firewall Device Manager CLI Console.

Step 2 Remove the unit from the network.

Step 3 Install the replacement unit and reconnect the interfaces.

Step 4 Complete the device setup wizard on the replacement unit.

Step 5 On the peer unit, go to the High Availability page and copy the configuration to the clipboard. Note whether the unit is the Primary or the Secondary unit.

If there are any pending changes, deploy them now and wait for deployment to complete before continuing.

Step 6 On the replacement unit, click **Configure** in the **High Availability** group, then select the opposite unit type from the peer. That is, if the peer is primary, select **Secondary**, if the peer is secondary, select **Primary**.

Step 7 Paste in the HA configuration from the peer, then enter the IPsec key if you use one. Click **Activate HA**.

Once deployment is complete, the unit will contact the peer and join the HA group. The active peer's configuration will be imported, and the replacement unit will be either the primary or secondary unit in the group, based on your selection. You can now verify that HA is operating correctly, and if desired, switch modes so that the new unit is the active unit.

Monitoring High Availability

The following topics explain how you can monitor high availability.

Note that the Event Viewer and dashboards show data related to the device you are logged into only. They do not show merged information for both devices.

Monitoring General Failover Status and History

You can monitor general high availability status and history using the following:

- On the Device Summary (click **Device**), the High Availability group shows unit status.



- On the High Availability page (click **Device** > **High Availability**), you can see the status of both units. If any failures have occurred, the last failure reason (from the failover history) is shown. Click the synchronization icon between them for additional status.



- From the High Availability page, click the **Failover History** link next to the status. The system opens the CLI Console and executes the **show failover history details** command. You can also enter this command directly in the CLI or CLI Console.

CLI Commands

From the CLI or CLI console, you can use the following commands:

- **show failover**

  Displays information about the failover state of the unit.

- **show failover history** [**details**]

  Displays the past failover state changes and the reason for the state change. Add the **details** keyword to display failover history from the peer unit. This information helps with troubleshooting.

- **show failover state**

  Displays the failover state of both units. The information includes the primary or secondary status of the unit, the Active/Standby status of the unit, and the last reported reason for failover.

- **show failover statistics**

  Displays the transmit (tx) and receive (rx) packet count of the failover interface. For example, if the output shows that the unit is sending packets, but not receiving any, then you have a problem with the link. This could be a bad wire, wrong IP addresses configured on the peers, or perhaps the units are connecting the failover interfaces to different subnets.

```
> show failover statistics
        tx:320875
        rx:0
```

- **show failover interface**

  Displays the configuration of the failover and stateful failover links. For example:

```
> show failover interface
        interface failover-link GigabitEthernet1/3
                System IP Address: 192.168.10.1 255.255.255.0
                My IP Address    : 192.168.10.1
                Other IP Address : 192.168.10.2
        interface stateful-failover-link GigabitEthernet1/4
                System IP Address: 192.168.11.1 255.255.255.0
                My IP Address    : 192.168.11.1
                Other IP Address : 192.168.11.2
```

- **show monitor-interface**

Displays information about the interfaces monitored for high availability. For details, see Monitoring Status for HA-Monitored Interfaces, on page 227.

- **show running-config failover**

  Displays the failover commands in the running configuration. These are the commands that configure high availability.

Monitoring Status for HA-Monitored Interfaces

If you enabled HA monitoring for any interface, you can check the status of the monitored interfaces in the CLI or CLI Console using the **show monitor-interface** command.

```
> show monitor-interface
 This host: Primary - Active
  Interface inside (192.168.1.13): Normal (Monitored)
  Interface outside (192.168.2.13): Normal (Monitored)
 Other host: Secondary - Standby Ready
  Interface inside (192.168.1.14): Normal (Monitored)
  Interface outside (192.168.2.14): Normal (Monitored)
```

Monitored interfaces can have the following status:

- (Waiting) coupled with any other status, such as Unknown (Waiting)—The interface has not yet received a hello packet from the corresponding interface on the peer unit.
- Unknown—Initial status. This status can also mean the status cannot be determined.
- Normal—The interface is receiving traffic.
- Testing—Hello messages are not heard on the interface for five poll times.
- Link Down—The interface or VLAN is administratively down.
- No Link—The physical link for the interface is down.
- Failed—No traffic is received on the interface, yet traffic is heard on the peer interface.

Monitoring HA-Related Syslog Messages

The system issues a number of syslog messages related to failover at priority level 2, which indicates a critical condition. The ranges of message IDs associated with failover are: 101xxx, 102xxx, 103xxx, 104xxx, 105xxx, 210xxx, 311xxx, 709xxx, 727xxx. For example, 105032 and 105043 indicate a problem with the failover link. For an explanation of the syslog messages, see the *Cisco Threat Defense Syslog Messages* guide at https://www.cisco.com/c/en/us/td/docs/security/firepower/Syslogs/b_fptd_syslog_guide.html.

Note During failover, the system logically shuts down and then brings up interfaces, generating syslog messages 411001 and 411002. This is normal activity.

To be able to see syslog messages, you must configure diagnostic logging on **Device** > **Logging Settings**. Set up an external syslog server so that you can monitor the messages reliably.

Remotely Executing CLI Commands on the Peer Unit

From the CLI, you can enter show commands on the peer device using the failover exec command without needing to log into the peer.

failover exec {**active** | **standby** | **mate**} *command*

You must indicate which unit should execute the command, either the active or the standby, or enter **mate** if you want to ensure the other unit responds instead of the unit that you are logged into.

For example, if you want to see the peer's interface configuration and statistics, you can enter:

```
> failover exec mate show interface
```

You cannot enter **configure** commands. This feature is for use with **show** commands.

Note If you are logged into the active unit, you can reload the standby unit using the **failover reload-standby** command.

You cannot enter the these commands through the Firewall Device Manager CLI Console.

Troubleshooting High Availability (Failover)

If the units in a high availability group are not performing as expected, consider the following steps for troubleshooting the configuration.

If the active unit shows the peer unit as Failed, see Troubleshooting Failed State for a Unit, on page 230.

Procedure

Step 1 From each device (primary and secondary):

- Ping the other device's IP address for the failover link.
- Ping the other device's IP address for the stateful failover link if you use a separate link.

If the ping fails, ensure that the interfaces on each device are connected to the same network segment. If you are using a direct cable connection, check the cable.

Step 2 Make the following general checks:

- Check for duplicate management IP addresses on the primary and secondary.
- Check for duplicate failover and stateful failover IP addresses on the units.
- Check that the equivalent interface port on each device is connected to the same network segment.

Step 3 Check the task list or audit log on the standby device. You should see a successful "Configuration import from Active node" task after every successful deployment on the active device. If the task fails, check the failover link and try deployment again.

Note
If the task list indicates there was a failed deployment task, there might have been a failover during the deployment job. If the standby device was the active unit when you started the deployment task, but failover occurred during the task, the deployment would fail. To resolve the issue, switch modes to make the standby unit the active unit again, then redeploy the configuration changes.

Step 4 Use the **show failover history** command to get detailed information on the state changes on a device.

Some things to look for:

- App Sync failures:

```
12:41:24 UTC Dec 6 2017

App Sync     Disabled     HA state progression failed due to APP SYNC timeout
```

  The Application Synchronization phase is where the configuration from the active device is transported to the standby device. An application synchronization failure puts the device in disabled state, and the device is no longer available to be made active.

  If the device is disabled due to an app sync problem, then you might be using different interfaces on the devices for the endpoints of the failover and stateful failover links. You must be using the same port number for each end of the link.

  If the show failover command shows the secondary device in Pseudo Standby state, this could indicate that you configured different IP addresses for the failover link on the secondary device than what you configured on the primary device. Ensure that you are using the same primary/secondary IP addresses on both devices for the failover link.

  The Pseudo Standby state might also indicate that you configured different IPsec keys on the primary and secondary.

  For additional app sync issues, see Troubleshooting HA App Sync Failures, on page 231.

- Abnormally frequent failovers (going from active to standby and back) might indicate problems with the failover link. In a worst-case scenario, both units might become active, which disrupts through traffic. Ping each end of the link to verify connectivity. You can also use **show arp** to check that the failover IP address and ARP mapping are proper.

  If the failover link is healthy and configured correctly, consider increasing the peer poll and hold time, the interface poll and holdtime, reducing the number of interfaces monitored for HA, or increasing the interface threshold.

- Failures due to interface checks. The Interface Check reason includes a list of the interfaces that were considered to have failed. Check those interfaces to ensure they are configured correctly and there are no hardware issues. Verify there are no issues with the switch configuration on the other end of the links. If there are no problems, consider disabling HA monitoring on those interfaces, Other options are to increase the interface failure threshold or timing.

```
06:17:51 UTC Jan 15 2017

Active    Failed    Interface check
```

```
This Host:3

   admin: inside

   ctx-1: ctx1-1

   ctx-2: ctx2-1

Other Host:0
```

Step 5 If the standby unit cannot be detected, and you cannot find a specific reason such as a bad LAN or cable connection on the failover link, try the following steps.

a) Log into the CLI on the standby unit and enter the **failover reset** command. This command should change a unit in failed state to unfailed state. Now, check the HA status on the active device. If the standby peer is now detected, you are done.

b) Log into the CLI on the active unit and enter the **failover reset** command. This should reset HA status on both the active and standby unit. Ideally, it will reestablish the link between the devices. Check the HA status; if it is not correct yet, continue.

c) Either from the CLI on the active device, or from the Firewall Device Manager, first suspend HA, then resume HA. The CLI commands are **configure high-availability suspend** and **configure high-availability resume**.

d) If these steps fail, **reboot** the standby device.

Troubleshooting Failed State for a Unit

If a unit is marked as Failed in the peer unit's high availability status (on the **Device** or **Device** > **High Availability** page), the following are the general possible reasons based on unit A being the active unit and unit B being the failed peer.

- If the unit B has not yet been configured for high availability (it is still in standalone mode), unit A shows it as Failed.
- If you suspend HA on unit B, then unit A will show it as Failed.
- If you reboot unit B, then unit A will show it as Failed until unit B completes the reboot and resumes communication over the failover link.
- If application synchronization (App Sync) fails on unit B, unit A will show it as Failed. See Troubleshooting HA App Sync Failures, on page 231.
- If unit B fails unit or interface health monitoring, then unit A marks it as failed. Check unit B for systemic problems. Try rebooting the device. If the unit is generally healthy, consider relaxing the unit or interface health monitoring settings. The **show failover history** output should provide information on interface health check failures.
- If both units become active, then each unit will show the peer as Failed. This usually indicates a problem with the failover link.

  It can also indicate a problem with licensing. The devices must have consistent licensing, either both in evaluation mode or both registered. If registered, the Smart License accounts used can be different, but both accounts must have the same selection for export controlled features, either enabled or disabled. If

you configure an IPsec encryption key with inconsistent settings for export controlled features, both devices will become active after you activate HA. This will impact routing on the supported network segments, and you will have to manually break HA on the secondary unit to recover.

Troubleshooting HA App Sync Failures

If the peer unit fails to join the HA group, or fails while you are deploying changes from the active unit, log into the failed unit, go to the **High Availability** page, and click the **Failover History** link. If the **show failover history** output indicates an App Sync failure, then there was a problem during the HA validation phase, where the system checks that the units can functioning correctly as a high availability group.

This type of failure might look like the following:

```
==========================================================================
From State                To State                  Reason
==========================================================================
16:19:34 UTC May 9 2018
Not Detected              Disabled                  No Error

17:08:25 UTC May 9 2018
Disabled                  Negotiation               Set by the config command

17:09:10 UTC May 9 2018
Negotiation               Cold Standby              Detected an Active mate

17:09:11 UTC May 9 2018
Cold Standby              App Sync                  Detected an Active mate

17:13:07 UTC May 9 2018
App Sync                  Disabled                  CD App Sync error is
High Availability State Link Interface Mismatch between Primary and Secondary Node
```

Ideally, you want to see the message “All validation passed” when the From State is App Sync, and the node will reach the Standby Ready state. Any validation failure will transition the peer to the Disabled (Failed) state. You must resolve the problems to make the peers function as a high availability group again. Note that if you fix an App Sync error by making changes to the active unit, you must deploy them and then resume HA for the peer node to join.

The following messages indicate failures, with an explanation of how you can resolve the issues. These errors can happen on node join and on each subsequent deployment. During node join, the system performs a check against the last deployed configuration on the active unit.

- License registration mode mismatch between Primary and Secondary Node.

  The license error indicates that one peer is registered while the other peer is in evaluation mode. The peers must both be registered or both in evaluation mode for them to join an HA group. Because you cannot return a registered device to evaluation mode, you must register the other peer from the **Device** > **Smart License** page.

  If the device you register is the active unit, after registering the device, perform a deployment. Deployment forces the units to refresh and synchronize configurations, which should allow the secondary unit to join the high availability group correctly.

- License export compliance mismatch between Primary and Secondary Node.

The license compliance error indicates that the devices are registered to different Cisco Smart Software Manager accounts, and one account is enabled for export-controlled functionality, whereas the other account is not. The devices must be registered with accounts that have the same setting, enabled or disabled, for export-controlled functionality. Change the device registration on the **Device** > **Smart License** page.

- Software version mismatch between Primary and Secondary Node.

  The software mismatch error indicates that the peers are running different versions of the Firewall Threat Defense software. The system allows a mismatch only temporarily, while you are installing software upgrades one device at a time. However, you cannot deploy configuration changes between upgrading the peers. To resolve this problem, upgrade the peer, then redo the deployment.

- Physical interfaces mismatch between Primary and Secondary Node.

  The standby unit in an HA group must have all of the physical interfaces that exist on the active unit, and these interfaces must have the same hardware names and types (such as GigabitEthernet1/1). This error indicates that the standby unit is missing some interfaces that are present on the active unit. You are allowed to have more interfaces on the standby unit than on the active unit, so either switch which unit is active, or choose another peer unit. However, mismatching interfaces should be temporary state, for example, if you are replacing an interface module on one unit and you need to run it without the module for a short time. For normal operations, both units should have the same number and type of interfaces.

- Failover link interface mismatch between Primary and Secondary Node.

  When you link the failover physical interface to the network on each unit, you must choose the same physical interface. For example, GigabitEthernet1/8 on each unit. This error indicates that you used different interfaces. To resolve the error, correct the cabling on the peer unit.

- Stateful failover link interface mismatch between Primary and Secondary Node.

  If you use a separate stateful failover link, when you link the stateful failover physical interface to the network on each unit, you must choose the same physical interface. For example, GigabitEthernet1/7 on each unit. This error indicates that you used different interfaces. To resolve the error, correct the cabling on the peer unit.

- Failover/Stateful failover link EtherChannel's member interfaces mismatch between Primary and Secondary Node

  If you select an EtherChannel interface for either the failover or stateful failover interfaces, the EtherChannels must have the same ID and member interfaces on each device. This error message indicates whether it is the failover or the stateful failover link that has the mismatch. To resolve the error, correct the configuration of the EtherChannel interfaces so they use the same ID and include the same interfaces on each device.

- Device Model Number mismatch between Primary and Secondary Node.

  For the peers to join an HA group, they must be devices of the exact same model. This error indicates that the peers are not the same device model. You must choose a different peer to configure HA.

- Active and Standby Nodes cannot be on the same chassis.

  You cannot configure high availability using devices that are hosted on the same hardware chassis. When configuring high availability on models that support multiple devices on the same chassis, you must select devices that reside on separate hardware.

- Unknown error occurred, please try again.

Something went wrong during the app sync, but the system could not identify the problem. Try deploying the configuration again.

- Rule package is corrupted. Please update the rule package and try again.

  There is an issue with the intrusion rules database. On the failed peer, go to **Device** > **Updates**, and click **Update Now** in the **Rule** group. Wait for the update to complete, and deploy changes. You can then retry the deployment from the active unit.

- Cloud Service enrollment status mismatch between Primary and Secondary Node.

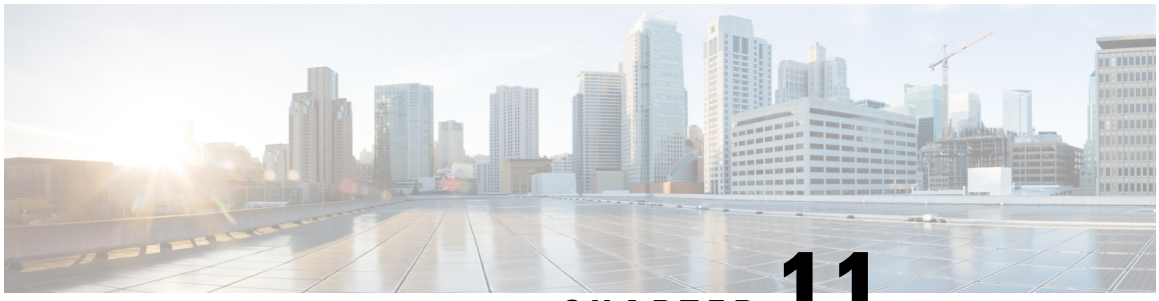
  One of the nodes is enrolled with the Cisco cloud, but the other is not enrolled. Either both nodes must be enrolled or neither can be enrolled to form a high availability group. Go to **Device** > **System Settings** > **Cloud Services** on the each device and ensure both devices are enrolled in the same Cloud Services region.

- Active and Standby Nodes cannot have different cloud regions.

  The devices are registered in different Cisco Cloud Services regions. Determine which region is correct, unregister the other device from Smart Licensing, and select the correct region during re-registration. If both devices have the wrong region, unregister both devices and re-register in the correct region.

- Deployment package is corrupted. Please try again.

  This is a system error. Try the deployment again, which should resolve the problem.



CHAPTER 11

Interfaces

The following topics explain how to configure the interfaces on your Firewall Threat Defense device.

- About Firewall Threat Defense Interfaces, on page 235
- Guidelines and Limitations for Interfaces, on page 238
- Configure a Physical Interface, on page 239
- Configure Bridge Groups, on page 244
- Configure EtherChannels, on page 248
- Configure VLAN Interfaces and Switch Ports, on page 259
- Configure VLAN Subinterfaces and 802.1Q Trunking, on page 270
- Configure Passive Interfaces, on page 275
- Configure Advanced Interface Options, on page 279
- Scan for Interface Changes, and Migrate an Interface, on page 283
- Manage the Network Module for the Secure Firewall 3100, on page 288
- Configure Hardware Bypass for Power Failure (ISA 3000), on page 296
- Monitoring Interfaces, on page 298
- Examples for Interfaces, on page 299

About Firewall Threat Defense Interfaces

Firewall Threat Defense includes data interfaces as well as a Management/Diagnostic interface.

When you attach a cable to an interface connection (physically or virtually), you need to configure the interface. At minimum, you need to name the interface and enable it for it to pass traffic. If the interface is a member of a bridge group, this is sufficient. For non-bridge group members, you also need to give the interface an IP address. If you intend to create VLAN subinterfaces rather than a single physical interface on a given port, you would typically configure the IP addresses on the subinterface, not on the physical interface. VLAN subinterfaces let you divide a physical interface into multiple logical interfaces that are tagged with different VLAN IDs, which is useful when you connect to a trunk port on a switch. You do not configure IP addresses on passive interfaces.

The **Interfaces** page include sub-pages for interface types: **Interfaces** (for physical interfaces), **Bridge Groups**, **Virtual Tunnel Interfaces**, **EtherChannels**, and **VLANs** (for the Firepower 1010). Note that Firepower 4100/9300 EtherChannels are listed on the **Interfaces** page and not on the **EtherChannel** page, because you can only modify EtherChannel parameters in FXOS, not in the Firewall Device Manager. Each page shows the available interfaces, their names, addresses, modes, and states. You can change the state of an interface, on or off, directly in the list of interfaces. The list shows the interface characteristics based on your

configuration. Use the open/close arrow on a bridge group, EtherChannel, or VLAN interface to view the member interfaces, which also appear by themselves in the appropriate list. You can also view subinterfaces for supported parent interfaces.

The following topics explain the limitations of configuring interfaces through the Firewall Device Manager as well as other interface management concepts.

Interface Modes

You can configure one of the following modes for each interface:

Routed

Each Layer 3 routed interface requires an IP address on a unique subnet. You would typically attach these interfaces to switches, a port on another router, or to an ISP/WAN gateway.

Passive

Passive interfaces monitor traffic flowing across a network using a switch SPAN (Switched Port Analyzer) 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 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.

Switch Port (Firepower 1010)

Switch ports forward traffic at Layer 2, using the switching function in hardware. Switch ports on the same VLAN can communicate with each other using hardware switching, and traffic is not subject to the Firewall Threat Defense security policy. Access ports accept only untagged traffic, and you can assign them to a single VLAN. Trunk ports accept untagged and tagged traffic, and can belong to more than one VLAN. You cannot configure the Management interface as a switch port.

BridgeGroupMember

A bridge group is a group of interfaces that the Firewall Threat Defense device bridges instead of routes. All interfaces are on the same network. The bridge group is represented by a Bridge Virtual Interface (BVI) that has an IP address on the bridge network.

You can route between routed interfaces and BVIs, if you name the BVI. In this case, the BVI acts as the gateway between member interfaces and routed interfaces. If you do not name the BVI, traffic on the bridge group member interfaces cannot leave the bridge group. Normally, you would name the interface so that you can route member interfaces to the internet.

One use for a bridge group in routed mode is to use extra interfaces on the Firewall Threat Defense device instead of an external switch. You can attach endpoints directly to bridge group member interfaces. You can also attach switches to add more endpoints to the same network as the BVI.

Management Interface

The physical port labeled Management (or for the Firewall Threat Defense, the Management0/0 virtual interface) actually has two separate interfaces associated with it.

- Management virtual interface—This IP address is used for system communication. This is the address the system uses for Smart Licensing and to retrieve database updates. You can open management sessions to it (Firewall Device Manager and CLI). You must configure a management address, which is defined on **System Settings** > **Management Interface**.

- Diagnostic virtual interface—You can use this interface to send syslog messages to an external syslog server. Configuring an IP address for the Diagnostic interface is optional. The main reason to configure the interface is if you want to use it for syslog messages. This interface appears, and is configurable, on the **Device** > **Interfaces** page. The Diagnostic interface only allows management traffic, and does not allow through traffic.

(Hardware devices.) One way to configure Management/Diagnostic is to not wire the physical port to a network. Instead, configure the Management IP address only, and configure it to use the data interfaces as the gateway for obtaining updates from the internet. Then, open the inside interfaces to HTTPS/SSH traffic (by default, HTTPS is enabled) and open the Firewall Device Manager using the inside IP address (see Configuring the Management Access List, on page 748).

For the Firewall Threat Defense Virtual, the recommended configuration is to attach Management0/0 to the same network as the inside interface, and use the inside interface as the gateway. Do not configure a separate address for Diagnostic.

Recommendations for Configuring a Separate Management Network

(Hardware devices.) If you want to use a separate management network, wire the Management interface to a switch or router.

For Firewall Threat Defense Virtual, attach Management0/0 to a separate network from any of the data interfaces. If you are still using the default IP addresses, you will need to change either the management IP address or the inside interface IP address, as they are on the same subnet.

Then, configure the following:

- Select **Device** > **System Settings** > **Management Interface** and configure IPv4 or IPv6 addresses (or both) on the attached network. If you want to, you can configure a DHCP server to provide IPv4 addresses to other endpoints on the network. If there is a router with a route to the internet on the management network, use that as the gateway. Otherwise, use the data interfaces as the gateway.
- Configure an address for the Diagnostic interface (on **Device** > **Interfaces**) only if you intend to send syslog messages through the interface to a syslog server. Otherwise, do not configure an address for Diagnostic; it is not needed. Any IP address you configure must be on the same subnet as the management IP address and cannot be the in DHCP server pool. For example, if you use 192.168.45.45 as the management address, and 192.168.45.46-192.168.45.254 as the DHCP pool, you can configure Diagnostic using any address from 192.168.45.1 to 192.168.45.44.

Security Zones

Each interface can be assigned to a single security zone. You then apply your security policy based on zones. 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.

Each zone has a mode that relates directly to the interface mode. You can add interfaces to the same mode security zone only.

For bridge groups, you add member interfaces to the zones, you cannot add the Bridge Virtual Interface (BVI).

You do not include the Management/Diagnostic interface in a zone. Zones apply to data interfaces only.

You can create security zones on the **Objects** page.

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, you configure the global address on the Bridge Virtual Interface (BVI), not on each member interface. You cannot specify any of the following as a global address.
  - Internally reserved IPv6 addresses: fd00::/56 (from=fd00:: to= fd00:0000:0000:00ff:ffff:ffff:ffff:ffff)
  - An unspecified address, such as ::/128
  - The loopback address, ::1/128
  - multicast addresses, ff00::/8
  - Link-local addresses, fe80::/10
- 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 Network Discovery functions such as address resolution and neighbor discovery. In a bridge group, enabling IPv6 on the BVI automatically configures link-local addresses for each bridge group member interface. Each interface must have its own address because the link-local address is only available on a segment, and is tied to the interface MAC 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. If you do not configure a global address, then you need to configure the link-local address, either automatically or manually.

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.

Guidelines and Limitations for Interfaces

The following topics cover some of the limitations for interfaces.

Limitations for Interface Configuration

When you use the Firewall Device Manager to configure the device, there are several limitations to interface configuration. If you need any of the following features, you must use the Firewall Management Center to configure the device.

- Routed firewall mode only is supported. You cannot configure transparent firewall mode interfaces.
- You can configure passive interfaces, but not ERSPAN interfaces.
- You cannot configure interfaces to be inline (in an inline set), or inline tap, for IPS-only processing. IPS-only mode interfaces bypass many firewall checks and only support IPS security policy. In comparison, 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 firewall mode traffic according to your security policy.
- You cannot configure redundant interfaces.
- The Firepower 4100/9300 supports EtherChannels, but you must perform all hardware configuration of EtherChannels in FXOS on the chassis. Firepower 4100/9300 EtherChannels appear in the Firewall Device Manager **Interfaces** page alongside single physical interfaces.
- You can only add one bridge group.
- Firewall Threat Defense supports IPv4 PPPoE on routed interfaces only. PPPoE is not supported on High Availability units.

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 1010	60
Firepower 1120	512
Firepower 1140, 1150	1024
Firepower 2100	1024
Secure Firewall 3100	1024
Firepower 4100	1024
Firepower 9300	1024
Firewall Threat Defense Virtual	50
ISA 3000	100

Configure a Physical Interface

At minimum, you must enable a physical interface to use it. You would also typically name it and configure IP addressing. You would not configure IP addressing if you intend to create VLAN subinterfaces, if you are configuring a passive mode interface, or if you intend to add the interface to a bridge group. Firepower

4100/9300 EtherChannels appear in the Firewall Device Manager **Interfaces** page alongside single physical interfaces, and this procedure also applies to those EtherChannels. You must perform all hardware configuration of Firepower 4100/9300 EtherChannels in FXOS on the chassis.

Note To configure physical interfaces as Firepower 1010 switch ports, see Configure VLAN Interfaces and Switch Ports, on page 259.

To configure physical interfaces as passive interfaces, see Configure a Physical Interface in Passive Mode, on page 278.

You can disable an interface to temporarily prevent transmission on the connected network. You do not need to remove the interface's configuration.

Procedure

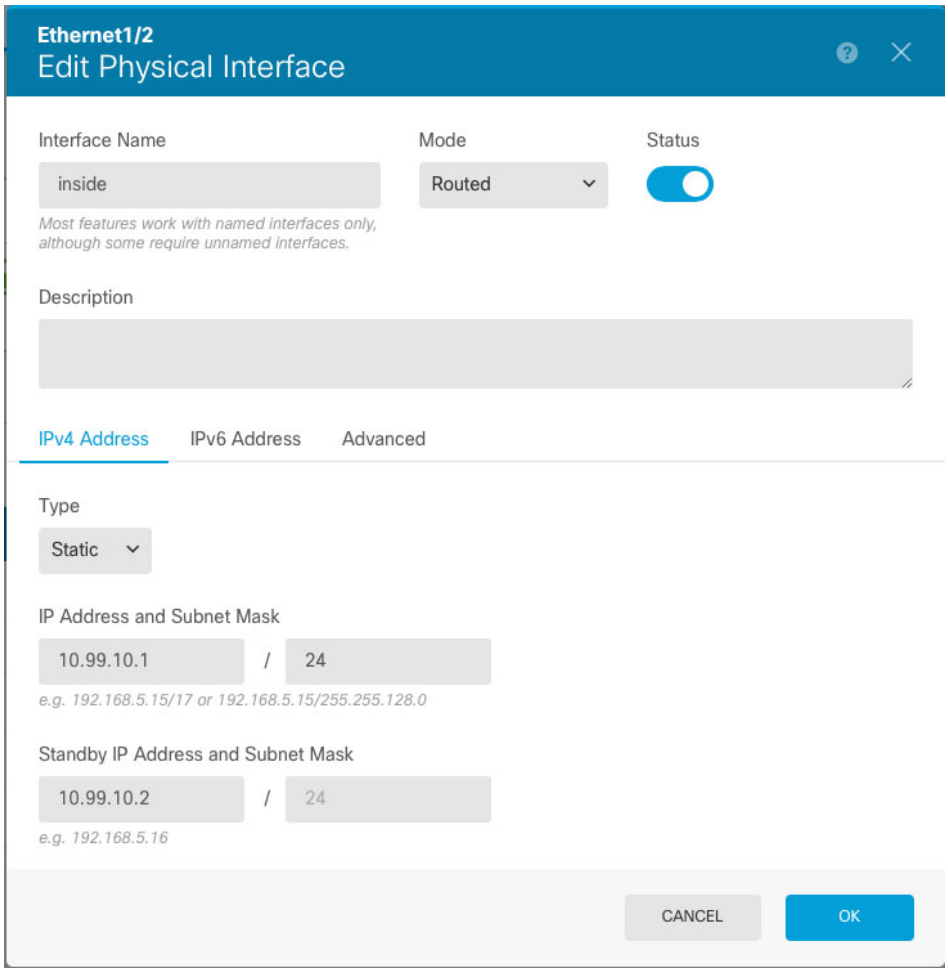
Step 1 Click **Device**, and then click the link in the **Interfaces** summary.

The **Interfaces** page is selected by default. The interfaces list shows physical interfaces, their names, addresses, and states.

Step 2 Click the edit icon ( ) for the physical interface you want to edit.

You cannot edit an interface that you are using as the failover or stateful failover link in a high availability configuration.

Step 3 Set the following:



a) Set the **Interface Name**.

Set the name for the interface, up to 48 characters. Alphabetic characters must be lower case. For example, **inside** or **outside**. Without a name, the rest of the interface configuration is ignored. Unless you configure subinterfaces, the interface should have a name. **Note:** Do not configure the name for an interface that you want to add to an EtherChannel.

Note
If you change the name, the change is automatically reflected everywhere you used the old name, including security zones, syslog server objects, and DHCP server definitions. However, you cannot remove the name until you first remove all configurations that use the name, because you typically cannot use an unnamed interface for any policy or setting.

b) Choose the **Mode**.

- **Routed**—Routed mode interfaces subject traffic to all firewall functions, including maintaining flows, tracking flow states at both IP and TCP layers, IP defragmentation, and TCP normalization, and your firewall policies. This is the normal interface mode.
- **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 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. If you select this mode, do not following the rest of this procedure. Instead, see Configure a Physical Interface in Passive Mode, on page 278. Note that you cannot configure IP addresses on passive interfaces.

- **Switch Port**—(Firepower 1010) Switch ports allow hardware switching between ports on the same VLAN. Switched traffic is not subject to the security policy. If you select this mode, do not following the rest of this procedure. Instead, see Configure VLAN Interfaces and Switch Ports, on page 259

If you later add this interface to a bridge group, the mode will automatically change to **BridgeGroupMember**. Note that you cannot configure IP addresses on bridge group member interfaces.

c) Set the **Status** slider to the enabled setting ().

For interfaces on a Firepower 4100/9300 device, you must also enable the interface in FXOS.

If you intend to configure subinterfaces for this physical interface, you are probably done. Click **Save** and continue with Configure VLAN Subinterfaces and 802.1Q Trunking, on page 270. Otherwise, continue.

Note
Even when configuring subinterfaces, it is valid to name the interface and supply IP addresses. This is not the typical setup, but if you know that is what you need, you can configure it.

d) (Optional) Set the **Description**.

The description can be up to 200 characters on a single line, without carriage returns.

Step 4 Click the **IPv4 Address** tab and configure the IPv4 address.

Select one of the following options from the **Type** field:

- **DHCP**—Choose this option if the address should be obtained from the DHCP server on the network. You cannot use this option if you configure high availability. Change the following options if necessary:
  - **Route Metric**—If you obtain the default route from the DHCP server, the administrative distance to the learned route, between 1 and 255. The default is 1.
  - **Obtain Default Route**—Whether to get the default route from the DHCP server. You would normally select this option, which is the default.
- **Static**—Choose this option if you want to assign an address that should not change. Type in the interface's IP address and the subnet mask for the network attached to the interface. For example, if you attach the 10.100.10.0/24 network, you could enter 10.100.10.1/24. Ensure that the address is not already used on the network.

  If you configured high availability, and you are monitoring this interface for HA, also configure a standby IP address on the same subnet. The standby address is used by this interface on the standby device. 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.

  Note
  If there is a DHCP server configured for the interface, you are shown the configuration. You can edit or delete the DHCP address pool. If you change the interface IP address to a different subnet, you must either delete the DHCP server, or configure an address pool on the new subnet, before you can save the interface changes. See Configuring the DHCP Server, on page 756.

- **PPPoE**—Choose this option if the address should be obtained using Point-to-point Protocol over Ethernet (PPPoE). PPPoE may be required if the interface is connected to a DSL modem, cable modem, or other connection to your ISP, and your ISP uses PPPoE to provide your IP address. You cannot use this option if you configure High Availability. Set the following values:
  - **Group Name**—Specify a group name of your choice to represent this connection.
  - **PPPoE Username**—Specify the username provided by your ISP.
  - **PPPoE Password**—Specify 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 Learned 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.
  - **Obtain Default Route from PPPoE**—Check this check box to enable obtaining the default route from the PPPoE server.
  - **IP Address Type**—Choose **Dynamic** to obtain the IP address from the PPPoE server. You can alternatively choose **Static** if you were assigned a static IP address from the ISP.

Step 5 (Optional.) Click the **IPv6 Address** tab and configure the IPv6 address.

- **State**—To enable IPv6 processing and to automatically configure the link-local address when you do not configure the global address, select **Enabled**. The link local address is generated based on the interface MAC addresses (*Modified* EUI-64 format).

  Note
  Disabling IPv6 does not disable IPv6 processing on an interface that is configured with an explicit IPv6 address or that is enabled for autoconfiguration.

- **Address Auto Configuration**—Select this option to have the address automatically configured. IPv6 stateless autoconfiguration will generate a global IPv6 address only if the link on which the device resides has a router configured to provide IPv6 services, including the advertisement of an IPv6 global prefix for use on the link. If IPv6 routing services are not available on the link, you will get a link-local IPv6 address only, which you cannot access outside of the device's immediate network link. The link local address is based on the Modified EUI-64 interface ID.

  Although RFC 4862 specifies that hosts configured for stateless autoconfiguration do not send Router Advertisement messages, the Firewall Threat Defense device does send Router Advertisement messages in this case. Select **Suppress RA** to suppress messages and conform to the RFC.
- **Static Address/Prefix**—If you do not use stateless autoconfiguration, enter the full static global IPv6 address and network prefix. For example, 2001:0DB8::BA98:0:3210/48. For more information on IPv6 addressing, see IPv6 Addressing, on page 238.

If you want to use the address as link local only, select the **Link - Local** option. Link local addresses are not accessible outside the local network. You cannot configure a link-local address on a bridge group interface.

Note
A link-local address should start with FE8, FE9, FEA, or FEB, for example fe80::20d:88ff:feee:6a82. 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.

- **Standby IP Address**—If you configure high availability, and you are monitoring this interface for HA, also configure a standby IPv6 address on the same subnet. The standby address is used by this interface on the standby device. 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.
- **Suppress RA**—Whether to suppress router advertisements. The Firewall Threat Defense can participate in router advertisements so that neighboring devices can dynamically learn a default router address. By default, router advertisement messages (ICMPv6 Type 134) are periodically sent out each IPv6 configured interface.

  Router advertisements are also 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 might want to suppress these messages on any interface for which you do not want the Firewall Threat Defense device to supply the IPv6 prefix (for example, the outside interface).

Step 6 (Optional.) Configure Advanced Options, on page 281.

The advanced settings have defaults that are appropriate for most networks. Edit them only if you are resolving network issues.

Step 7 Click **OK**.

What to do next

- Add the interfaces to the appropriate security zones. See Configuring Security Zones, on page 138.
- Register a fully-qualified domain name (FQDN) with your dynamic DNS service provider and configure DDNS to have the DNS server updated with the interface addresses, both IPv4 and IPv6. See Configuring Dynamic DNS, on page 759.

Configure Bridge Groups

A bridge group is a virtual interface that groups one or more interfaces. The main reason to group interfaces is to create a group of switched interfaces. Thus, you can attach workstations or other endpoint devices directly to the interfaces included in the bridge group. You do not need to connect them through a separate physical switch, although you can also attach a switch to a bridge group member.

The group members do not have IP addresses. Instead, all member interfaces share the IP address of the Bridge Virtual Interface (BVI). If you enable IPv6 on the BVI, member interfaces are automatically assigned unique link-local addresses.

You enable and disable the member interfaces individually. Thus, you can disable any unused interfaces without needing to remove them from the bridge group. The bridge group itself is always enabled.

You typically configure a DHCP server on the bridge group interface (BVI), which provides IP addresses for any endpoints connected through member interfaces. However, you can configure static addresses on the endpoints connected to the member interfaces if you prefer. All endpoints within the bridge group must have IP addresses on the same subnet as the bridge group IP address.

Guidelines and Limitations

- You can add one bridge group.
- Firewall Device Manager-defined EtherChannels are not supported as bridge group members. EtherChannels on the Firepower 4100/9300 can be bridge group members.
- You cannot configure bridge groups on Firepower 2100 series devices.
- For the Firepower 1010, you cannot mix logical VLAN interfaces and physical firewall interfaces in the same bridge group.
- The ISA 3000 comes pre-configured with bridge group BVI1 (not named, which means it does not participate in routing). BVI1 includes all data interfaces: GigabitEthernet1/1 (outside1), GigabitEthernet1/2 (inside1), GigabitEthernet1/3 (outside2), and GigabitEthernet1/4 (inside2). You must set the BVI1 IP address to match your network.

Before you begin

Configure the interfaces that will be members of the bridge group. Specifically, each member interface must meet the following requirements:

- The interface must have a name.
- The interface cannot have any IPv4 or IPv6 addresses defined for it, either static or served through DHCP. If you need to remove the address from an interface that you are currently using, you might also need to remove other configurations for the interface, such as static routes, DHCP server, or NAT rules, that depend on the interface having an address.
- You must remove the interface from its security zone (if it is in a zone), and delete any NAT rules for the interface, before you can add it to a bridge group.

Procedure

Step 1 Click **Device**, click the link in the **Interfaces** summary, then click **Bridge Groups**.

The bridge groups list shows existing bridge groups. Click the open/close arrow to view the member interfaces for each bridge group. Member interfaces also appear separately on the **Interfaces** or **VLANs** pages.

Step 2 Do one of the following:

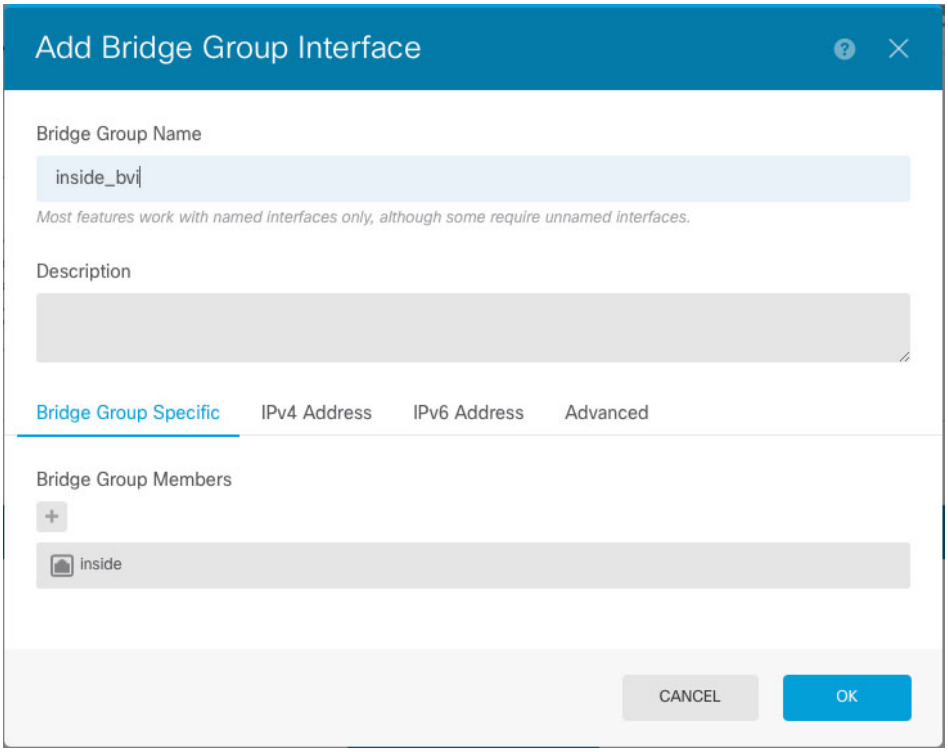
- Click the edit icon () for the BVI1 bridge group.

- Click **Create Bridge Group** or the plus icon (+) to create a new group.

  Note
  You can have a single bridge group. If you already have a bridge group defined, you should edit that group instead of trying to create a new one. If you need to create a new bridge group, you must first delete the existing bridge group.

- Click the delete icon ( ) for the bridge group if you no longer need it. When you delete a bridge group, its members become standard routed interfaces, and any NAT rules or security zone membership are retained. You can edit the interfaces to give them IP addresses. If you want to add them to a new bridge group, first you need to remove the NAT rules and remove the interface from its security zone.

Step 3 Configure the following:



a) (Optional) Set the **Interface Name**.

Set the name for the bridge group, up to 48 characters. Alphabetic characters must be lower case. For example, **inside** or **outside**. Set the name if you want this BVI to participate in routing between it and other named interfaces.

Note
If you change the name, the change is automatically reflected everywhere you used the old name, including security zones, syslog server objects, and DHCP server definitions. However, you cannot remove the name until you first remove all configurations that use the name, because you typically cannot use an unnamed interface for any policy or setting.

b) (Optional) Set the **Description**.

The description can be up to 200 characters on a single line, without carriage returns.

c) Edit the **Bridge Group Members** list.

You can add up to 64 interfaces or subinterfaces to a single bridge group.

- Add an interface—Click the plus icon (+) , click one or more interfaces, and then click **OK**.
- Remove an interface—Mouse over an interface and click the **x** on the right side.

Step 4 Click the **IPv4 Address** tab and configure the IPv4 address.

Select one of the following options from the **Type** field:

- **Static**—Choose this option if you want to assign an address that should not change. Type in the bridge group's IP address and the subnet mask. All attached endpoints will be on this network. Ensure that the address is not already used on the network.

  If you configured high availability, and you are monitoring this interface for HA, also configure a standby IP address on the same subnet. The standby address is used by this interface on the standby device. 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.

  Note
  If there is a DHCP server configured for the interface, you are shown the configuration. You can edit or delete the DHCP address pool. If you change the interface IP address to a different subnet, you must either delete the DHCP server, or configure an address pool on the new subnet, before you can save the interface changes. See Configuring the DHCP Server, on page 756.

- **Dynamic** (DHCP)—Choose this option if the address should be obtained from the DHCP server on the network. This is not the typical option for bridge groups, but you can configure it if needed. You cannot use this option if you configure high availability.Change the following options if necessary:
  - **Route Metric**—If you obtain the default route from the DHCP server, the administrative distance to the learned route, between 1 and 255. The default is 1.
  - **Obtain Default Route**—Whether to get the default route from the DHCP server. You would normally select this option, which is the default.

Step 5 (Optional.) Click the **IPv6 Address** tab and configure the IPv6 address.

- **State**—To enable IPv6 processing and to automatically configure the link-local address when you do not configure the global address, select **Enabled**. The link local address is generated based on the interface MAC addresses (*Modified* EUI-64 format).

  Note
  Disabling IPv6 does not disable IPv6 processing on an interface that is configured with an explicit IPv6 address or that is enabled for autoconfiguration.

- **Static Address/Prefix**—If you do not use stateless autoconfiguration, enter the full static global IPv6 address and network prefix. For example, 2001:0DB8::BA98:0:3210/48. For more information on IPv6 addressing, see IPv6 Addressing, on page 238.

  If you want to use the address as link local only, select the **Link - Local** option. Link local addresses are not accessible outside the local network. You cannot configure a link-local address on a bridge group interface.

  Note

A link-local address should start with FE8, FE9, FEA, or FEB, for example fe80::20d:88ff:feee:6a82. 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.

- **Standby IP Address**—If you configure high availability, and you are monitoring this interface for HA, also configure a standby IPv6 address on the same subnet. The standby address is used by this interface on the standby device. 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.
- **Suppress RA**—Whether to suppress router advertisements. The Firewall Threat Defense can participate in router advertisements so that neighboring devices can dynamically learn a default router address. By default, router advertisement messages (ICMPv6 Type 134) are periodically sent out each IPv6 configured interface.

  Router advertisements are also 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 might want to suppress these messages on any interface for which you do not want the Firewall Threat Defense device to supply the IPv6 prefix (for example, the outside interface).

Step 6 (Optional.) Configure Advanced Options, on page 281.

You configure most advanced options on bridge group *member* interfaces, but some are available for the bridge group interface.

The advanced settings have defaults that are appropriate for most networks. Edit them only if you are resolving network issues.

Step 7 Click **OK**.

What to do next

- Ensure that all member interfaces that you intend to use are enabled.
- Configure a DHCP server for the bridge group. See Configuring the DHCP Server, on page 756.
- Add the member interfaces to the appropriate security zones. See Configuring Security Zones, on page 138.
- Ensure that policies, such as identity, NAT, and access, supply the required services for the bridge group and member interfaces.

Configure EtherChannels

This section describes EtherChannels and how to configure them.

Note You can add EtherChannels in the Firewall Device Manager to the following models:

- Firepower 1000
- Firepower 2100
- Secure Firewall 3100
- ISA 3000

You cannot use Firepower 1010 or Secure Firewall 1210/1220 switch ports or VLAN interfaces in EtherChannels.

The Firepower 4100/9300 supports EtherChannels, but you must perform all hardware configuration of EtherChannels in FXOS on the chassis. Firepower 4100/9300 EtherChannels appear in the Firewall Device Manager Interfaces page alongside single physical interfaces. You also cannot configure EtherChannels in the Firewall Device Manager on other models, such as the Firewall Threat Defense Virtual.

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.

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.

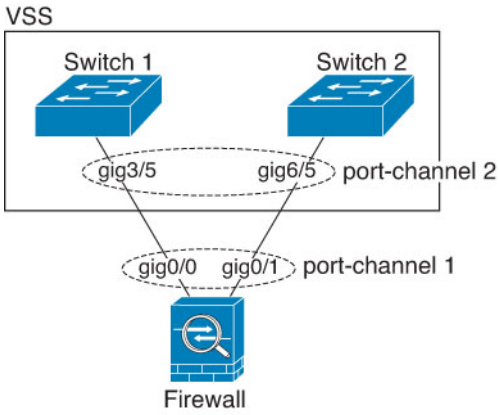
The EtherChannel aggregates the traffic across all the available active interfaces in the channel. The interface is selected using a proprietary hash algorithm, based on source or destination MAC addresses, IP addresses, TCP and UDP port numbers and VLAN numbers.

Connecting to an EtherChannel on Another Device

The device to which you connect the Firewall Threat Defense 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 Firewall Threat Defense 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.

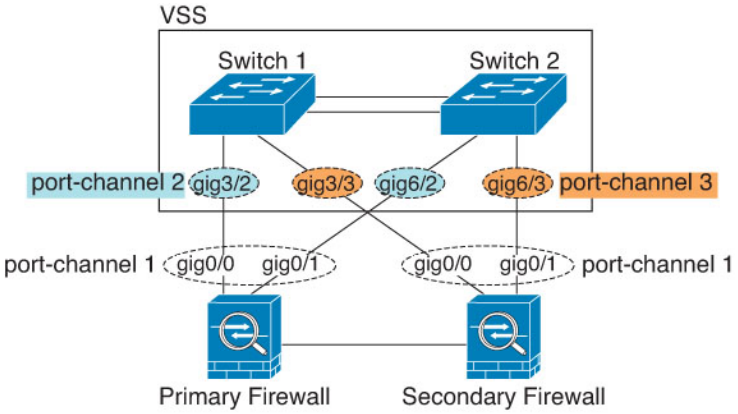
Figure 7: Connecting to a VSS/vPC



Note If the Firewall Threat Defense device is in transparent firewall mode, and you place the Firewall Threat Defense device between two sets of VSS/vPC switches, then be sure to disable Unidirectional Link Detection (UDLD) on any switch ports connected to the Firewall Threat Defense 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 Firewall Threat Defense device in an Active/Standby failover deployment, then you need to create separate EtherChannels on the switches in the VSS/vPC, one for each Firewall Threat Defense device. On each Firewall Threat Defense device, a single EtherChannel connects to both switches. Even if you could group all switch interfaces into a single EtherChannel connecting to both Firewall Threat Defense devices (in this case, the EtherChannel will not be established because of the separate Firewall Threat Defense system IDs), a single EtherChannel would not be desirable because you do not want traffic sent to the standby Firewall Threat Defense device.

Figure 8: 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.
- On—The EtherChannel is always on, and LACP is not used. An "on" EtherChannel can only establish a connection with another "on" EtherChannel.

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 Firewall Threat Defense 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 and Secure Firewall 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.

Guidelines for EtherChannels

Bridge Group

Firewall Device Manager-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.
- 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 can add EtherChannels in the Firewall Device Manager to the following models:
  - Firepower 1000
  - Firepower 2100
  - Secure Firewall 3100
  - ISA 3000

  The Firepower 4100/9300 supports EtherChannels, but you must perform all hardware configuration of EtherChannels in FXOS on the chassis. Firepower 4100/9300 EtherChannels appear in the Firewall Device Manager Interfaces page alongside single physical interfaces.
- You cannot use Firepower 1010 switch ports or VLAN interfaces in EtherChannels.

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.
- When you add the first member interface, it sets the required hardware properties of all member interfaces.

- The media type of member interfaces can be either RJ-45 or SFP; SFPs of different types (copper and fiber) can be mixed. You cannot mix RJ-45 and SFP interfaces.
- All interfaces must be set to the same speed and duplex.
- The first interface sets the speed *capacity*, which cannot be changed later.
  - For SFP Detect interfaces—You can include interfaces with different speed capacities as long as they have a common speed. When you set the speed to SFP Detect (the default), the speed will be dynamically set to the highest common speed. If you later change the member interfaces so that the common speed is now higher, the EtherChannel speed will also be higher automatically.

    You can set a specific speed, but only speeds that are available on the first member interface. For example, if your first interface is 1/10GB, then the available speeds for the EtherChannel will be 1GB, 10GB, and SFP Detect. If you later remove the 1/10GB interfaces and replace them with 1/10/25GB interfaces, you cannot manually set the speed to 25GB. In this case, you can use SFP Detect to use the 25GB speed.
  - For non-SFP Detect interfaces—All additional interfaces must have the same speed capacity. For example, if your first interface speed capacity is 10MB/100MB/1GB, you must add other 10MB/100MB/1GB interfaces. You can set the EtherChannel (and its member interfaces) to any of those speeds. You cannot later add 1/10GB interfaces to the EtherChannel, even if you remove the lower capacity interfaces. You also 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 Firewall Threat Defense EtherChannel must also support 802.3ad EtherChannels.
- The Firewall Threat Defense 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 Firewall Threat Defense device will drop the tagged LACPDUs. Be sure to disable native VLAN tagging on the neighboring switch.
- The LACP rate depends on the model. When you set the rate (normal or fast), the device requests that rate from the connecting switch. In return, the device will send at the rate requested by the connecting switch. We recommend that you set the same rate on both sides.
  - Firepower 4100/9300—The LACP rate is set to fast by default in FXOS, but you can configure it as normal (also known as slow).
  - Secure Firewall 3100—The LACP rate is set to normal (slow) by default, but you can configure it as fast on the device.
  - All other models—The LACP rate set to normal (also known as slow), and it is not configurable, which means the device will always request a slow rate from the connecting switch. We recommend setting the rate on the switch to slow, so both sides send LACP messages at the same rate.
- In Cisco IOS software versions earlier than 15.1(1)S2, Firewall Threat Defense did not support connecting an EtherChannel to a switch stack. With default switch settings, if the Firewall Threat Defense 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 Firewall Threat Defense configuration refers to the logical EtherChannel interface instead of the member physical interfaces.

Add an EtherChannel

Add an EtherChannel and assign member interfaces to it.

Note You can add EtherChannels in the Firewall Device Manager to the following models:

- Firepower 1000
- Firepower 2100
- Secure Firewall 3100
- ISA 3000

You cannot use switch ports or VLAN interfaces in EtherChannels.

The Firepower 4100/9300 supports EtherChannels, but you must perform all hardware configuration of EtherChannels in FXOS on the chassis. Firepower 4100/9300 EtherChannels appear in the Firewall Device Manager Interfaces page alongside single physical interfaces.

Before you begin

- When you add the first member interface, it sets the required hardware properties of all member interfaces. See Guidelines for EtherChannels, on page 252 for more details about member interface requirements.
- Member interfaces cannot be named.

Caution If you are using an interface already in your configuration, removing the name will clear any configuration that refers to the interface.

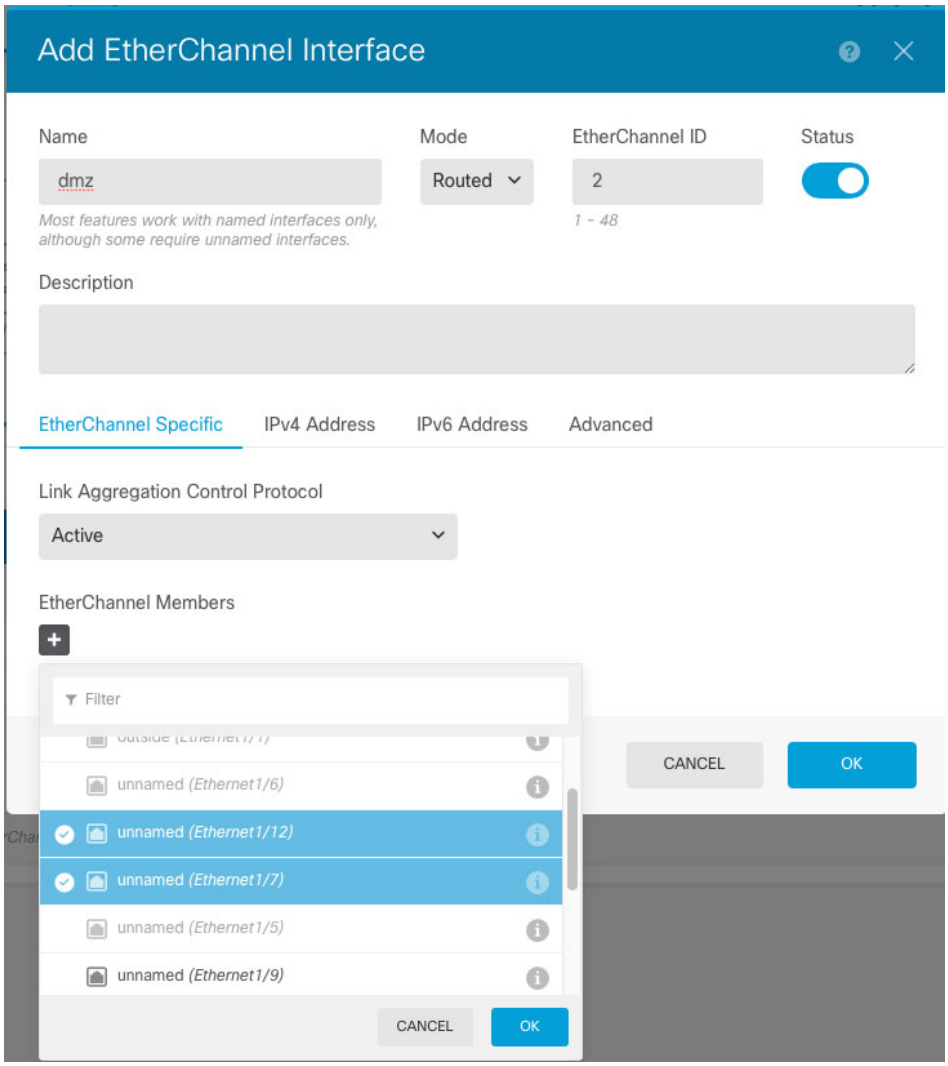
Procedure

Step 1 Click **Device**, click the link in the **Interfaces** summary, and then click **EtherChannels**.

The EtherChannels list shows existing EtherChannels, their names, addresses, and states. Click the open/close arrow to view the member interfaces for each EtherChannel. Member interfaces also appear separately on the **Interfaces** page.

Step 2 Click **Create EtherChannel** (if there are no current EtherChannels) or the plus icon ( ) then **EtherChannel** to create a new EtherChannel.

Step 3 Configure the following:



a) Set the **Interface Name**.

Set the name for the EtherChannel, up to 48 characters. Alphabetic characters must be lower case. For example, **inside** or **outside**.

Note
If you change the name, the change is automatically reflected everywhere you used the old name, including security zones, syslog server objects, and DHCP server definitions. However, you cannot remove the name until you first remove all configurations that use the name, because you typically cannot use an unnamed interface for any policy or setting.

b) Set the **Mode**.

- **Routed**—Routed mode interfaces subject traffic to all firewall functions, including maintaining flows, tracking flow states at both IP and TCP layers, IP defragmentation, and TCP normalization, and your firewall policies. Use this mode if you intend for traffic to go through the interface. This is the normal interface mode.

- **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 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. If you select this mode, do not following the rest of this procedure. Instead, see Configure a Physical Interface in Passive Mode, on page 278.

c) Set the **EtherChannel ID**, between 1 and 48 (1 and 8 for the Firepower 1010).

d) Set the **Status** slider to the enabled setting ( ).

e) (Optional) Set the **Description**.

The description can be up to 200 characters on a single line, without carriage returns.

f) Choose the **EtherChannel Mode**.

- **Active**—(Recommended) 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.
- **On**—The EtherChannel is always on, and LACP is not used. An "on" EtherChannel can only establish a connection with another "on" EtherChannel.

g) (Secure Firewall 3100 only) Choose the **EtherChannel Rate**. You should match the setting on the connected switch.

- **Default**—The default is **Normal** (slow, every 30 seconds).
- **Normal**—Receives LACP data units every 30 seconds.
- **Fast**—Receives LACP data units every second.

h) Add **EtherChannel Members**.

You can add up to 8 (unnamed) interfaces to the EtherChannel.

- Add an interface—Click the plus icon ( ) , click one or more interfaces, and then click **OK**.
- Remove an interface—Mouse over an interface and click the **x** on the right side.

Step 4 Click the **IPv4 Address** tab and configure the IPv4 address.

Select one of the following options from the **Type** field:

- **DHCP**—Choose this option if the address should be obtained from the DHCP server on the network. You cannot use this option if you configure high availability. Change the following options if necessary:
  - **Route Metric**—If you obtain the default route from the DHCP server, the administrative distance to the learned route, between 1 and 255. The default is 1.
  - **Obtain Default Route**—Whether to get the default route from the DHCP server. You would normally select this option, which is the default.
- **Static**—Choose this option if you want to assign an address that should not change. Type in the interface's IP address and the subnet mask for the network attached to the interface. For example, if you attach the

10.100.10.0/24 network, you could enter 10.100.10.1/24. Ensure that the address is not already used on the network.

If you configured high availability, and you are monitoring this interface for HA, also configure a standby IP address on the same subnet. The standby address is used by this interface on the standby device. 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.

Note
If there is a DHCP server configured for the interface, you are shown the configuration. You can edit or delete the DHCP address pool. If you change the interface IP address to a different subnet, you must either delete the DHCP server, or configure an address pool on the new subnet, before you can save the interface changes. See Configuring the DHCP Server, on page 756.

- **PPPoE**—Choose this option if the address should be obtained using Point-to-point Protocol over Ethernet (PPPoE). PPPoE may be required if the interface is connected to a DSL modem, cable modem, or other connection to your ISP, and your ISP uses PPPoE to provide your IP address. You cannot use this option if you configure High Availability. Set the following values:
  - **Group Name**—Specify a group name of your choice to represent this connection.
  - **PPPoE Username**—Specify the username provided by your ISP.
  - **PPPoE Password**—Specify 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 Learned 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.
  - **Obtain Default Route from PPPoE**—Check this check box to enable obtaining the default route from the PPPoE server.
  - **IP Address Type**—Choose **Dynamic** to obtain the IP address from the PPPoE server. You can alternatively choose **Static** if you were assigned a static IP address from the ISP.

Step 5 (Optional.) Click the **IPv6 Address** tab and configure the IPv6 address.

- **State**—To enable IPv6 processing and to automatically configure the link-local address when you do not configure the global address, select **Enabled**. The link local address is generated based on the interface MAC addresses (*Modified* EUI-64 format).

  Note
  Disabling IPv6 does not disable IPv6 processing on an interface that is configured with an explicit IPv6 address or that is enabled for autoconfiguration.

- **Address Auto Configuration**—Select this option to have the address automatically configured. IPv6 stateless autoconfiguration will generate a global IPv6 address only if the link on which the device resides has a router configured to provide IPv6 services, including the advertisement of an IPv6 global prefix

for use on the link. If IPv6 routing services are not available on the link, you will get a link-local IPv6 address only, which you cannot access outside of the device's immediate network link. The link local address is based on the Modified EUI-64 interface ID.

Although RFC 4862 specifies that hosts configured for stateless autoconfiguration do not send Router Advertisement messages, the Firewall Threat Defense device does send Router Advertisement messages in this case. Select **Suppress RA** to suppress messages and conform to the RFC.

- **Static Address/Prefix**—If you do not use stateless autoconfiguration, enter the full static global IPv6 address and network prefix. For example, 2001:0DB8::BA98:0:3210/48. For more information on IPv6 addressing, see IPv6 Addressing, on page 238.

  If you want to use the address as link local only, select the **Link - Local** option. Link local addresses are not accessible outside the local network. You cannot configure a link-local address on a bridge group interface.

  Note
  A link-local address should start with FE8, FE9, FEA, or FEB, for example fe80::20d:88ff:feee:6a82. 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.

- **Standby IP Address**—If you configure high availability, and you are monitoring this interface for HA, also configure a standby IPv6 address on the same subnet. The standby address is used by this interface on the standby device. 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.

- **Suppress RA**—Whether to suppress router advertisements. The Firewall Threat Defense can participate in router advertisements so that neighboring devices can dynamically learn a default router address. By default, router advertisement messages (ICMPv6 Type 134) are periodically sent out each IPv6 configured interface.

  Router advertisements are also 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 might want to suppress these messages on any interface for which you do not want the Firewall Threat Defense device to supply the IPv6 prefix (for example, the outside interface).

Step 6 Set the speed of member interfaces by clicking **Advanced**, and setting the speed.

You can also configure other Advanced options. See Configure Advanced Options, on page 281.

Step 7 Click **OK**.

What to do next

- Add the EtherChannels to the appropriate security zones. See Configuring Security Zones, on page 138.

Configure VLAN Interfaces and Switch Ports

For models with an internal switch, you can configure each interface to run as a regular firewall interface or as a Layer 2 hardware switch port. This section includes tasks for starting your switch port configuration, including enabling or disabling the switch mode and creating VLAN interfaces and assigning switch ports to VLANs. This section also describes how to customize Power over Ethernet (PoE) on supported interfaces.

Understanding switch ports and interfaces

Ports and interfaces

For each physical interface, you can set its operation as a firewall interface or as a switch port. See the following information about physical interface and port types as well as logical VLAN interfaces to which you assign switch ports:

- Physical firewall interface—In routed mode, these interfaces forward traffic between networks at Layer 3, using the configured security policy to apply firewall and VPN services. In routed mode, you can also use Integrated Routing and Bridging with some interfaces as bridge group members and others as Layer 3 interfaces. By default, the Ethernet 1/1 interface is configured as a firewall interface. You can also configure these interfaces to be IPS-only (passive interfaces).
- Physical switch port—Switch ports forward traffic at Layer 2, using the switching function in hardware. Switch ports on the same VLAN can communicate with each other using hardware switching, and traffic is not subject to the Firewall Threat Defense security policy. Access ports accept only untagged traffic, and you can assign them to a single VLAN. Trunk ports accept untagged and tagged traffic, and can belong to more than one VLAN. By default, Ethernet 1/2 and higher are configured as access switch ports on VLAN 1. You cannot configure the Management interface as a switch port.
- Logical VLAN interface—These interfaces operate the same as physical firewall interfaces, with the exception being that you cannot create subinterfaces, IPS-only interfaces (inline sets and passive interfaces), or EtherChannel interfaces. When a switch port needs to communicate with another network, then the Firewall Threat Defense device applies the security policy to the VLAN interface and routes to another logical VLAN interface or firewall interface. You can even use Integrated Routing and Bridging with VLAN interfaces as bridge group members. Traffic between switch ports on the same VLAN are not subject to the Firewall Threat Defense security policy, but traffic between VLANs in a bridge group are subject to the security policy, so you may choose to layer bridge groups and switch ports to enforce the security policy between certain segments.

Power Over Ethernet

PoE is available on Ethernet 1/7 and Ethernet 1/8 using IEEE 802.3af (PoE) and 802.3at (PoE+) up to 30 watts per port.

PoE+ uses Link Layer Discovery Protocol (LLDP) to negotiate the power level. Power is only supplied when needed.

If you shut down the interface, then you disable power to the device.

Prerequisites for switch ports

Model support

- Firepower 1010

Guidelines for switch ports

High availability

- You should not use the switch port functionality when using High availability. Because the switch ports operate in hardware, they continue to pass traffic on both the active *and* the standby units. High availability is designed to prevent traffic from passing through the standby unit, but this feature does not extend to switch ports. In a normal High availability network setup, active switch ports on both units will lead to network loops. We suggest that you use external switches for any switching capability. Note that VLAN interfaces can be monitored by failover, while switch ports cannot. Theoretically, you can put a single switch port on a VLAN and successfully use High availability, but a simpler setup is to use physical firewall interfaces instead.
- You can only use a firewall interface as the failover link.

Logical VLAN interfaces (SVIs)

- If you also use VLAN subinterfaces on a firewall interface, you cannot use the same VLAN ID as for a logical VLAN interface.
- MAC Addresses:
  - All VLAN interfaces share a MAC address. Ensure that any connected switches can support this scenario. If the connected switches require unique MAC addresses, you can manually assign MAC addresses. See Configure Advanced Options, on page 281.

Bridge groups

You cannot mix logical VLAN interfaces and physical firewall interfaces in the same bridge group.

VLAN interface and switch port unsupported features

VLAN interfaces and switch ports do not support:

- Dynamic routing
- Multicast routing
- Equal-Cost Multi-Path routing (ECMP)
- Passive interfaces
- EtherChannels—Switch ports cannot be part of an EtherChannel. PoE is also not supported on a port in an EtherChannel.
- Failover and state link

Other Guidelines and Limitations

- You can configure a maximum of 60 named interfaces.
- You cannot configure the Management interface as a switch port.

Default Settings

- Ethernet 1/1 is a firewall interface.
- Ethernet 1/2 through Ethernet 1/8 are switch ports assigned to VLAN 1.
- Default Speed and Duplex—By default, the speed and duplex are set to auto-negotiate.

Configure a VLAN Interface

This section describes how to configure VLAN interfaces for use with associated switch ports. You must first configure a VLAN interface for each VLAN you intend to assign to a switch port.



Note If you only want to enable switching between switch ports on a particular VLAN, and you do not want to route between the VLAN and other VLANs or firewall interfaces, then leave the VLAN interface name empty. In this case, you also do not need to configure an IP address; any IP configuration is ignored.

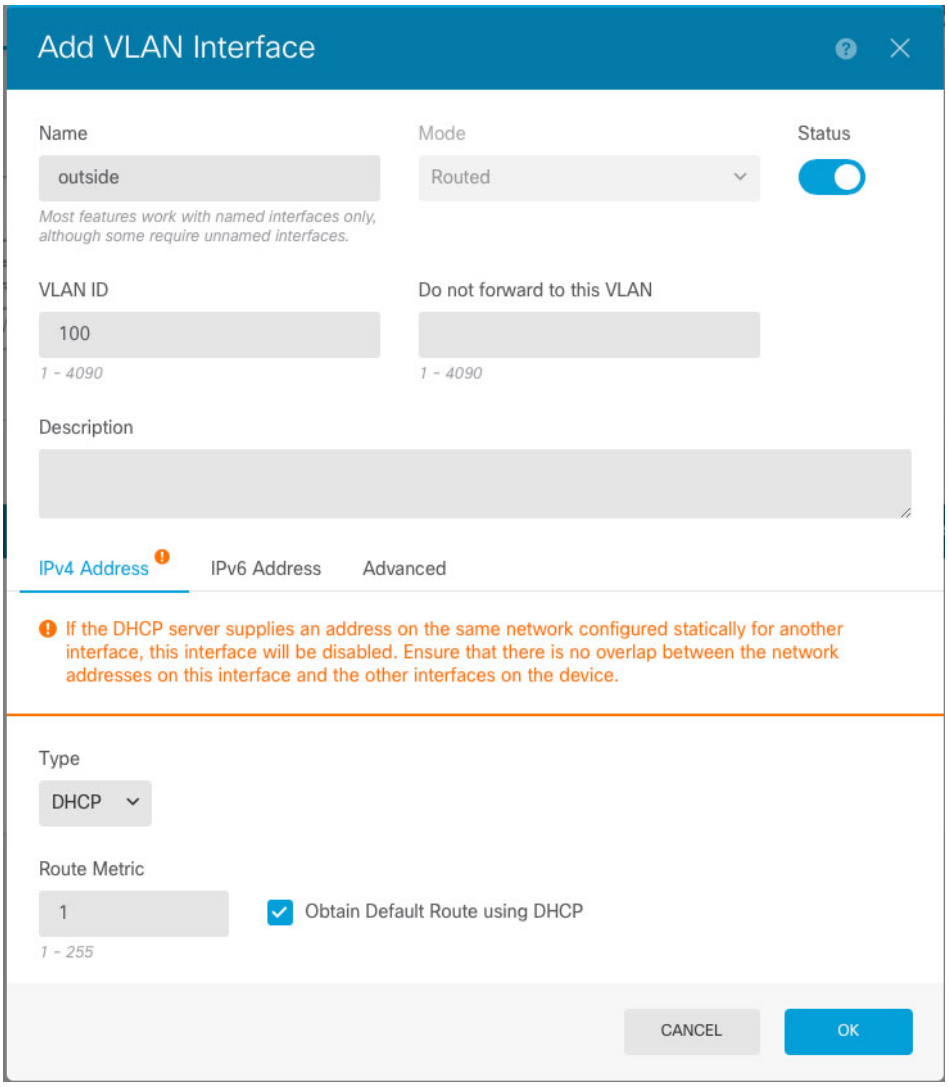
Procedure

Step 1 Click **Device**, click the link in the **Interfaces** summary, then click **VLANs**.

The VLANs list shows existing VLAN interfaces. Click the open/close arrow to view the switch ports associated with each VLAN. Switch ports also appear separately on the **Interfaces** page.

Step 2 Click **Create VLAN Interface** (if there are no current VLANs) or the plus icon (+) to create a new VLAN interface.

Step 3 Configure the following:



a) Set the **Interface Name**.

Set the name for the VLAN, up to 48 characters. Alphabetic characters must be lower case. For example, **inside** or **outside**.

If you do not want to route between the VLAN and other VLANs or firewall interfaces, then leave the VLAN interface name empty.

Note
If you change the name, the change is automatically reflected everywhere you used the old name, including security zones, syslog server objects, and DHCP server definitions. However, you cannot remove the name until you first remove all configurations that use the name, because you typically cannot use an unnamed interface for any policy or setting.

b) Leave the **Mode** as **Routed**.

If you later add this VLAN interface to a bridge group, then the mode will automatically change to **BridgeGroupMember**. You cannot configure IP addresses on bridge group member interfaces.

c) Set the **Status** slider to the enabled setting ().

d) Set the **VLAN ID**, between 1 and 4070.

You cannot change the VLAN ID after you save the interface; the VLAN ID is both the VLAN tag used, and the interface ID in your configuration.

e) (Optional) In the **Do not forward to this VLAN** field, enter a VLAN ID to which this VLAN interface cannot initiate traffic.

For example, you have one VLAN assigned to the outside for internet access, one VLAN assigned to an inside business network, and a third VLAN assigned to your home network. The home network does not need to access the business network, so you can use the Block Traffic From this Interface to option on the home VLAN; the business network can access the home network, but the home network cannot access the business network.

f) (Optional) Set the **Description**.

The description can be up to 200 characters on a single line, without carriage returns.

Step 4 Click the **IPv4 Address** tab and configure the IPv4 address.

Select one of the following options from the **Type** field:

- **DHCP**—Choose this option if the address should be obtained from the DHCP server on the network. You cannot use this option if you configure high availability. Change the following options if necessary:
  - **Route Metric**—If you obtain the default route from the DHCP server, the administrative distance to the learned route, between 1 and 255. The default is 1.
  - **Obtain Default Route**—Whether to get the default route from the DHCP server. You would normally select this option, which is the default.

- **Static**—Choose this option if you want to assign an address that should not change. Type in the interface's IP address and the subnet mask for the network attached to the interface. For example, if you attach the 10.100.10.0/24 network, you could enter 10.100.10.1/24. Ensure that the address is not already used on the network.

  If you configured high availability, and you are monitoring this interface for HA, also configure a standby IP address on the same subnet. The standby address is used by this interface on the standby device. 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.

  Note
  If there is a DHCP server configured for the interface, you are shown the configuration. You can edit or delete the DHCP address pool. If you change the interface IP address to a different subnet, you must either delete the DHCP server, or configure an address pool on the new subnet, before you can save the interface changes. See Configuring the DHCP Server, on page 756.

- **PPPoE**—Choose this option if the address should be obtained using Point-to-point Protocol over Ethernet (PPPoE). PPPoE may be required if the interface is connected to a DSL modem, cable modem, or other connection to your ISP, and your ISP uses PPPoE to provide your IP address. You cannot use this option if you configure High Availability. Set the following values:
  - **Group Name**—Specify a group name of your choice to represent this connection.
  - **PPPoE Username**—Specify the username provided by your ISP.

- **PPPoE Password**—Specify 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 Learned 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.
- **Obtain Default Route from PPPoE**—Check this check box to enable obtaining the default route from the PPPoE server.
- **IP Address Type**—Choose **Dynamic** to obtain the IP address from the PPPoE server. You can alternatively choose **Static** if you were assigned a static IP address from the ISP.

Step 5 (Optional.) Click the **IPv6 Address** tab and configure the IPv6 address.

- **State**—To enable IPv6 processing and to automatically configure the link-local address when you do not configure the global address, select **Enabled**. The link local address is generated based on the interface MAC addresses (*Modified* EUI-64 format).

  Note
  Disabling IPv6 does not disable IPv6 processing on an interface that is configured with an explicit IPv6 address or that is enabled for autoconfiguration.

- **Address Auto Configuration**—Select this option to have the address automatically configured. IPv6 stateless autoconfiguration will generate a global IPv6 address only if the link on which the device resides has a router configured to provide IPv6 services, including the advertisement of an IPv6 global prefix for use on the link. If IPv6 routing services are not available on the link, you will get a link-local IPv6 address only, which you cannot access outside of the device's immediate network link. The link local address is based on the Modified EUI-64 interface ID.

  Although RFC 4862 specifies that hosts configured for stateless autoconfiguration do not send Router Advertisement messages, the Firewall Threat Defense device does send Router Advertisement messages in this case. Select **Suppress RA** to suppress messages and conform to the RFC.
- **Static Address/Prefix**—If you do not use stateless autoconfiguration, enter the full static global IPv6 address and network prefix. For example, 2001:0DB8::BA98:0:3210/48. For more information on IPv6 addressing, see IPv6 Addressing, on page 238.

  If you want to use the address as link local only, select the **Link - Local** option. Link local addresses are not accessible outside the local network. You cannot configure a link-local address on a bridge group interface.

  Note
  A link-local address should start with FE8, FE9, FEA, or FEB, for example fe80::20d:88ff:feee:6a82. 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.

- **Standby IP Address**—If you configure high availability, and you are monitoring this interface for HA, also configure a standby IPv6 address on the same subnet. The standby address is used by this interface on the standby device. 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.
- **Suppress RA**—Whether to suppress router advertisements. The Firewall Threat Defense can participate in router advertisements so that neighboring devices can dynamically learn a default router address. By default, router advertisement messages (ICMPv6 Type 134) are periodically sent out each IPv6 configured interface.

  Router advertisements are also 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 might want to suppress these messages on any interface for which you do not want the Firewall Threat Defense device to supply the IPv6 prefix (for example, the outside interface).

Step 6 (Optional.) Configure Advanced Options, on page 281.

The advanced settings have defaults that are appropriate for most networks. Edit them only if you are resolving network issues.

Step 7 Click **OK**.

What to do next

- Add the VLANs to the appropriate security zones. See Configuring Security Zones, on page 138.

Configure Switch Ports as Access Ports

To assign a switch port to a single VLAN, configure it as an access port. By default, Ethernet1/2 through Ethernet 1/8 switch ports are enabled and assigned to VLAN 1.

Note The Firepower 1010 does not support Spanning Tree Protocol for loop detection in the network. Therefore you must ensure that any connection with the Firewall Threat Defense device does not end up in a network loop.

Before you begin

Add a VLAN interface for the VLAN ID to which you want to assign the access port. Access ports accept only untagged traffic. See Configure a VLAN Interface, on page 261.

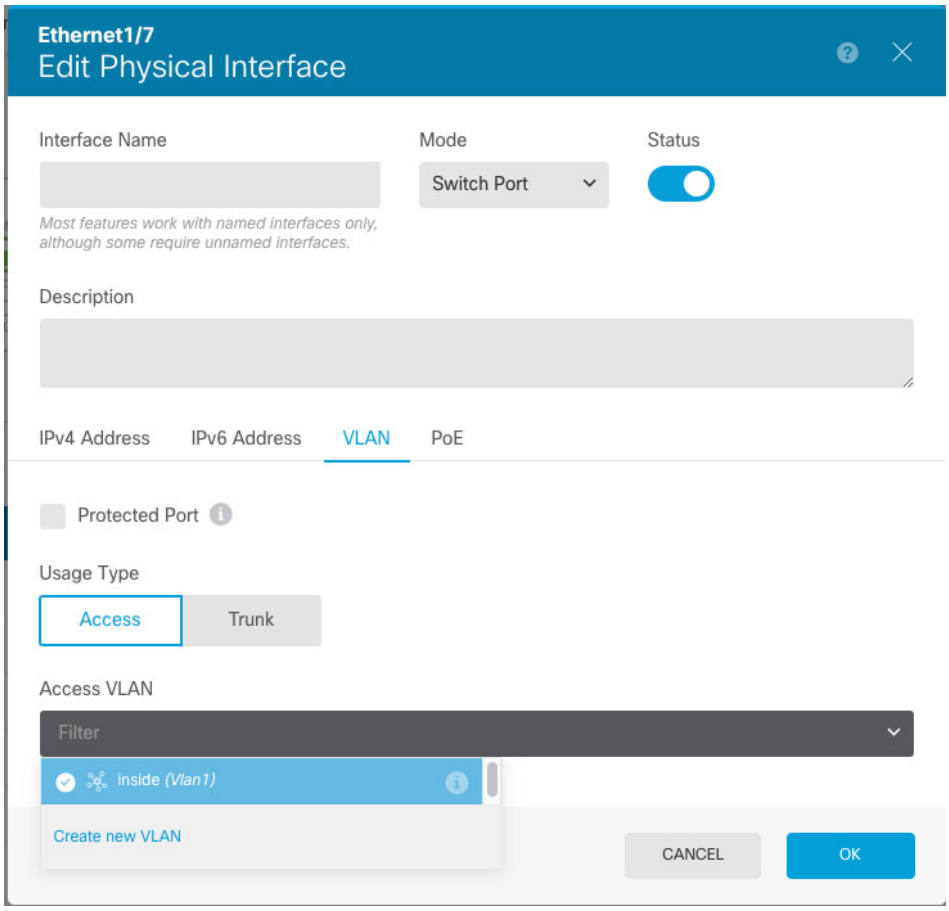
Procedure

Step 1 Click **Device**, and then click the link in the **Interfaces** summary.

The **Interfaces** page is selected by default. The interfaces list shows physical interfaces, their names, addresses, and states.

Step 2 Click the edit icon ( ) for the physical interface you want to edit.

Step 3 Set the following:



a) Do not set the **Interface Name** for switch ports; only the associated VLAN interface is a named interface.

b) Set the **Mode** to **Switch Port**.

c) Set the **Status** slider to the enabled setting ( ).

d) (Optional) Set the **Description**.

The description can be up to 200 characters on a single line, without carriage returns.

Step 4 Click **VLAN** to set the following:

a) (Optional) Check the **Protected Port** check box to set this switch port as protected, so you can prevent the switch port from communicating with other protected switch ports on the same VLAN.

You might want to prevent switch ports from communicating with each other if: the devices on those switch ports are primarily accessed from other VLANs; you do not need to allow intra-VLAN access; and you want to isolate the devices from each other in case of infection or other security breach. For example, if you have a DMZ that hosts three web servers, you can isolate the web servers from each other if you apply this option to each switch port. The inside and outside networks can both communicate with all three web servers, and vice versa, but the web servers cannot communicate with each other.

b) For the **Usage Type**, click **Access**.
c) For the **Access VLAN**, click the down arrow to choose one of the existing VLAN interfaces.

You can add a new VLAN interface by clicking **Create new VLAN**. See Configure a VLAN Interface, on page 261.

Step 5 Click **OK**.

Configure Switch Ports as Trunk Ports

This procedure describes how to create a trunk port that can carry multiple VLANs using 802.1Q tagging. Trunk ports accept untagged and tagged traffic. Traffic on allowed VLANs pass through the trunk port unchanged.

When the trunk receives untagged traffic, it tags it to the native VLAN ID so that the device can forward the traffic to the correct switch ports, or can route it to another firewall interface. When the device sends native VLAN ID traffic out of the trunk port, it removes the VLAN tag. Be sure to set the same native VLAN on the trunk port on the other switch so that the untagged traffic will be tagged to the same VLAN.

Before you begin

Add a VLAN interface for each VLAN ID to which you want to assign the trunk port. See Configure a VLAN Interface, on page 261.

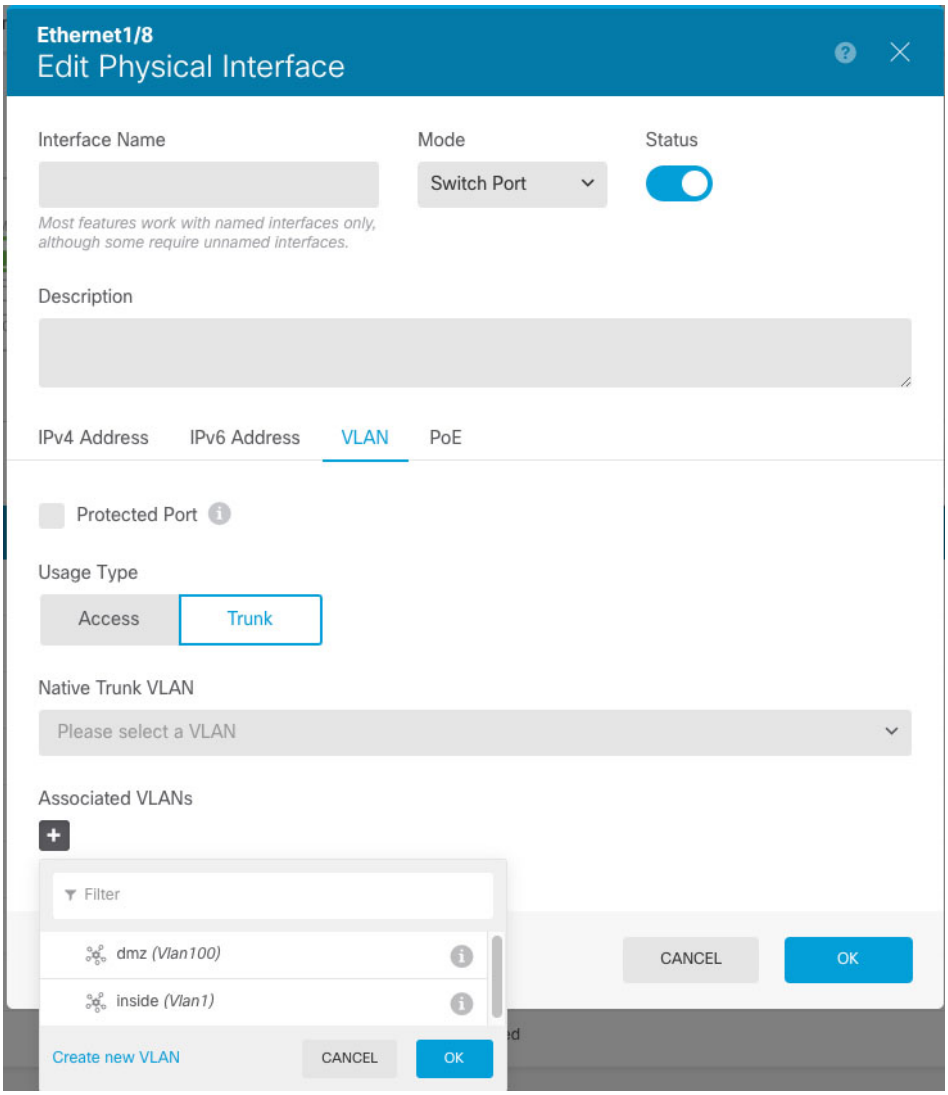
Procedure

Step 1 Click **Device**, and then click the link in the **Interfaces** summary.

The **Interfaces** page is selected by default. The interfaces list shows physical interfaces, their names, addresses, and states.

Step 2 Click the edit icon ( ) for the physical interface you want to edit.

Step 3 Set the following:



a) Do not set the **Interface Name** for switch ports; only the associated VLAN interface is a named interface.
b) Set the **Mode** to **Switch Port**.
c) Set the **Status** slider to the enabled setting ( ).
d) (Optional) Set the **Description**.

   The description can be up to 200 characters on a single line, without carriage returns.

Step 4 Click **VLAN** to set the following:

a) (Optional) Check the **Protected Port** check box to set this switch port as protected, so you can prevent the switch port from communicating with other protected switch ports on the same VLAN.

   You might want to prevent switch ports from communicating with each other if: the devices on those switch ports are primarily accessed from other VLANs; you do not need to allow intra-VLAN access; and you want to isolate the devices from each other in case of infection or other security breach. For example, if you have a DMZ that hosts three web servers, you can isolate the web servers from each other if you apply this option to each switch port. The inside and outside networks can both communicate with all three web servers, and vice versa, but the web servers cannot communicate with each other.

b) For the **Usage Type**, click **Trunk**.

c) (Optional) For the **Native Trunk VLAN**, click the down arrow to choose one of the existing VLAN interfaces for the native VLAN.

The default native VLAN ID is 1.

Each port can only have one native VLAN, but every port can have either the same or a different native VLAN.

You can add a new VLAN interface by clicking **Create new VLAN**. See Configure a VLAN Interface, on page 261.

d) For the **Associated VLANs**, click the plus icon ( ) to select one or more existing VLAN interfaces.

If you include the native VLAN in this field, it is ignored; the trunk port always removes the VLAN tagging when sending native VLAN traffic out of the port. Moreover, it will not receive traffic that still has native VLAN tagging.

You can add a new VLAN interface by clicking **Create new VLAN**. See Configure a VLAN Interface, on page 261.

Step 5 Click **OK**.

Configure Power Over Ethernet

Power over Ethernet (PoE) ports provide power for devices such as IP phones or wireless access points. PoE is enabled by default. This procedure describes how to disable and enable PoE and how to set optional parameters.

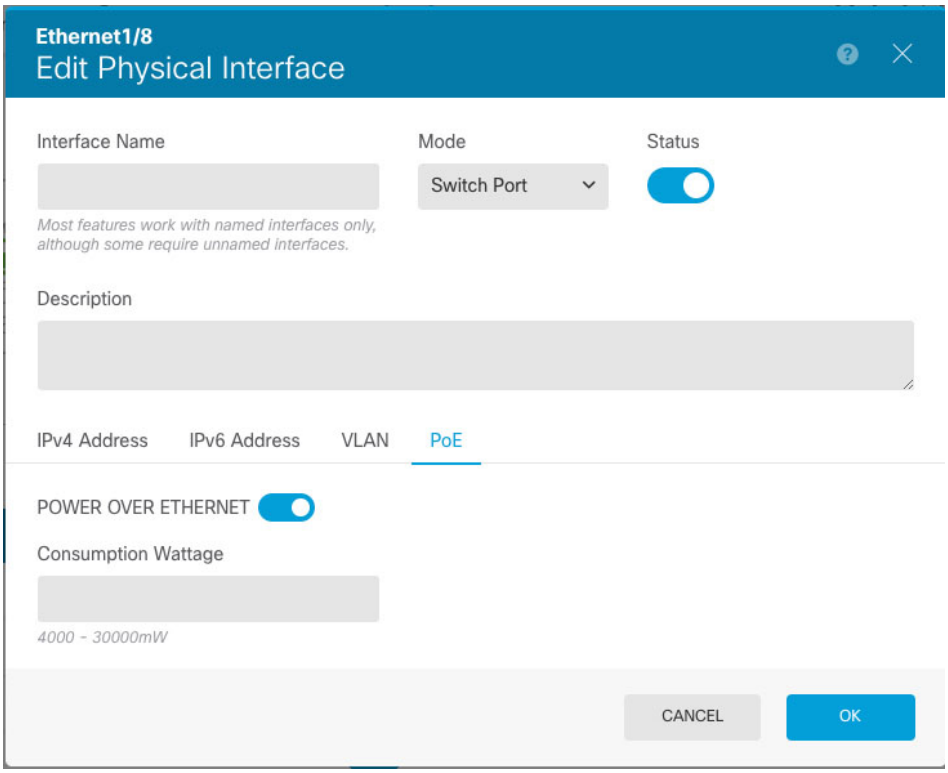
Procedure

Step 1 Click **Device**, and then click the link in the **Interfaces** summary.

The **Interfaces** page is selected by default. The interfaces list shows physical interfaces, their names, addresses, and states.

Step 2 Click the edit icon ( ) for Ethernet1/7 or 1/8 .

Step 3 Click **PoE**, and set the following:



a) To enable **POWER OVER ETHERNET**, click the slider ( ) so that it is enabled.

PoE is enabled by default.

b) (Optional) Enter the **Consumption Wattage** if you know the exact wattage you need.

To set the consumption manually, specify the wattage in milliwatts, from 4000 to 30000. Use this option if you want to set the watts manually and disable LLDP negotiation. For manual allocation, the class will show in **show power inline** output as **n/a** because the class is not used to decide the power consumption.

By default, PoE delivers power automatically to the powered device using a wattage appropriate to the class of the powered device. The firewall uses LLDP to further negotiate the correct wattage. When you connect a device of a certain class, it will be provisioned up to the maximum for that class in case it ever needs to use more power. For example, if you add a class 4 device that requests 12.95W, it will be allocated 30W even if it doesn't currently use all that power. Some devices can renegotiate power needs. If you know your device needs less power than is allocated, you can instead set the **Consumption Wattage** manually to free up the power for other devices.

Step 4 Click **OK**.

Configure VLAN Subinterfaces and 802.1Q Trunking

VLAN subinterfaces let you divide a physical 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 allow you to 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.

Create subinterfaces if you attach the physical interface to a trunk port on a switch. Create a subinterface for each VLAN that can appear on the switch trunk port. If you attach the physical interface to an access port on the switch, there is no point in creating a subinterface.

Guidelines and Limitations

- 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. Because the physical interface must be enabled for the subinterface to pass traffic, ensure that the physical interface does not pass traffic by not naming the interface. If you want to let the physical interface pass untagged packets, you can name the interface as usual.
- 1010—Subinterfaces are not supported on switch ports or VLAN interfaces.
- You cannot configure IP addresses on bridge group member interfaces, although you can modify advanced settings as needed.
- All subinterfaces on the same parent interface must be either bridge group members or routed interfaces; you cannot mix and match.
- Firewall Threat Defense 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 Firewall Threat Defense device, 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 Firewall Threat Defense device.

Procedure

Step 1 Click **Device**, and then click the link in the **Interfaces** summary.

The **Interfaces** page is selected by default. To add a subinterface to an EtherChannel, click **EtherChannel**. The interfaces list shows physical interfaces, their names, addresses, and states.

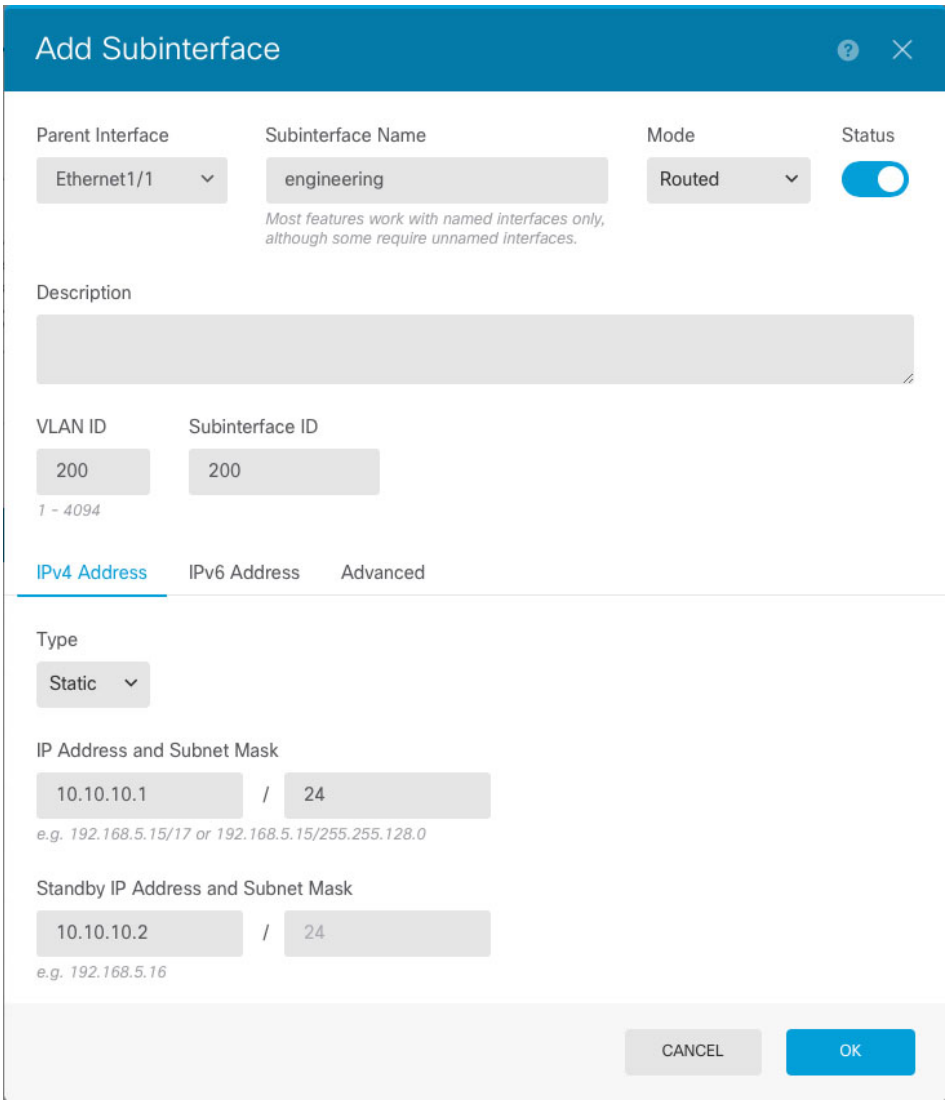
Step 2 Do one of the following:

- On the **Interfaces** page, click the plus icon ( ) to create a new subinterface.
- On the **EtherChannel** page, click the plus and down arrow icon ( ), and choose **Subinterface**.
- Click the edit icon ( ) for the subinterface you want to edit.

If you no longer need a subinterface, click the delete icon ( ) for the subinterface to delete it.

Step 3 Set the **Status** slider to the enabled setting ( ).

Step 4 Configure the parent interface, name, and description:



a) Choose the **Parent Interface**.

The parent interface is the physical interface to which you want to add the subinterface. You cannot change the parent interface after you create the subinterface.

b) Set the **Subinterface Name**, up to 48 characters.

Alphabetic characters must be lower case. For example, **inside** or **outside**. Without a name, the rest of the interface configuration is ignored.

Note
If you change the name, the change is automatically reflected everywhere you used the old name, including security zones, syslog server objects, and DHCP server definitions. However, you cannot remove the name until you first remove all configurations that use the name, because you typically cannot use an unnamed interface for any policy or setting.

c) Set the **Mode** to **Routed**.

If you later add this interface to a bridge group, then the mode will automatically change to **BridgeGroupMember**. Note that you cannot configure IP addresses on bridge group member interfaces.

d) (Optional) Set a **Description**.

The description can be up to 200 characters on a single line, without carriage returns.

e) Set the **VLAN ID**.

Enter the VLAN ID between 1 and 4094 that will be used to tag the packets on this subinterface.

f) Set the **Subinterface ID**.

Enter the subinterface ID as an integer between 1 and 4294967295. This ID is appended to the interface ID; for example Ethernet1/1.100. You can match the VLAN ID for convenience, but it is not required. You cannot change the ID after you create the subinterface.

Step 5 Click the **IPv4 Address** tab and configure the IPv4 address.

Select one of the following options from the **Type** field:

- **DHCP**—Choose this option if the address should be obtained from the DHCP server on the network. You cannot use this option if you configure high availability. Change the following options if necessary:
  - **Route Metric**—If you obtain the default route from the DHCP server, the administrative distance to the learned route, between 1 and 255. The default is 1.
  - **Obtain Default Route**—Whether to get the default route from the DHCP server. You would normally select this option, which is the default.
- **Static**—Choose this option if you want to assign an address that should not change. Type in the interface's IP address and the subnet mask for the network attached to the interface. For example, if you attach the 10.100.10.0/24 network, you could enter 10.100.10.1/24. Ensure that the address is not already used on the network.

  If you configured high availability, and you are monitoring this interface for HA, also configure a standby IP address on the same subnet. The standby address is used by this interface on the standby device. 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.

  Note
  If there is a DHCP server configured for the interface, you are shown the configuration. You can edit or delete the DHCP address pool. If you change the interface IP address to a different subnet, you must either delete the DHCP server, or configure an address pool on the new subnet, before you can save the interface changes. See Configuring the DHCP Server, on page 756.
- **PPPoE**—Choose this option if the address should be obtained using Point-to-point Protocol over Ethernet (PPPoE). PPPoE may be required if the interface is connected to a DSL modem, cable modem, or other connection to your ISP, and your ISP uses PPPoE to provide your IP address. You cannot use this option if you configure High Availability. Set the following values:
  - **Group Name**—Specify a group name of your choice to represent this connection.
  - **PPPoE Username**—Specify the username provided by your ISP.
  - **PPPoE Password**—Specify 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 Learned 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.
- **Obtain Default Route from PPPoE**—Check this check box to enable obtaining the default route from the PPPoE server.
- **IP Address Type**—Choose **Dynamic** to obtain the IP address from the PPPoE server. You can alternatively choose **Static** if you were assigned a static IP address from the ISP.

Step 6 (Optional.) Click the **IPv6 Address** tab and configure the IPv6 address.

- **State**—To enable IPv6 processing and to automatically configure the link-local address when you do not configure the global address, select **Enabled**. The link local address is generated based on the interface MAC addresses (*Modified* EUI-64 format).

  Note
  Disabling IPv6 does not disable IPv6 processing on an interface that is configured with an explicit IPv6 address or that is enabled for autoconfiguration.

- **Address Auto Configuration**—Select this option to have the address automatically configured. IPv6 stateless autoconfiguration will generate a global IPv6 address only if the link on which the device resides has a router configured to provide IPv6 services, including the advertisement of an IPv6 global prefix for use on the link. If IPv6 routing services are not available on the link, you will get a link-local IPv6 address only, which you cannot access outside of the device's immediate network link. The link local address is based on the Modified EUI-64 interface ID.

  Although RFC 4862 specifies that hosts configured for stateless autoconfiguration do not send Router Advertisement messages, the Firewall Threat Defense device does send Router Advertisement messages in this case. Select **Suppress RA** to suppress messages and conform to the RFC.

- **Static Address/Prefix**—If you do not use stateless autoconfiguration, enter the full static global IPv6 address and network prefix. For example, 2001:0DB8::BA98:0:3210/48. For more information on IPv6 addressing, see IPv6 Addressing, on page 238.

  If you want to use the address as link local only, select the **Link - Local** option. Link local addresses are not accessible outside the local network. You cannot configure a link-local address on a bridge group interface.

  Note
  A link-local address should start with FE8, FE9, FEA, or FEB, for example fe80::20d:88ff:feee:6a82. 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.

- **Standby IP Address**—If you configure high availability, and you are monitoring this interface for HA, also configure a standby IPv6 address on the same subnet. The standby address is used by this interface

on the standby device. 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.

- **Suppress RA**—Whether to suppress router advertisements. The Firewall Threat Defense can participate in router advertisements so that neighboring devices can dynamically learn a default router address. By default, router advertisement messages (ICMPv6 Type 134) are periodically sent out each IPv6 configured interface.

  Router advertisements are also 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 might want to suppress these messages on any interface for which you do not want the Firewall Threat Defense device to supply the IPv6 prefix (for example, the outside interface).

Step 7 (Optional.) Configure Advanced Options, on page 281.

The advanced settings have defaults that are appropriate for most networks. Edit them only if you are resolving network issues.

Step 8 Click **OK**.

What to do next

- Add the subinterfaces to the appropriate security zones. See Configuring Security Zones, on page 138.
- Register a fully-qualified domain name (FQDN) with your dynamic DNS service provider and configure DDNS to have the DNS server updated with the interface addresses, both IPv4 and IPv6. See Configuring Dynamic DNS, on page 759.

Configure Passive Interfaces

Passive interfaces monitor traffic flowing across a network using a switch SPAN (Switched Port Analyzer) 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 configured in a passive deployment, the system cannot take certain actions such as blocking traffic. Passive interfaces receive all traffic unconditionally and no traffic received on these interfaces is retransmitted.

You use a passive interface to monitor the traffic on the network to gather information about the traffic. For example, you can apply intrusion policies to identify the types of threats that afflict the network, or to see the URL categories for the web requests users are making. You can implement various security policies and rules to see what the system would do if deployed actively, so that it could drop traffic based on your access control and other rules.

However, because passive interfaces cannot impact traffic, there are many configuration limitations. These interfaces are merely letting the system peek at the traffic: no packets that enter a passive interface ever leave the device.

The following topics explain more about passive interfaces and how to configure them.

Why Use Passive Interfaces?

The main purpose of passive interfaces is to provide a simple demonstration mode. You can set up the switch to monitor a single source port, then use a workstation to send test traffic that is monitored by the passive interface. Thus, you can see how the Firewall Threat Defense system evaluates connections, identifies threats, and so forth. Once you are satisfied with how the system performs, you can then deploy it actively in your network and remove the passive interface configuration.

However, you can also use passive interfaces in a production environment to provide the following services:

- Pure IDS deployment—If you do not want to use the system as a firewall or IPS (intrusion prevention system), you can deploy it passively as an IDS (intrusion detection system). In this deployment method, you would use an access control rule to apply an intrusion policy to all traffic. You would also have the system monitor multiple source ports on the switch. Then, you would be able to use the dashboards to monitor the threats seen on the network. However, in this mode, the system can do nothing to prevent these threats.
- Mixed deployment—You can mix active routed interfaces with passive interfaces on the same system. Thus, you can deploy the Firewall Threat Defense device as a firewall in some networks, while configuring one or more passive interfaces to monitor traffic in other networks.

Limitations for Passive Interfaces

Any physical interface that you define as a passive mode interface has the following restrictions:

- You cannot configure subinterfaces for the passive interface.
- You cannot include the passive interface in a bridge group.
- You cannot configure IPv4 or IPv6 addresses on the passive interface.
- You cannot select the Management Only option for a passive interface.
- You can include the interface in a passive mode security zone only, you cannot include it in a routed security zone.
- You can include passive security zones in the source criteria of access control or identity rules. You cannot use passive zones in the destination criteria. You also cannot mix passive and routed zones in the same rule.
- You cannot configure management access rules (HTTPS or SSH) for a passive interface.
- You cannot use passive interfaces in NAT rules.
- You cannot configure static routes for passive interfaces. You also cannot use a passive interface in the configuration of a routing protocol.
- You cannot configure a DHCP server on a passive interface. You also cannot use a passive interface to obtain DHCP settings through auto configuration.
- You cannot use a passive interface in a syslog server configuration.
- You cannot configure any type of VPN on a passive interface.

Configure the Switch for a Hardware Firewall Threat Defense Passive Interface

A passive interface on a hardware Firewall Threat Defense device works only if you configure the network switch correctly. The following procedure is based on a Cisco Nexus 5000 series switch. If you have a different type of switch, the commands might be different.

The basic idea is to configure a SPAN (Switched Port Analyzer) or mirror port, connect the passive interface to that port, and configure a monitoring session on the switch to send copies of traffic from one or more source ports to the SPAN or mirror port.

Procedure

Step 1 Configure a port on the switch as a monitor (SPAN or mirror) port.

```
switch(config)# interface Ethernet1/48
switch(config-if)# switchport monitor
switch(config-if)#
```

Step 2 Define a monitoring session to identify the ports to monitor.

Ensure that you define the SPAN or mirror port as the destination port. In the following example, two source ports are monitored.

```
switch(config)# monitor session 1
switch(config-monitor)# source interface ethernet 1/7
switch(config-monitor)# source interface ethernet 1/8
switch(config-monitor)# destination interface ethernet 1/48
switch(config-monitor)# no shut
```

Step 3 (Optional.) Verify the configuration using **show monitor session** command.

The following example shows the brief output for session 1.

```
switch# show monitor session 1 brief
   session 1
---------------
type              : local
state             : up
source intf       :
    rx            : Eth1/7        Eth1/8
    tx            : Eth1/7        Eth1/8
    both          : Eth1/7        Eth1/8
source VSANs      :
destination ports : Eth1/48

Legend: f = forwarding enabled, l = learning enabled
```

Step 4 Physically connect the cable from the Firewall Threat Defense passive interface to the destination port on the switch.

You can configure the interface in passive mode either before or after you make the physical connection. See Configure a Physical Interface in Passive Mode, on page 278.

Configure the VLAN for a Firewall Threat Defense Virtual Passive Interface

A passive interface on a Firewall Threat Defense Virtual device works only if you configure the VLAN on the virtual network correctly. Ensure that you do the following:

- Connect the Firewall Threat Defense Virtual interface to a VLAN that you have configured in promiscuous mode. Then, configure the interface as explained in Configure a Physical Interface in Passive Mode, on page 278. The passive interface will see a copy of all traffic on the promiscuous VLAN.
- To the same VLAN, connect one or more endpoint devices, such as virtual Windows systems. You can use a single device if there is a connection from the VLAN to the Internet. Otherwise, you need at least two devices so that you can pass traffic between them. To get data for URL categories, you need to have an Internet connection.

Configure a Physical Interface in Passive Mode

You can configure an interface in passive mode. When acting passively, the interface simply monitors the traffic from the source ports in a monitoring session configured on the switch itself (for hardware devices) or on the promiscuous VLAN (for Firewall Threat Defense Virtual). For detailed information on what you need to configure in the switch or virtual network, see the following topics:

- Configure the Switch for a Hardware Firewall Threat Defense Passive Interface, on page 277
- Configure the VLAN for a Firewall Threat Defense Virtual Passive Interface, on page 278

Use passive mode when you want to analyze the traffic coming through the monitored switch ports without impacting the traffic. For an end-to-end example of using passive mode, see How to Passively Monitor the Traffic on a Network, on page 72.

Procedure

Step 1 Click **Device**, and then click the link in the **Interfaces** summary, and then click **Interfaces** or **EtherChannel**.

Step 2 Click the edit icon ( ) for the physical interface or EtherChannel you want to edit.

Pick a currently-unused interface. If you intend to convert an in-use interface to a passive interface, you need to first remove the interface from any security zone and remove all other configurations that use the interface.

Step 3 Set the **Status** slider to the enabled setting ( ).

Step 4 Configure the following:

- **Interface Name**—The name for the interface, up to 48 characters. Alphabetic characters must be lower case. For example, monitor.
- **Mode**—Select **Passive**.

- (Optional.) **Description**—The description can be up to 200 characters on a single line, without carriage returns.

Note
You cannot configure IPv4 or IPv6 addresses. On the Advanced tab, you can change the MTU, duplex, and speed settings only.

Step 5 Click **OK**.

What to do next

Creating a passive interface is not sufficient for populating the dashboards with information about the traffic seen on the interface. You must also do the following. The use case covers these steps. See How to Passively Monitor the Traffic on a Network, on page 72.

- Create a passive security zone and add the interface to it. See Configuring Security Zones, on page 138.
- Create access control rules that use the passive security zone as the source zone. Typically, you would apply intrusion policies in these rules to implement IDS (intrusion detection system) monitoring. See Configuring the Access Control Policy, on page 489.
- Optionally, create SSL decryption and identity rules for the passive security zone, and enable the Security Intelligence policy.

Configure Advanced Interface Options

Advanced options include setting the MTU, hardware settings, management only, MAC address, and other settings.

About MAC Addresses

You can manually configure Media Access Control (MAC) addresses to override the default.

For a high availability configuration, you can configure both the active and standby MAC address for an interface. 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.

Default MAC Addresses

For native instances:

Default MAC address assignments depend on the type of interface.

- Physical interfaces—The physical interface uses the burned-in MAC address.
- VLAN interfaces (Firepower 1010)—All VLAN interfaces share a MAC address. Ensure that any connected switches can support this scenario. If the connected switches require unique MAC addresses, you can manually assign MAC addresses. See Configure Advanced Options, on page 281.
- 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.

- 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 Firewall Threat Defense.

About the MTU

The MTU specifies the maximum frame *payload* size that the Firewall Threat Defense 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 Firewall Threat Defense 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.

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 UDP or ICMP, the application should take the MTU into account to avoid fragmentation.



Note The Firewall Threat Defense 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 Firewall Threat Defense 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—A jumbo frame is an Ethernet packet larger than the standard maximum of 1522 bytes (including Layer 2 header and VLAN header), up to 9216 bytes. You can set the MTU to 9000 bytes or higher to accommodate jumbo frames. The maximum depends on the model.

Note Increasing the MTU assigns more memory for jumbo frames, which might limit the maximum usage of other features, such as access rules. If you increase the MTU above the default 1500 on Firewall Threat Defense Virtual, you must reboot the system. If the device is configured for high availability, you must also reboot the standby device. You do not need to reboot other models, where jumbo frame support is always enabled.

Configure Advanced Options

Advanced interface options have default settings that are appropriate for most networks. Configure them only if you are resolving networking problems, or if you configure high availability.

The following procedure assumes the interface is already defined. You can also edit these settings while initially editing or creating the interface.

Limitations

- For bridge groups, you configure most of these options on the member interfaces. Except for DAD attempts and Enable for HA Monitoring, these options are not available for the Bridge Virtual Interface (BVI).
- You cannot set MTU, duplex, or speed for the Management interface.
- Advanced options are not available for Firepower 1010 switch ports.
- You cannot set duplex or speed for interfaces on the Firepower 4100/9300. Set these features for the interface using FXOS.
- For passive interfaces, you can set the MTU, duplex, and speed only. You cannot make the interface management only.

Procedure

Step 1 Click **Device**, click the link in the **Interfaces** summary, and then click the interfaces type to view the list of interfaces.

Step 2 Click the edit icon ( ) for the interface you want to edit.

Step 3 Click **Advanced Options**.

Step 4 Select **Enable for HA Monitoring** if you want the health of the interface to be a factor when the system decides whether to fail over to the peer unit in a high availability configuration.

This option is ignored if you do not configure high availability. It is also ignored if you do not configure a name for the interface.

Step 5 To make a data interface management only, select **Management Only**.

A management only interface does not allow through traffic, so there is very little value in setting a data interface as management only. You cannot change this setting for the Management/Diagnostic interface, which is always management only.

Step 6 Change the **MTU** (maximum transmission unit) to the desired value.

The default MTU is 1500 bytes. The minimum and maximum depend on your platform. Set a high value if you typically see jumbo frames on your network.

Note
If you increase MTU above 1500 on the following devices, you must reboot the device: ISA 3000 series devices,Firewall Threat Defense Virtual. If the device is configured for high availability, you must also reboot the standby device. You do not need to reboot other models, where jumbo frame support is always enabled.

Step 7 (Physical interface only.) Modify the speed and duplex settings.

The default is that the interface negotiates the best duplex and speed with the interface at the other end of the wire, but you can force a specific duplex or speed if necessary. The options listed are only those supported by the interface. Before setting these options for interfaces on a network module, please read Limitations for Interface Configuration, on page 238.

- **Duplex**—Choose **Half** or **Full**. SFP interfaces only support **Full** duplex.
- **Speed**—The exact options depend on the model and interface type. Choose a speed, **Auto**, **No Negotiate**, or **Detect SFP**. For the Firepower 1100 or 2100 fiber ports, **No Negotiate** sets the speed to 1000 Mbps and disables link negotiation for flow-control parameters and remote fault information. (Secure Firewall 3100 only) Choose **Detect SFP** to detect the speed of the installed SFP module and use the appropriate speed. Duplex is always Full, and auto-negotiation is always enabled. This option is useful if you later change the network module to a different model, and want the speed to update automatically.
- (Secure Firewall 3100 only) **Auto Negotiation**— Depending on the type of interface, set the interface to negotiate the link status for flow-control parameters and remote fault information.
- **Forward Error Correction Mode**—(Secure Firewall 3100 only) For 25 Gbps and higher interfaces, enable Forward Error Correction (FEC). For an EtherChannel member interface, you must configure Forward Error Correction before you add it to the EtherChannel. The setting chosen when you use **Auto** depends on the transceiver type and whether the interface is fixed (built-in) or on a network module.

Table 8: Default FEC for Auto Setting

Transceiver Type	Fixed Port Default FEC (Ethernet 1/9 through 1/16)	Network Module Default FEC
25G-SR	Clause 74 FC-FEC	Clause 108 RS-FEC
25G-LR	Clause 74 FC-FEC	Clause 108 RS-FEC
10/25G-CSR	Clause 74 FC-FEC	Clause 74 FC-FEC
25G-AOC*x*M	Clause 74 FC-FEC	Clause 74 FC-FEC
25G-CU2.5/3M	Auto-Negotiate	Auto-Negotiate
25G-CU4/5M	Auto-Negotiate	Auto-Negotiate

Step 8 Modify the **IPv6 Configuration** settings.

- **Enable DHCP for IPv6 address configuration**—Whether to set the Managed Address Configuration flag in the IPv6 router advertisement packet. This flag informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain addresses, in addition to the derived stateless autoconfiguration address.
- **Enable DHCP for IPv6 non-address configuration**—Whether to set the Other Address Configuration flag in the IPv6 router advertisement packet. This flag informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain additional information from DHCPv6, such as the DNS server address.
- **DAD Attempts**—How often the interface performs Duplicate Address Detection (DAD), from 0 - 600. The default is 1. During the stateless autoconfiguration process, DAD verifies the uniqueness of new unicast IPv6 addresses before the addresses are assigned to interfaces. 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. The interface uses neighbor solicitation messages to perform Duplicate Address Detection. Set the value to 0 to disable duplicate address detection (DAD) processing.

Step 9 (Optional, recommended for subinterfaces and high availability units.) Configure the MAC address.

By default, the system uses the MAC address burned into the network interface card (NIC) for the interface. Thus, all subinterfaces on an interface use the same MAC address, so you might want to create unique addresses per subinterface. Manually configured active/standby MAC addresses are also recommended if you configure high availability. Defining the MAC addresses helps maintain consistency in the network in the event of failover.

- **MAC Address**—The Media Access Control in H.H.H format, where H is a 16-bit hexadecimal digit. For example, you would enter the MAC address 00-0C-F1-42-4C-DE 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.)
- **Standby 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 10 Click **OK**.

Scan for Interface Changes, and Migrate an Interface

When you change interfaces on the device, the device informs the Firewall Device Manager that a change has occurred. You will not be able to deploy your configuration until you perform an interface scan.The Firewall Device Manager supports migrating an interface in your security policy with another interface, so removing an interface can be almost seamless.

About Interface Scanning and Migrating

Scanning

When you change interfaces on the device, the device informs the Firewall Device Manager that a change has occurred. You will not be able to deploy your configuration until you perform an interface scan. After a scan, which detects any added, removed, or restored interfaces, you can deploy your configuration; however, the parts of the configuration that refer to removed interfaces will not be deployed.

Interface changes that require a scan include adding or removing interfaces. For example: network module change; allocated interface change on the Firepower 4100/9300 chassis; interface change on the Firewall Threat Defense Virtual.

The following changes do not block deployment after a scan:

- Security zone membership
- EtherChannel interface membership
- Firepower 1010 VLAN interface switch port membership
- Bridge group interface membership, for policies that refer to the BVI



Note A syslog server egress interface change will not block deployment, although you should fix the syslog server configuration, either manually or using the interface replacement feature.

Migrating

Adding a new interface, or removing an unused interface, has minimal impact on the Firewall Threat Defense configuration. However, removing an interface that is used in your security policy will impact the configuration. Interfaces can be referenced directly in many places in the Firewall Threat Defense configuration, including security zones, NAT, VPN, routing, DHCP server, and so on.

Firewall Device Manager supports migrating an interface in your security policy to another interface, so removing an interface can be almost seamless.



Note The migrate feature does *not* copy the name, IP address, and other configuration from one interface to another; rather, this feature changes the security policy to refer to the new interface instead of the old interface. You need to manually configure the new interface settings before migrating.

If you need to remove an interface, we recommend that you add the new interface and migrate the old interface *before* you remove it. If you add and remove interfaces at the same time, the migration process will still work; however, you cannot *manually* edit removed interfaces or policies referring to them, so you may find it easier to perform the migration in stages.

If you replace an interface with the same type (for example, you need to RMA a network module), you can: 1. Remove the old module from the chassis; 2. Perform a scan; 3. Deploy changes unrelated to the removed interfaces; 4. Replace the module; 5. Perform a new scan; 6. Deploy your configuration, including any changes related to the interfaces. You do not need to perform a migration if the new interface has the same interface ID and characteristics as the old interface.

Guidelines and Limitations for Interface Scanning and Migrating

Unsupported Interface Migrations

- Physical interface to BVI
- Passive interface to Firewall interface

- Bridge group members
- EtherChannel interface members
- ISA 3000 Hardware Bypass members
- Firepower 1010 VLAN interfaces or switch ports
- Diagnostic interface
- HA failover and state links
- Migration of interfaces of different types, for example migrating a bridge group interface to a feature that requires a physical interface

Additional Guidelines

- If you need to remove an interface, we recommend that you add the new interface and migrate the old interface *before* you remove it.
- For Firewall Threat Defense Virtual, only add and remove interfaces at the end of the interface list. If you add or remove an interface anywhere else, then the hypervisor will renumber your interfaces, causing the interface IDs in your configuration to line up with the wrong interfaces.
- If a scan/migration goes bad, restore the original interfaces on the chassis, and re-scan to get back to the original state.
- For backups, be sure to create a new backup with the new interfaces. A restore with the old configuration will restore the old interface information, and you will have to perform the scan/replace again.
- For HA, make the same interface changes on both units before you perform the interface scan on the active unit. You only need to perform the scan/migration on the active unit. Configuration changes are replicated to the standby unit.

Scan and Migrate Interfaces

Scan for interface changes in the Firewall Device Manager, and migrate interface configurations from removed interfaces. If you only want to migrate an interface configuration (and no scan is required), ignore the steps in the following procedure related to scanning.



Note The migrate feature does *not* copy the name, IP address, and other configuration from one interface to another; rather, this feature changes the security policy to refer to the new interface instead of the old interface. You need to manually configure the new interface settings before migrating.

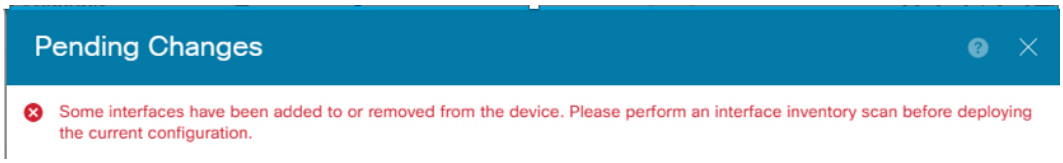
Procedure

Step 1 Add or remove interfaces on the chassis.

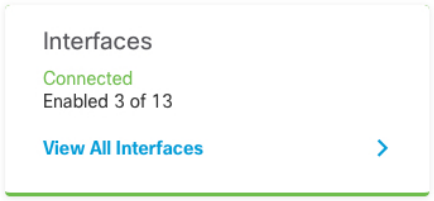
If you need to remove an interface, we recommend that you add the new interface and perform a replacement of the old interface *before* you remove it.

Step 2 Scan for interface changes.

You will not be able to deploy your configuration until you perform an interface scan. If you try to deploy before a scan, you see the following error:

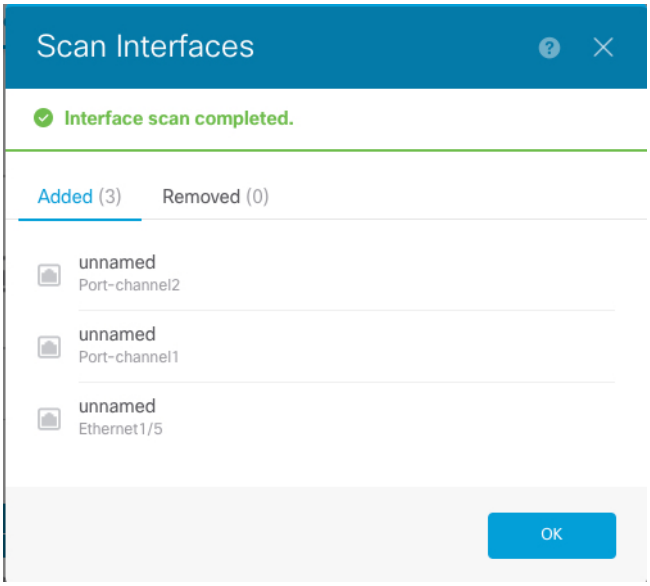


a) Click **Device**, then click the **View All Interfaces** link in the **Interfaces** summary.

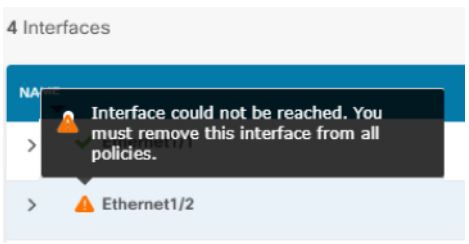


b) Click the Scan Interfaces icon ( ).

c) Wait for the interfaces to scan, and then click **OK**.



After the scan, removed interfaces show on the Interfaces page with caution symbols:



Step 3 To migrate an existing interface to a new one:

a) Configure the new interface with a name, IP address, and so on.

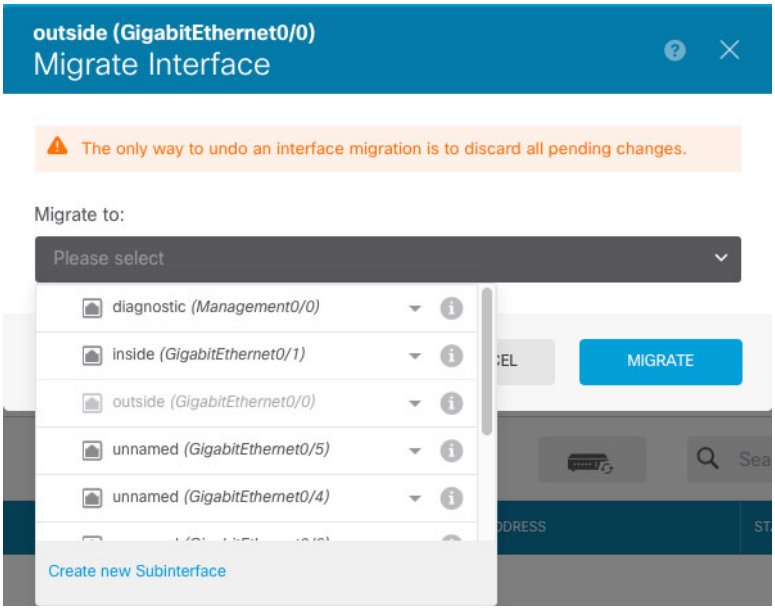
If you want to use the existing IP address and name of an interface that you want to remove, then you need to first reconfigure the old interface with a dummy name and IP address so that you can use those settings on the new interface.

b) Click the Migrate icon for the old interface.

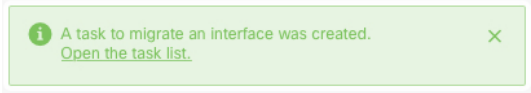


This process migrates the old interface to the new interface in all configuration settings that refer to the interface.

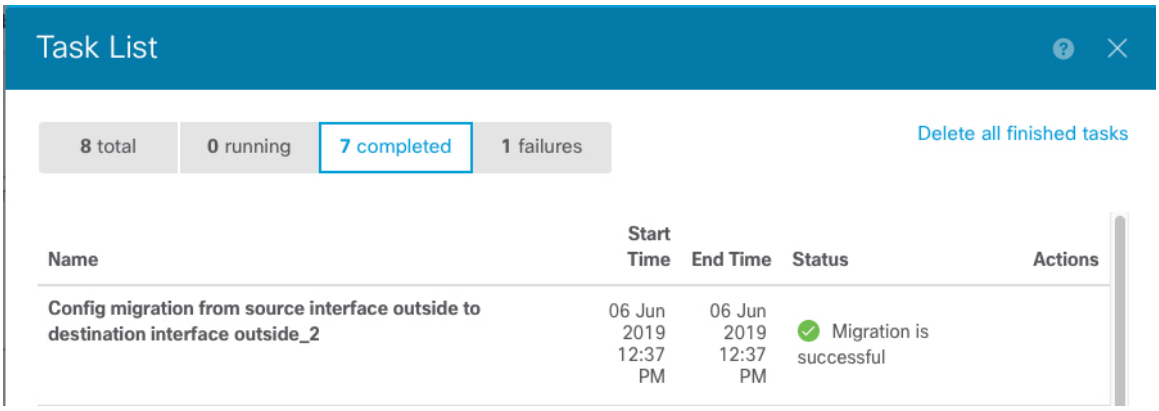
c) Choose the new interface from the **Migrate to:** drop-down list.



d) A message appears on the **Interfaces** page. Click the link in the message.



e) Check the **Task List** to ensure that the migration was successful.

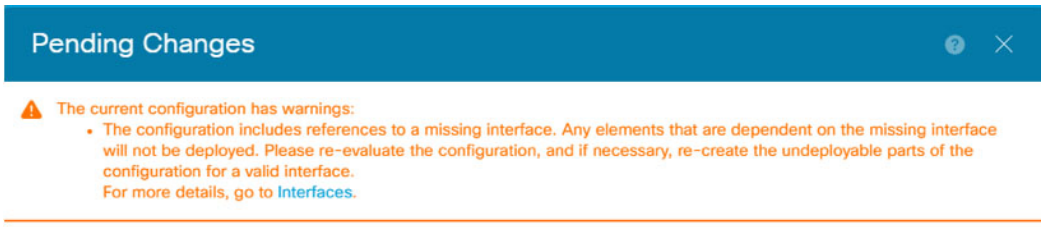


f) If the migration fails, you can view the reasons in the API Explorer.

To open API Explorer, click the more options button (⋮) and choose **API Explorer**. Choose **Interface** > **GET /jobs/interfacemigrations**, and then click **Try it Out!**.

Step 4 Deploy your configuration.

The parts of the configuration that refer to removed inerfaces will not be deployed, in which case you will see the following message:



Step 5 Remove the old interfaces on the chassis, and perform another scan.

Removed interfaces that are no longer used in your policy will be removed from the **Interfaces** page.

Step 6 Redeploy your configuration to remove the unused interfaces from your configuration.

Manage the Network Module for the Secure Firewall 3100

If you install a network module before you first power on the firewall, no action is required; the network module is enabled and ready for use.

If you need to make changes to your network module installation after initial bootup, then see the following procedures.

Add a Network Module

To add a network module to a firewall after initial bootup, perform the following steps. Adding a new module requires a reboot.

Procedure

Step 1 Install the network module according to the hardware installation guide.

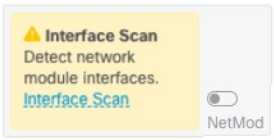
For High Availability, install the network module on both units.

Step 2 Reboot the firewall; see Rebooting or Shutting Down the System, on page 815. For High Availability, reboot the standby unit, and then perform the rest of this procedure on the standby unit.

Step 3 Click **Device**, then click the **View All Interfaces** link in the **Interfaces** summary.

The graphic shows that an interface scan is required.

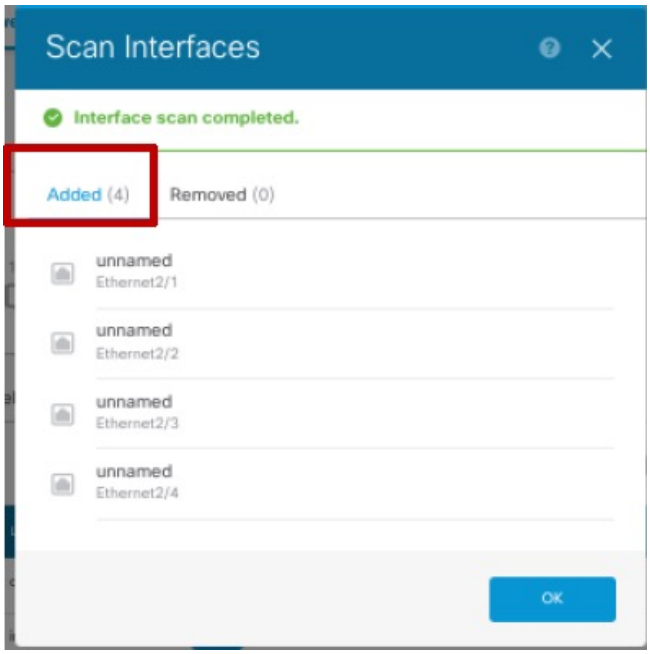
Figure 9: Interface Scan Required



Step 4 Click **Interface Scan** to update the page with the new network module details.

Wait for the interfaces to scan, and then click **OK**.

Figure 10: Scan Interfaces



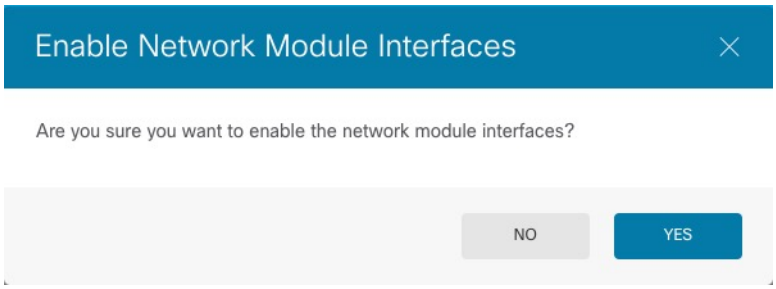
Step 5 On the interfaces graphic, click the slider ( ) to enable the network module.

Figure 11: Enable Network Module



Step 6 You are prompted to confirm that you want to enable the network module. Click **Yes**.

Figure 12: Confirm Enable



Step 7 For High Availability, change the active unit (see Switching the Active and Standby Peers (Forcing Failover), on page 219), and then perform the above steps for the new standby unit.

Hot Swap the Network Module

You can hot swap a network module for a new module of the same type without having to reboot. However, you must shut down the current module to remove it safely. This procedure describes how to shut down the old module, install a new module, and enable it.

Before you begin

For High Availability, you cannot disable a network module if the failover link is on the module. You will have to break High Availability (see Breaking High Availability, on page 217). After you hot swap the module, you can reform High Availability.

Procedure

Step 1 For High Availability, ensure the unit you want to perform the hot swap on is the standby node. See Switching the Active and Standby Peers (Forcing Failover), on page 219.

Step 2 Click **Device**, then click the **View All Interfaces** link in the **Interfaces** summary.

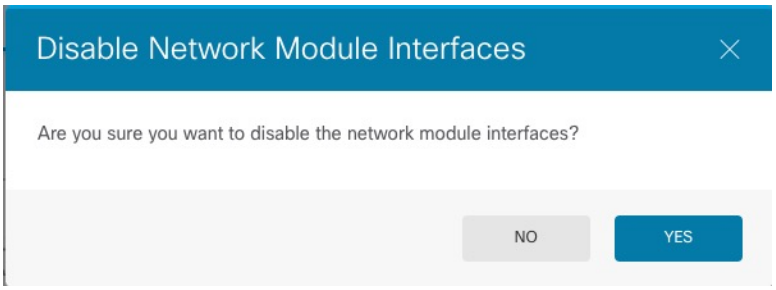
Step 3 On the interfaces graphic, click the slider ( ) to disable the network module.

Figure 13: Disable Network Module



Step 4 You are prompted to confirm that you want to disable the network module. Click **Yes**.

Figure 14: Confirm Disable



Step 5 Install the network module according to the hardware installation guide.

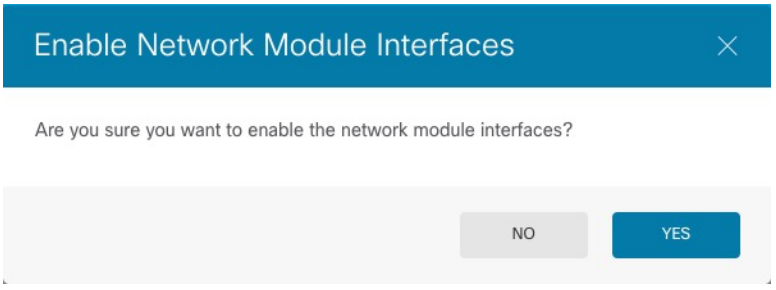
Step 6 On the interfaces graphic, click the slider ( ) to enable the network module.

Figure 15: Enable Network Module



Step 7 You are prompted to confirm that you want to enable the network module. Click **Yes**.

Figure 16: Confirm Enable



Replace the Network Module with a Different Type

If you replace a network module with a different type, then a reboot is required. If the new module has fewer interfaces than the old module, you will have to manually remove any configuration related to interfaces that will no longer be present.

Before you begin

For high availability, you cannot disable a network module if the failover link is on the module. You will have to break high availability (see Breaking High Availability, on page 217), which means you will have downtime when you reboot the active unit. After the units finish rebooting, you can reform high availability.

Procedure

Step 1 Click **Device**, then click the **View All Interfaces** link in the **Interfaces** summary. For high availability, perform this procedure on the standby unit first.

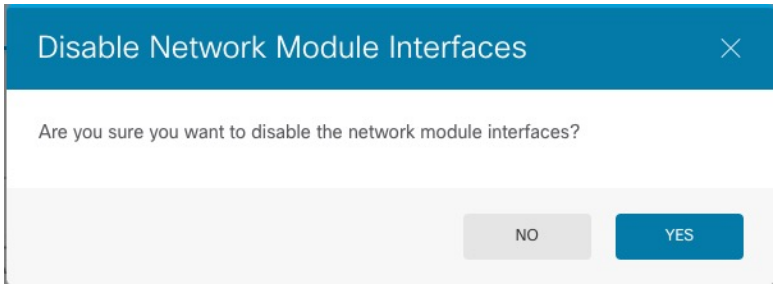
Step 2 On the interfaces graphic, click the slider ( ) to disable the network module.

Figure 17: Disable Network Module



Step 3 You are prompted to confirm that you want to disable the network module. Click **Yes**.

Figure 18: Confirm Disable



Step 4 On the device, remove the old network module and replace it with the new network module according to the hardware installation guide.

Step 5 Reboot the firewall; see Rebooting or Shutting Down the System, on page 815.

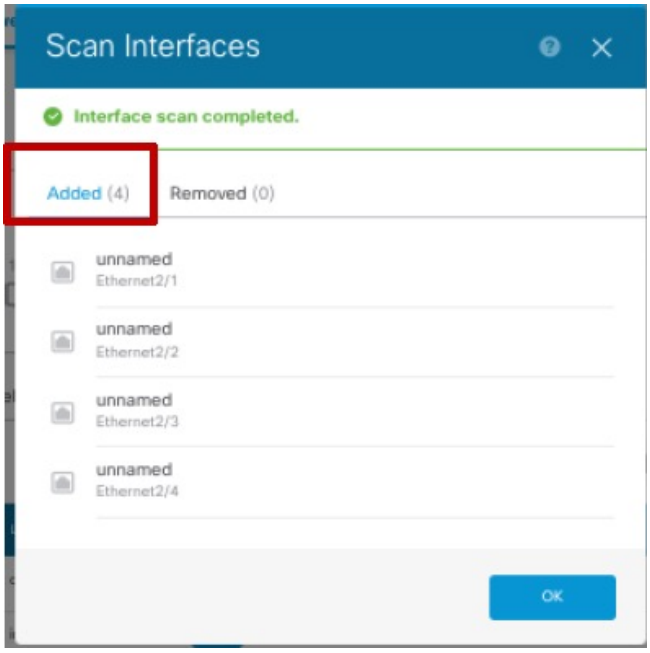
Step 6 On the **Interfaces** page, the graphic shows that an interface scan is required. Click **Interface Scan** to update the page with the new network module details.

Figure 19: Interface Scan Required



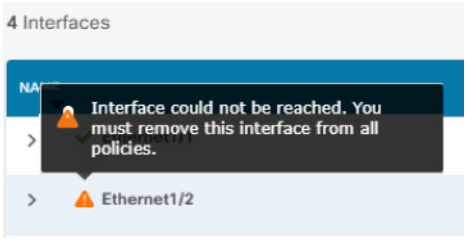
Step 7 Wait for the interfaces to scan, and then click **OK**.

Figure 20: Scan Interfaces



After the scan, removed interfaces show on the **Interfaces** page with caution symbols:

Figure 21: Removed Interfaces



Step 8 If the network module has *fewer* interfaces, you need to remove any configuration that directly refers to the removed interfaces.

Policies that refer to security zones are not affected. You can optionally migrate the configuration to a different interface. See Scan and Migrate Interfaces, on page 285.

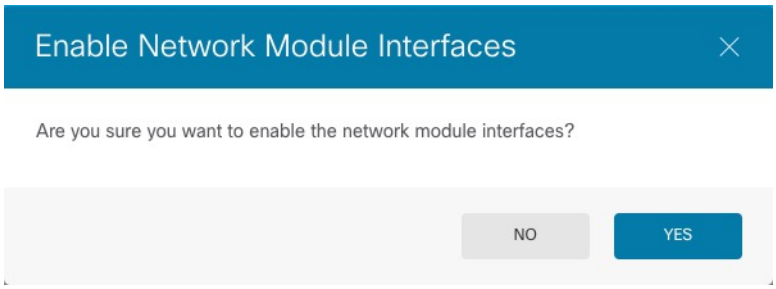
Step 9 On the interfaces graphic, click the slider ( ) to enable the network module.

Figure 22: Enable Network Module



Step 10 You are prompted to confirm that you want to enable the network module. Click **Yes**.

Figure 23: Confirm Enable



Step 11 To change the interface speed, see Configure Advanced Options, on page 281.

The default speed is set to Detect SFP, which detects the correct speed from the SFP installed. You only need to fix the speed if you manually set the speed to a particular value and you now need a new speed.

Step 12 If you had to change any configuration, click the **Deployment** icon.

You do not need to deploy just to save the network module changes.

Step 13 For high availability, change the active unit (see Switching the Active and Standby Peers (Forcing Failover), on page 219), and then perform the above steps for the new standby unit.

Remove the Network Module

If you want to permanently remove the network module, follow these steps. Removing a network module requires a reboot.

Before you begin

For High Availability, make sure the failover link is not on the network module.

Procedure

Step 1 Click **Device**, then click the **View All Interfaces** link in the **Interfaces** summary. For High Availability, perform this procedure on the standby unit first.

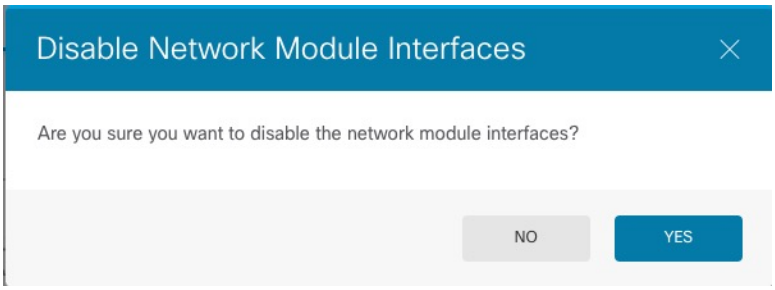
Step 2 On the interfaces graphic, click the slider ( ) to disable the network module.

Figure 24: Disable Network Module



Step 3 You are prompted to confirm that you want to disable the network module. Click **Yes**.

Figure 25: Confirm Disable

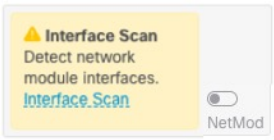


Step 4 On the firewall, remove the network module.

Step 5 Reboot the firewall; see Rebooting or Shutting Down the System, on page 815.

Step 6 On the **Interfaces** page, the graphic shows that an interface scan is required. Click **Interface Scan** to update the page with the correct network module details.

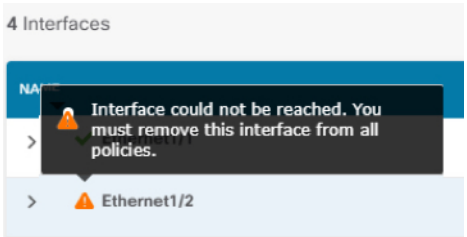
Figure 26: Interface Scan Required



Step 7 Wait for the interfaces to scan, and then click **OK**.

After the scan, removed interfaces show on the **Interfaces** page with caution symbols:

Figure 27: Removed Interfaces



Step 8 You need to remove any configuration that directly refers to the removed interfaces.

Policies that refer to security zones are not affected. You can optionally migrate the configuration to a different interface. See Scan and Migrate Interfaces, on page 285.

Step 9 If you had to change any configuration, click the **Deployment** icon.

You do not need to deploy just to save the network module changes.

Step 10 For High Availability, change the active unit (see Switching the Active and Standby Peers (Forcing Failover), on page 219), and then perform the above steps for the new standby unit.

Configure Hardware Bypass for Power Failure (ISA 3000)

You can enable hardware bypass so that traffic continues to flow between an interface pair during a power outage. Supported interface pairs are copper interfaces GigabitEthernet 1/1 and 1/2; and GigabitEthernet 1/3 and 1/4. If you have a fiber Ethernet model, only the copper Ethernet pair (GigabitEthernet 1/1 and 1/2) supports hardware bypass. By default, hardware bypass is enabled for both interface pairs if supported.

When hardware bypass is active, traffic passes between these interface pairs at layer 1. Both the Firewall Device Manager and the Firewall Threat Defense CLI will see the interfaces as being down. No firewall functions are in place, so make sure you understand the risks of allowing traffic to pass through the device.

We suggest that you disable TCP sequence number randomization (as described in this procedure). By default, the ISA 3000 rewrites the initial sequence number (ISN) of TCP connections passing through it to a random number. When hardware bypass is activated, the ISA 3000 is no longer in the data path and does not translate the sequence numbers. The receiving client receives an unexpected sequence number and drops the connection, so the TCP session needs to be re-established. Even with TCP sequence number randomization disabled, some TCP connections will have to be re-established because of the link that is temporarily down during the switchover.

In CLI Console or an SSH session, use the **show hardware-bypass** command to monitor the operational status.

Before you begin

For hardware bypass to work:

- You must place the interface pairs in the same bridge group.
- You must attach the interfaces to access ports on the switch. Do not attach them to trunk ports.

Procedure

Step 1 Click **Device**, then click the link in the Interfaces summary.

The **Hardware Bypass** section at the top of the page shows the current configuration on the allowed interface pairs for this device.

However, you must ensure the pairs are configured in the same bridge group before you can enable hardware bypass.

Step 2 Click **Edit** to configure hardware bypass.

The **Hardware Bypass Configuration** dialog box appears.

Step 3 To configure automatic hardware bypass behavior, for each interface pair choose one of the following options in the **Hardware Bypass during Power Down** area.

- **Disable**—Disables hardware bypass. Traffic will not pass through the device during a power outage.
- **Enable**—Activates hardware bypass during a power outage. Hardware bypass ensures that traffic is not interrupted during a power outage. Note that bypassed traffic is not inspected, and security policies are not applied. After power is restored, hardware bypass is automatically disabled so traffic can flow through normally, with inspection. Note that there may be a brief traffic interruption when hardware bypass is disabled.
- **Enable with Persistence**—Activates hardware bypass during a power outage, and keeps it enabled after power restoration. After power is restored, you must disable hardware bypass using the **Manual Hardware Bypass** slider. This option lets you control when the brief interruption in traffic occurs.

Step 4 (Optional.) To manually enable or disable hardware bypass, click the **Manual Hardware Bypass** slider.

For example, you might want to test the system, or temporarily bypass the device for some reason. You also need to manually disable hardware bypass when you use the **Enable with Persistence** option.Note that you must deploy the configuration to change the state of hardware bypass; simply changing the settings is not sufficient.

When you manually enable/disable hardware bypass, you will see the following syslog messages, where *pair* is 1/1-1/2 or 1/3-1/4.

- %FTD-6-803002: no protection will be provided by the system for traffic over GigabitEthernet *pair*
- %FTD-6-803003: User disabled bypass manually on GigabitEthernet *pair*

Step 5 Click **OK**.

The change is not immediate. You must deploy the configuration.

Step 6 (Optional.) Create the FlexConfig object and policy needed to disable TCP sequence number randomization.

a) Click **View Configuration** in **Device** > **Advanced Configuration**.
b) Click **FlexConfig** > **FlexConfig Objects** in the Advanced Configuration table of contents.
c) Click the + button to create a new object.
d) Enter a name for the object. For example, **Disable_TCP_Randomization**.
e) In the **Template** editor, enter the commands to disable TCP sequence number randomization.

The command is **set connection random-sequence-number disable**, but you must configure it for a specific class within a policy map. By far, the easiest approach is to disable random sequence numbers globally, which requires the following commands:

```
policy-map global_policy
 class default_class
  set connection random-sequence-number disable
```

f) In the **Negate Template** editor, enter the lines required to undo this configuration.

For example, if you disable TCP sequence number randomization globally, the negate template would be the following:

```
policy-map global_policy
 class default_class
  set connection random-sequence-number enable
```

g) Click **OK** to save the object.

You now need to add the object to the FlexConfig policy. Creating the object is not enough.

h) Click **FlexConfig Policy** in the table of contents.
i) Click + in the Group List.
j) Select the Disable_TCP_Randomization object and click **OK**.

The preview should update with the commands in the template. Verify you are seeing the expected commands.

k) Click **Save**.

You can now deploy the policy.

Monitoring Interfaces

You can view some basic information about interfaces in the following areas:

- **Device**. Use the port graphic to monitor the current state of the interfaces. Mouse over a port to see its IP addresses, EtherChannel membership, and enabled and link statuses. The IP addresses can be statically assigned or obtained using DHCP.

  Interface ports use the following color coding:

  - Green—The interface is configured, enabled, and the link is up.
  - Gray—The interface is not enabled.
  - Orange/Red—The interface is configured and enabled, but the link is down. If the interface is wired, this is an error condition that needs correction. If the interface is not wired, this is the expected status.

- **Monitoring** > **System**. The **Throughput** dashboard shows information on traffic flowing through the system. You can view information on all interfaces, or you can select a specific interface to examine.
- **Monitoring** > **Zones**. This dashboard shows statistics based on security zones, which are composed of interfaces. You can drill into this information for more detail.

Monitoring Interfaces in the CLI

You can also open the CLI console or log into the device CLI and use the following commands to get more detailed information about interface-related behavior and statistics.

- **show interface** displays interface statistics and configuration information. This command has many keywords you can use to get to the information you need. Use ? as a keyword to see the available options.
- **show ipv6 interface** displays IPv6 configuration information about the interfaces.
- **show bridge-group** displays information about Bridge Virtual Interfaces (BVI), including member information and IP addresses.
- **show conn** displays information about the connections currently established through the interfaces.

- **show traffic** displays statistics about traffic flowing through each interface.
- **show ipv6 traffic** displays statistics about IPv6 traffic flowing through the device.
- **show dhcpd** displays statistics and other information about DHCP usage on the interfaces, particularly about the DHCP servers configured on interfaces.
- **show switch vlan** displays the VLAN–to–switch port association.
- **show switch mac-address-table** shows the static and dynamic MAC address entries.
- **show arp** shows dynamic, static, and proxy ARP entries.
- **show power inline** shows PoE status.
- **show vpdn group** shows PPPoE groups and the usernames and authentication configured.
- **show vpdn username** shows PPPoE usernames and passwords.
- **show vpdn session pppoe state** shows the status of the PPPoE session.

Examples for Interfaces

The use case chapter includes the following interface-related examples:

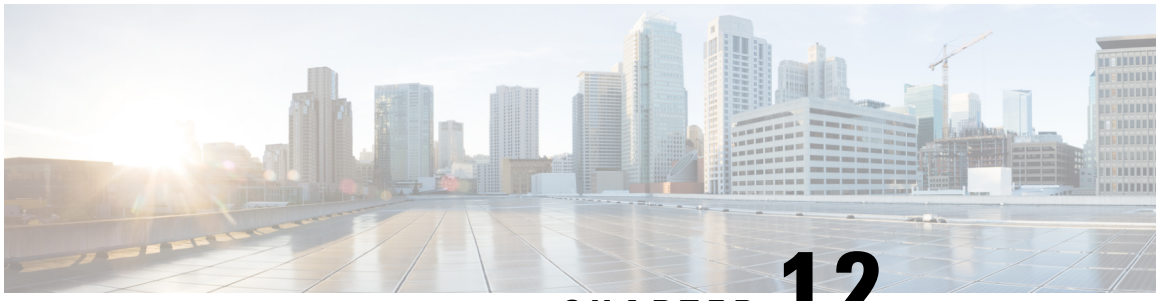
- How to Configure the Device in Firewall Device Manager, on page 39
- How to Add a Subnet, on page 67
- How to Passively Monitor the Traffic on a Network, on page 72



PART IV

Routing

- Routing Basics and Static Routes, on page 303
- Virtual Routers, on page 319
- Route Maps and Other Objects for Route Tuning, on page 345
- Open Shortest Path First (OSPF), on page 363
- Enhanced Interior Gateway Routing Protocol (EIGRP), on page 383
- Border Gateway Protocol (BGP), on page 401



CHAPTER 12

Routing Basics and Static Routes

The system uses a routing table to determine the egress interface for packets entering the system. The following topics explain routing basics and how to configure static routing on the device.

- Best Practices for Routing, on page 303
- Routing Overview, on page 303
- Static Routes, on page 309
- Monitoring Routing, on page 317

Best Practices for Routing

Designing the routing processes in a network can be a complex process. This chapter assumes that you are configuring the Firewall Threat Defense device to work within an existing network and to participate in the routing processes that you have already established in the network.

If you are instead creating a new network, please take the time to read elsewhere about routing protocols and how to design an effective routing plan that works for your network. This chapter does not go into recommendations for choosing protocols, nor does it go into depth on how the protocols work.

If your network is very small, and you are simply linking up to an ISP, you might need nothing more than a few static routes, and not need to implement routing protocols at all.

But if you are setting up a large network that will include many routers, you probably need to implement at least one routing protocol for interior routing, such as OSPF, and possibly one for external routing, such as BGP. Your service provider can help you understand what external routing might be needed, if any. If this is your situation, first understand which routing protocols you can configure using Firewall Threat Defense, then plan your network, and finally configure the Firewall Threat Defense device according to your plan.

Routing Overview

The following topics describe how routing behaves within the Firewall Threat Defense device. Routing is the act of moving information across a network from a source to a destination. Along the way, at least one intermediate node is typically encountered. Routing involves two basic activities: determining optimal routing paths and transporting packets through a network.

Supported Routing Protocols

The following table explains the routing protocols and technologies you can configure on a Firewall Threat Defense device using the Firewall Device Manager, and the method you need to use to complete the configuration.

Table 9: Supported Routing Protocols

Routing Feature	Configuration Method	Notes
BGP	Smart CLI	Configure BGP Smart CLI objects from the **Device** > **Routing** page. Configure objects used in BGP, such as route maps, using Smart CLI objects from the **Device** > **Advanced Configuration** page.
Bi-directional forwarding detection (BFD)	FlexConfig	Configure BFD using FlexConfig objects from the **Device** > **Advanced Configuration** page. BFD is supported with BGP only.
EIGRP	Smart CLI	Configure EIGRP Smart CLI objects from the **Device** > **Routing** page. Configure objects used in EIGRP, such as route maps, using Smart CLI objects from the **Device** > **Advanced Configuration** page.
IS-IS	FlexConfig	Configure IS-IS using FlexConfig objects from the **Device** > **Advanced Configuration** page.
Multicast routing	FlexConfig	Configure multicast routing using FlexConfig objects from the **Device** > **Advanced Configuration** page.
OSPFv2	Smart CLI	Configure OSPFv2 Smart CLI objects from the **Device** > **Routing** page. Configure objects used in OSPFv2, such as route maps, using Smart CLI objects from the **Device** > **Advanced Configuration** page.
OSPFv3	—	OSPFv3 configuration is not supported.
Policy-based routing (PBR)	FlexConfig	Configure policy-based routing (PBR) using FlexConfig objects from the **Device** > **Advanced Configuration** page.
RIP	FlexConfig	Configure RIP using FlexConfig objects from the **Device** > **Advanced Configuration** page.
Static routes	Firewall Device Manager	Configure static routes globally or per virtual router from the **Device** > **Routing** page.
Virtual routers, VRF	Firewall Device Manager	Configure virtual routers from the **Device** > **Routing** page.

Route Types

There are two main types of route: static or dynamic.

Static routes are those that you define explicitly. These are stable, normally high-priority routes, that you would use to ensure traffic to the route destination is always sent out the correct interface. For example, you would create a default static route to cover all traffic not already covered by any other route, that is, 0.0.0.0/0 for IPv4 or ::/0 for IPv6. Another example would be a static route to an internal syslog server that you always want to use.

Dynamic routes are those learned through the operation of a routing protocol, such as OSPF, BGP, EIGRP, IS-IS, or RIP. You do not define the routes directly. Instead, you configure the routing protocol, and the system then communicates with neighbor routers, transmitting routing updates to them and receiving routing updates in turn.

Dynamic routing protocols adjust the routing table to changing network circumstances by analyzing incoming routing update messages. If the message indicates that a network change has occurred, the system 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.

Static routing is simple and serves the purpose of basic routing. It works well in environments where network traffic is relatively predictable and where network design is relatively simple. However, because static routes cannot change unless you edit them, they cannot react to changes in the network.

Unless you have a small network, you would typically combine static routes with one or more dynamic routing protocol. You would define at least one static route, as the default route for any traffic that does not match an explicit route.

The Routing Table and Route Selection

When NAT translations (xlates) and rules do not determine the egress interface, the system uses the routing table to determine the path for a packet.

Routes in the routing table include a metric called “administrative distance” that provides a relative priority to a given route. If a packet matches more than one route entry, the one with the lowest distance is used. Directly connected networks (those defined on an interface) have the distance 0, so they are always preferred. Static routes have a default distance of 1, but you can create them with any distance between 1-254.

Routes that identify a specific destination take precedence over the default route (the route whose destination is 0.0.0.0/0 or ::/0).

How the Routing Table Is Populated

The Firewall Threat Defense routing table can be populated by statically defined routes, directly connected routes, and routes discovered by the dynamic routing protocols. Because the Firewall Threat Defense 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 Firewall Threat Defense 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 Firewall Threat Defense 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 Firewall Threat Defense 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 Firewall Threat Defense.

Table 10: 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

Route Source	Default Administrative Distance
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 Firewall Threat Defense 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 Firewall Threat Defense 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 Firewall Threat Defense 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 Firewall Threat Defense 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 Firewall Threat Defense. 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.

Routing Table for Management Traffic

As a standard security practice, it is often necessary to segregate and isolate management (from-the-device) traffic from data traffic. To achieve this isolation, the Firewall Threat Defense device uses a separate routing table for management-only traffic vs. data traffic. Separate routing tables means that you can create separate default routes for data and management as well.

Types of Traffic for Each Routing Table

Through-the-device traffic always uses the data routing table.

From-the-device traffic, depending on the type, uses either the management-only routing table or the data routing table by default. If a match is not found in the default routing table, it checks the other routing table.

- Management-only table from-the-device traffic includes AAA server communications.
- Data table from-the-device traffic includes DNS server lookups and DDNS. An exception is if you only specify the Diagnostic interface for DNS, then the Firewall Threat Defense device will only use the management-only table.

Interfaces Included in the Management-Only Routing Table

Management-only interfaces include any the Management x/x interfaces as well as any interfaces that you have configured to be management-only.



Note The Management virtual interface uses its own Linux routing table that is not part of the Firewall Threat Defense route lookup. Traffic originating on the Management interface includes the Firewall Device Manager management sessions, licensing communication, and database updates. The Diagnostic logical interface, on the other hand, uses the management-only routing table described in this section.

Fallback to the Other Routing Table

If a match is not found in the default routing table, it checks the other routing table.

Using the Non-Default Routing Table

If you need from-the-box traffic to go out an interface that isn't in its default routing table, then you might need to specify that interface when you configure it, rather than relying on the fall back to the other table. The Firewall Threat Defense will only check routes for the specified interface. For example, if you need to communicate with a RADIUS server on a data interface, then specify that interface in the RADIUS configuration. Otherwise, if there is a default route in the management-only routing table, then it will match the default route and never fall back to the data routing table.

Equal-Cost Multi-Path (ECMP) Routing

The Firewall Threat Defense 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.

ECMP Across Multiple Interfaces Using Traffic Zones

If you configure traffic zones to contain a group of interfaces, you can have up to 8 equal cost static or dynamic routes across up to 8 interfaces within each zone. For example, you can configure multiple default routes across three interfaces in the zone:

```
route for 0.0.0.0 0.0.0.0 through outside1 to 10.1.1.2
route for 0.0.0.0 0.0.0.0 through outside2 to 10.2.1.2
route for 0.0.0.0 0.0.0.0 through outside3 to 10.3.1.2
```

Similarly, your dynamic routing protocol can automatically configure equal cost routes. The Firewall Threat Defense load-balances traffic across the interfaces with a more robust load balancing mechanism.

When a route is lost, the device seamlessly moves the flow to a different route.

Static Routes

You can create static routes to provide basic routing for your network.

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 Firewall Threat Defense 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 the Firewall Threat Defense device 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, 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.

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 Firewall Threat Defense.
- You are using a feature that does not support dynamic routing protocols.
- Virtual routers use static routes to create route leaks. Route leaks enable flow of traffic from an interface of a virtual router to another interface in another virtual router. For more information, see Virtual router interconnection.

Backup Static Routes and 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. Static routes 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 goes down.

By implementing route tracking, using a Service Level Agreement (SLA) monitor, you can track the availability of a static route and automatically install 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.

When you use route tracking, you associate a target IP address on the destination network to the tracked route. The system then uses ICMP echo requests to periodically verify that the address can be reached. If the system does not receive an echo reply within the time period you specify, the host is considered unreachable, and the system removes the associated route from the routing table. The system can then use an untracked backup route with a higher metric in place of the removed route.

Thus, to use a backup static route for a given destination, including for a default route, you must do the following:

1. Create an SLA monitor that will monitor a reliable IP address on the destination network, such as a gateway or an always-up server (such as a web server or syslog server). Do not monitor the IP address of a system that might be taken off-line while the destination network remains healthy and available. See Configure SLA Monitor Objects, on page 313.
2. Create the primary route to the destination and select the SLA monitor for the route. The metric for this route should typically be 1. See Configuring Static Routes, on page 312.
3. Create the backup static route that will be used if the primary route fails. This route must have a larger metric than the primary route. For example, if the primary route is 1, the backup route could be 10. You would also normally select a different interface for the backup route.

Guidelines for Static Routing

Bridge Groups

- In routed mode, you must specify the BVI as the gateway; you cannot specify the member interface.
- For traffic that originates on the Firewall Threat Defense (such as syslog or SNMP) that 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 Firewall Threat Defense knows out of which bridge group member interface to send traffic. If you have servers that cannot all be reached through a single default route, then you must configure static routes.
- Static route tracking is not supported for bridge group member interfaces or on the BVI.

IPv6

- Static route tracking (SLA monitor) is not supported for IPv6.

Equal-Cost Multi-Path (ECMP) Traffic Zones

- Keep the member interfaces of a ECMP traffic zone in the same security zone to prevent different access, SSL, or identity rules being applied to those interfaces.
- You can have up to 8 equal-cost routes for a network in a given ECMP traffic zone.
- You can create up to 256 ECMP traffic zones, with up to 8 interfaces per zone.
- ECMP traffic zones can contain named physical interfaces, subinterfaces, and Etherchannels. They cannot contain the following:
  - A bridge group (BVI) or its members
  - Etherchannel member interfaces
  - HA interfaces (failover or state links)
  - Management-only interfaces
  - Interfaces used for site-to-site VPN or remote access VPN connections.

- Virtual tunnel interfaces (VTI) or their source interfaces.
- Interfaces configured for VPN management access.

- You cannot enable DHCP Relay on an interface in a zone.

Configuring Static Routes

Define static routes to tell the system where to send packets that are not bound for networks that are directly connected to the interfaces on the system.

You need at least one static route, the default route, for network 0.0.0.0/0. This route defines where to send packets whose egress interface cannot be determined by existing NAT xlates (translations) or static NAT rules, or other static routes.

You might need other static routes if the default gateway cannot be used to get to all networks. For example, the default route is usually an upstream router on the outside interface. If there are additional inside networks that are not directly connected to the device, and they cannot be accessed through the default gateway, you need static routes for each of those inside networks.

You cannot define static routes for the networks that are directly connected to system interfaces. The system automatically creates these routes.

Procedure

Step 1 Click **Device**, then click the link in the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring a static route.

Step 3 On the **Static Routing** page, do one of the following:

- To add a new route, click **+**.
- Click the edit icon ( ) for the route you want to edit.

If you no longer need a route, click the trash can icon for the route to delete it.

Step 4 Configure the route properties

- **Name**—A display name for the route.
- **Description**—An optional description of the purpose for the route.
- **Interface**—Select the interface through which you want to send traffic. The gateway address needs to be accessible through this interface.

  For bridge groups, you configure the route for the bridge group interface (BVI), not for the member interfaces.

  If you have enabled virtual routing and forwarding, you can select an interface that belongs to a different virtual router. If you create a static route in a virtual router for an interface in a different virtual router, the route will cross virtual router boundaries, with the risk that traffic from this virtual router will leak into another virtual router. This might be the desired result, but carefully determine that you need this

route leak. As you select interfaces, the name of the virtual router to which the interface belongs will be shown to the right of the interface.

- **Protocol**—Select whether the route is for an **IPv4** or **IPv6** address.
- **Networks**—Select the network objects that identify the destination networks or hosts that should use the gateway in this route.

  To define a default route, use the pre-defined any-ipv4 or any-ipv6 network objects, or create an object for the 0.0.0.0/0 (IPv4) or ::/0 (IPv6) network.
- **Gateway**—Select the host network object that identifies the IP address for the gateway. Traffic is sent to this address. You cannot use the same gateway for routes on more than one interface.

  If you are defining a route in a virtual router, and the interface belongs to a different virtual router, you must leave the gateway empty. The system will route traffic to these networks to the other virtual router, and then use the target virtual router's routing table to determine the gateway.
- **Metric**—The administrative distance for the route, between 1 and 254. The default for static routes is 1. If there are additional routers between the interface and the gateway, enter the number of hops as the administrative distance.

  Administrative distance is a parameter used to compare routes. The lower the number, the higher precedence the route is given. Connected routes (networks directly connected to an interface on the device) always take precedence over static routes.

Step 5 (Optional; IPv4 routes only.) Select the **SLA Monitor** that should track this route's viability.

An SLA Monitor can verify that an always-available host on the target network is reachable. If it becomes unreachable, the system can install a backup route. Thus, if you configure an SLA Monitor, you should also configure another static route, with a larger metric, for this network. For example, if this route has the metric 1, create a backup route with the metric 10. For more information, see Backup Static Routes and Static Route Tracking, on page 310.

If the SLA Monitor object does not yet exist, click the **Create SLA Monitor** link at the bottom of the list and create it now.

Note
If a monitored route is removed because the monitored address cannot be pinged, the route is indicated in the static route table with a warning that the route is unreachable. Determine if the problem is temporary or if you need to reconfigure the route. Consider the possibility that the route is viable but that the monitored address is not sufficiently dependable.

Step 6 Click **OK**.

Configure SLA Monitor Objects

Configure Service Level Agreement (SLA) Monitor objects for use with static routes. By using an SLA monitor, you can track the health of a static route and automatically replace a failed route with a new one. For more information on route tracking, see Backup Static Routes and Static Route Tracking, on page 310.

When selecting a monitoring target, you need to make sure that it can respond to ICMP echo requests. The target can be any IP address, defined in a host network object, but you should consider using the following:

- The ISP gateway address, for dual ISP support.

- 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 system needs to communicate with.
- A persistent IP address on the destination network. A workstation that might be shut down at night is not a good choice.

Procedure

Step 1 Select **Objects**, then select **SLA Monitors** from the table of contents.

Step 2 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon () for the object.

To delete an unreferenced object, click the trash can icon () for the object.

Step 3 Enter a Name for the object and optionally, a description.

Step 4 Define the SLA monitor required options:

- **Monitor Address**—Select the host network object that defines the address to be monitored on the destination network. You can click **Create New Network** if the object you require does not exist.

  This address is monitored only if you attach the SLA monitor to a static route.
- **Target Interface**—Select the interface through which to send the echo request packets. This is normally the interface on which you will define the static route. The interface source address is used as the source address in the echo request packets.

Step 5 (Optional.) Adjust the **IP ICMP Echo Options**.

All of the ICMP options have defaults that are appropriate in most cases, but you can tune them to fit your requirements.

- **Threshold**—The number of milliseconds for a rising threshold to be declared, from 0 to 2147483647. The default is 5000 (5 seconds). This value should not be larger than the value set for the timeout. The threshold value is only used to indicate over threshold events, which do not affect reachability. You can use the frequency of threshold events to evaluate the setting for timeout.
- **Timeout**—The amount of time, in milliseconds, that the route monitoring operation should wait for a response from the request packets, from 0 to 604800000 milliseconds (7 days). The default value is 5000 milliseconds (5 seconds). If the monitor does not get a response to at least one echo request during this period, the process installs the backup route.
- **Frequency**—The number of milliseconds between SLA probes, from 1000 to 604800000, in multiples of 1000. You cannot set a frequency that is less than the timeout value. The default is 60000 milliseconds (60 seconds).
- **Type of Service**—An integer that defines the Type of Service (ToS) type in the IP header of the ICMP echo request packet, from 0 to 255. The default is 0.

- **Number of Packets**—The number of packets to be sent with each poll, from 1 to 100. The default is 1 packet.
- **Data Size**—The size of the data payload to use in the echo request packets, from 0 to 16384 bytes. The default value is 28. This setting specifies the size of the payload only; it does not specify the size of the entire packet.

Step 6 Click **OK**.

You can now use the SLA monitor object in a static route.

Configuring ECMP Traffic Zones

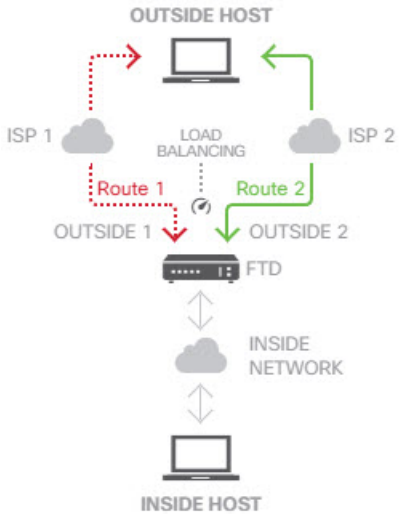
Normally, to configure multiple routes for a given network prefix, using the same route metric, you need to configure the routes on the same interface. Thus, the system uses Equal-Cost Multi-Path (ECMP) routing calculations to load balance traffic sent through the interface to the gateways.

For example, you can configure multiple default routes on the outside interface that specify different gateways, and this configuration is allowed without additional changes:

```
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
```

You can also use ECMP to balance traffic across multiple interfaces (within a virtual router) for the same network prefix and route metric. This configuration is necessary if the gateways are accessible through separate interfaces. For example, say you have two ISPs, and you want to load balance among the ISPs, but you do not want to divide your internal address spaces between the ISP gateways. One ISP is accessible through the outside1 interface, the other through the outside2 interface. To accomplish this, you would need to create a routing traffic zone that contains the outside1 and outside2 interfaces.

```
isp-zone containing outside1 and outside2
route for 0.0.0.0 0.0.0.0 through outside1 to 10.1.1.2
route for 0.0.0.0 0.0.0.0 through outside2 to 10.1.1.3
```



Note An ECMP routing traffic zone is not related to security zones. Creating a security zone that contains the outside1 and outside2 interfaces does not implement a traffic zone for ECMP routing purposes.

The following procedure explains how to configure ECMP zones to take advantage of ECMP processing across interfaces.

Procedure

Step 1 Click **Device**, then click the link in the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring a static route.

Step 3 Click the **ECMP Traffic Zones** tab.

Step 4 On the **ECMP Traffic Zones** page, do one of the following:

- To add a new zone, click +, or **Add ECMP Traffic Zone**.
- Click the edit icon ( ) for the zone you want to edit.

If you no longer need a zone, click the trash can icon for the zone to delete it. You must remove any static routes that depend on a zone before you can delete the zone.

Step 5 Enter a **Name** for the zone and optionally, a description.

Step 6 Select up to 8 **Interfaces** to include in the zone:

- Click + to add an interface.
- Click **x** to the right of an interface to remove it.

While selecting interfaces, keep the following restrictions in mind:

- You can select physical interfaces, subinterfaces, and Etherchannels.
- An ECMP traffic zone cannot include the following types of interface: a bridge group (BVI) or its members, Etherchannel member interfaces, HA interfaces (failover or state links), management-only interfaces, or interfaces configured for VPN management access.
- You cannot include interfaces that are used in remote access or site-to-site VPN connections.
- You cannot select interfaces that are enabled for DHCP relay, either as server or agent.
- The interfaces must be assigned to the same virtual router.
- An interface can be in one traffic zone only.

Step 7 Click **OK**.

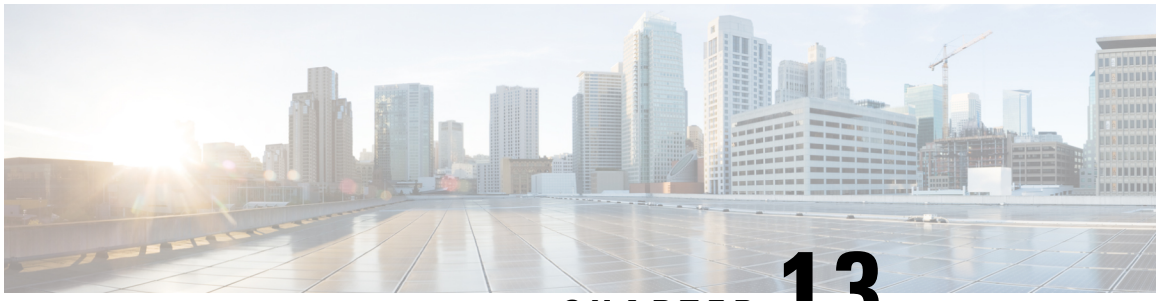
What to do next

You can now go to the Static Routes tab and create equal-cost routes through these interfaces for the same destination. Alternately, your dynamic routing protocols can automatically configure equal-cost routes if they are distributed through the system.

Monitoring Routing

To monitor and troubleshoot routing, open the CLI console or log into the device CLI and use the following commands. You can also select some of these commands from the **Commands** menu on the Routing page.

- **show route** displays the routing table for the data interfaces, including routes for directly-connected networks.
- **show ipv6 route** displays the IPv6 routing table for the data interfaces, including routes for directly-connected networks.
- **show network** displays the configuration for the virtual Management interface, including the management gateway. Routing through the virtual Management interface is not handled by the data interface routing table, unless you specify data-interfaces as the management gateway.
- **show network-static-routes** displays static routes configured for the virtual Management interface using the **configure network static-routes** command. Normally, there will not be any static routes, as the management gateway suffices for management routing in most cases. These routes are not available to traffic on the data interfaces. This command is not available in the CLI console.
- **show ospf** displays information about the OSPF processes and learned routes. Use **show ospf ?** to get a list of options you can include to view specific information about OSPF.
- **show bgp** displays information about the BGP processes and learned routes. Use**show bgp ?** to get a list of options you can include to view specific information about BGP.
- **show eigrp** *option* displays information about the EIGRP processes and learned routes. Use **show eigrp ?** to get a list of options you can include; you must supply an option.
- **show isis** *option* displays information about the IS-IS processes and learned routes. Use **show isis ?** to get a list of options you can include; you must supply an option.
- **show rip database** displays information about the RIP processes and learned routes.
- **show vrf** displays information about the virtual routers defined on the system.
- **show zone** displays information about ECMP traffic zones, including the interfaces that are part of each zone.



CHAPTER 13

Virtual Routers

You can create virtual routers to isolate the traffic on subsets of interfaces from each other.

- Virtual routers and virtual routing and forwarding (VRF), on page 319
- Guidelines for Virtual Routers, on page 322
- Managing Virtual Routers, on page 324
- Examples for Virtual Routers, on page 327
- Monitoring Virtual Routers, on page 343

Virtual routers and virtual routing and forwarding (VRF)

A virtual router is a network virtualization feature that

- maintains separate routing tables for groups of interfaces
- provides clean separation in the traffic flowing through the device, and
- implements the "light" version of Virtual Routing and Forwarding, or VRF-Lite, which does not support Multiprotocol Extensions for BGP (MBGP).

Virtual router capabilities and configuration

Virtual routers provide support to two or more distinct customers over a common set of networking equipment. You can also use virtual routers to provide more separation for elements of your own network, for example, by isolating a development network from your general purpose corporate network.

When you create a virtual router, you assign interfaces to the router. You can assign a given interface to one, and only one, virtual router. You would then define static routes, and configure routing protocols such as OSPF or BGP, for each virtual router. You would also configure separate routing processes over your entire network, so that routing tables on all participating devices are using the same per-virtual-router routing process and tables. Using virtual routers, you create logically-separated networks over the same physical network to ensure the privacy of the traffic that runs through each virtual router.

Because the routing tables are separate, you can use the same, or overlapping, address spaces across the virtual routers. For example, you could use the 192.168.1.0/24 address space for two separate virtual routers, supported by two separate physical interfaces.

Note that there are separate management and data routing tables per virtual router. For example, if you assign a management-only interface to a virtual router, then the routing table for that interface is separate from the data interfaces assigned to the virtual router.

Virtual-router-aware policies

Virtual-router-aware policies are security policies that

- apply to interfaces from a single virtual router only
- use security zones whose memberships are constrained to interfaces assigned to a specific virtual router, and
- prevent policies from affecting all interfaces across all virtual routers when that is not desired.

Virtual router policy behavior

When you create a virtual router, the routing table for that virtual router is automatically separated from the global virtual router or any other virtual router. However, security policies are not automatically virtual-router-aware.

For example, if you write an access control rule that applies to "any" source or destination security zone, then the rule will apply to all interfaces across all virtual routers. This might in fact be exactly what you want. For example, all of your customers might want to block access to the same list of objectionable URL categories.

But, if you need to apply a policy to one of the virtual routers but not others, you need to create security zones that contain interfaces from that single virtual router only. Then, use the virtual-router-constrained security zones in the source and destination criteria of the security policy.

By using security zones whose memberships are constrained to the interfaces assigned to a single virtual router, you can write virtual-router-aware rules in these policies:

- Access control policy.
- Intrusion and file policies.
- SSL decryption policy.
- Identity policy and user-to-IP address mappings. If you use overlapping address spaces in virtual routers, ensure that you create separate realms for each virtual router and apply them correctly in the identity policy rules.

If you use overlapping address spaces in your virtual routers, you should use security zones to ensure that the right policies get applied. For example, if you use the 192.168.1.0/24 address space in two separate virtual routers, an access control rule that simply specifies the 192.168.1.0/24 network will apply to traffic in both virtual routers. If that is not the desired outcome, you can limit the application of the rule by also specifying the source/destination security zones for just one of the virtual routers.

Routing Between Virtual Routers

You can configure static routes to route traffic between virtual routers.

For example, if you have the outside interface in the global virtual router, you can set up static default routes in each of the other virtual routers to send traffic to the outside interface. Then, any traffic that cannot be routed within a given virtual router gets sent to the global router for subsequent routing.

Static routes between virtual routers are known as route leaks, because you are leaking traffic to a different virtual router. When you are leaking routes, say, VR1 routes to VR2, you can initiate connections from VR2 to VR1 only. For traffic to flow from VR1 to VR2, you must configure the reverse route. When you create a

static route to an interface in another virtual router, you do not need to specify a gateway address. Simply select the destination interface.

For inter-virtual-router routes, the system does destination interface look-up in the source virtual router. Then, it looks up the MAC address of the next hop in the destination virtual router. Thus, the destination virtual router must have either a dynamic (learned) or static route for the selected interface for the destination address.

Configuring NAT rules that use source and destination interfaces in different virtual routers can also allow traffic to route between virtual routers. If you do not select the option for NAT to do a route lookup, the rule will simply send traffic out the destination interface with a NATed address whenever destination translation happens. However, the destination virtual router should have a route for the translated destination IP address so that next-hop lookup can succeed.

Maximum Number of Virtual Routers By Device Model

The maximum number of virtual routers you can create depends on the device model. The following table shows the maximum supported virtual routers per model. You can confirm the maximum for your system by entering the **show vrf counters** command, which displays the maximum number of user-defined virtual routers for your platform (not including the global virtual router). The numbers below include both user and global routers. For Firepower 4100/9300 devices, these values apply to native mode.

For platforms that support multi-instance capability, such as the Firepower 4100/9300, determine the maximum number of virtual routers per container instance by dividing the maximum virtual routers by the number of cores on the device, and then multiplying by the number of cores assigned to the instance, rounding down to the nearest whole number. For example, if the platform supports a maximum of 100 virtual routers, and it has 70 cores, then each core would support a maximum of 1.43 virtual routers (rounded). Thus, an instance assigned 6 cores would support 8.58 virtual routers, which rounds down to 8, and an instance assigned 10 cores would support 14.3 virtual routers (rounding down, 14).

Device Model	Maximum Virtual Routers
Firepower 1010	Virtual routers are not supported on this model.
Firepower 1120	5
Firepower 1140	10
Firepower 1150	10
Firepower 2110	10
Firepower 2120	20
Firepower 2130	30
Firepower 2140	40
Secure Firewall 3110	15
Secure Firewall 3120	25
Secure Firewall 3130	50
Secure Firewall 3140	100

Device Model	Maximum Virtual Routers
Firepower 4110	60
Firepower 4112	60
Firepower 4115	80
Firepower 4120	80
Firepower 4125	100
Firepower 4140	100
Firepower 4145	100
Firepower 4150	100
Firepower 9300 appliance, all models	100
Firewall Threat Defense Virtual, all platforms	30
ISA 3000	10

Guidelines for Virtual Routers

Device Model Guidelines

You can configure virtual routers on all support device models except the following:

- Firepower 1010

Additional Guidelines

- You can configure the following features on the global virtual router only:
  - RIP
  - EIGRP
  - IS-IS
  - BGPv6
  - Multicast Routing
  - Policy Based Routing
  - VPN
- You can configure the following features separately for each virtual router:

  - Static routes and their SLA monitors.
  - OSPFv2
  - BGPv4

- The following features are used by the system when querying or communicating with the remote system (from-the-box traffic). These features use interfaces in the global virtual router only. If you configure an interface for the feature, it must belong to the global virtual router. As a general rule, any time the system must look up a route to reach an external server for its own management purposes, it does the route lookup in the global virtual router.
  - DNS server when used to resolve fully-qualified names used in access control rules, or for resolving names for the **ping** command. If you specify **any** as the interface for a DNS server, the system considers interfaces in the global virtual router only.
  - AAA server or identity realm when used with VPN. You can configure VPN on interfaces in the global virtual router only, so the external AAA servers used for VPN, such as Active Directory, must be reachable through an interface in the global virtual router.
  - Syslog server.
  - SNMP.
- In NAT, if you specify source and destination interfaces that are assigned to different virtual routers, the NAT rule diverts traffic from one virtual router through another virtual router. Ensure that you do not mix interfaces in NAT rules unintentionally. Normally, the source and destination interfaces are used and the routing table is ignored, including for destination translations in manual NAT. However, if the NAT rule does need to do a route lookup, it does the lookup in the VRF table for the inbound interface only. If necessary, define static routes in the source virtual router for the destination interface. Note that if you leave the interface as **any**, the rule applies to all interfaces regardless of virtual router membership. When using virtual routers, carefully test your NAT rules to ensure you get the expected behavior. If you neglect to define a needed route leak, in some cases the rule will not match all of the traffic you expect it to match, and the translation will not be applied.
- If you configure inter-virtual-router routes, for example, leaking a route from one virtual router to a second virtual router, the system does destination interface look-up in the source virtual router. Then, it looks up the MAC address of the next hop in the destination virtual router. Thus, the destination virtual router must have either a dynamic (learned) or static route for the selected interface for the destination address.
- When using inter-virtual-router routes (leaked routes), for example, from virtual router 1 to virtual router 2, you do not need to configure a mirror (reverse) route in virtual router 2 to allow return traffic. However, if you want to allow connections to be initiated in both directions, ensure that you leak the route in both directions, from virtual router 1 to 2, and from virtual router 2 to 1.
- If you move an interface from one virtual router to another, all features configured for the interface are retained. Examine the configuration to ensure that static routes, IP addresses, and other policies make sense within the context of the new virtual router.
- If you use overlapping address spaces in multiple virtual routers, please be aware that static security group tag (SGT) to IP address mappings downloaded from Cisco Identity Services Engine (ISE) are not virtual-router-aware. Set up separate identity realms per virtual router if you need to create different SGT mappings per virtual router. This is not necessary if you intend to map the same IP addresses to the same SGT number in each virtual router.

- If you use overlapping address spaces in multiple virtual routers, dashboard data can be misleading. Connections for the same IP address are aggregated, so it will appear there was more traffic to/from a given address when it is shared by two or more endpoints. If you carefully construct your identity policies using separate identity realms, user-based statistics should be more accurate.
- You cannot use overlapping DHCP address pools in separate virtual routers.
- You can use DHCP server auto configuration on an interface in the global virtual router only. Auto configuration is not supported for interfaces assigned to a user-defined virtual router.
- If you move an interface from one virtual router to another, including from the global virtual router to a new router, any existing connections through the interface are dropped.
- The Security Intelligence policy is not virtual-router-aware. If you add an IP address, URL, or DNS name to the block list, it is blocked for all virtual routers.

Managing Virtual Routers

You can create multiple virtual routing and forwarding (VRF) instances, called virtual routers, to maintain separate routing tables for groups of interfaces. Because each virtual router has its own routing table, you can provide clean separation in the traffic flowing through the device.

Thus, you can provide support to two or more distinct customers over a common set of networking equipment. You can also use virtual routers to provide more separation for elements of your own network, for example, by isolating a development network from your general-purpose corporate network.

By default, virtual routing is disabled. The entire device uses a single set of global routing tables, for data (through) and management (to/from the box) traffic.

When you enable virtual routing, the initial routing page is a list of the virtual routers defined on the system. If you do not enable virtual routers, the initial routing page is a list of static routes defined on the system.

There is always a global virtual router. The global router holds all interfaces that you have not assigned to individual virtual routers.

Procedure

Step 1 Click **Device**, then click the link in the **Routing** summary.

Step 2 If you have not yet enabled virtual routers, click the **Add Multiple Virtual Routers** link, then click **Create First Custom Virtual Router**.

Creating the first virtual router is essentially the same as creating any additional virtual routers. For more information, see Create a Virtual Router or Edit Interface Assignments, on page 325.

Step 3 Do any of the following:

- To configure global BGP settings, which apply to all virtual routers, click the **BGP Global Settings** button. You configure these settings using Smart CLI, which is explained in Configuring Smart CLI Objects, on page 840. Configure global BGP settings only if you configure BGP in one or more virtual router.
- To create a new virtual router, click the + button above the table.

- To edit the routing properties of a virtual router, for example, to create static routes or to define routing processes, click the view icon ( ) in the Action cell for the virtual router.
- To edit the name, description, or interface assignments for a virtual router, click the view icon ( ) in the Action cell for the virtual router, then select the **Virtual Router Properties** tab.
- To switch between virtual routers when viewing them, click the down arrow next to the virtual router name (above the routing table) and select the desired virtual router. You can return to the listing page by clicking the **Go Back to Virtual Routers** arrow ( ).
- To delete a virtual router, click the delete icon ( ) in the Action cell for the virtual router, or the delete icon next to the virtual router name when viewing the virtual router's contents. When you delete the last virtual router (other than the global router, which you cannot delete), VRF is disabled.
- To monitor routing in a virtual router, click the link for one of the **show** commands in the table for that virtual router. Clicking the command opens the CLI Console, where you can examine the output of the CLI command. You can show information on routes, OSPF, and OSPF neighbors. Note that the command output is based on the deployed configuration; you will not see anything related to undeployed edits.

  You can also execute these commands by selecting them from the **Commands** drop-down list when viewing the virtual router.

Create a Virtual Router or Edit Interface Assignments

Before you can configure static routes or routing process on a virtual router, you must create the router and assign interfaces to it.

Before you begin

Go to the **Interfaces** page and ensure that each interface you want to add to the virtual router has a name. You cannot add an interface to a virtual router until it has a name.

Procedure

Step 1 Click **Device** > **Routing**.

Step 2 Do one of the following:

- If you have not yet created a virtual router, click the **Add Multiple Virtual Routers** link, then click **Create First Custom Virtual Router**
- Click the + button above the list of virtual routers to create a new one.
- Click the edit icon ( ) for a virtual router to edit its properties and interface list.
- When viewing a virtual router, click the **Virtual Router Properties** tab to edit the properties of the virtual router you are viewing.
- When viewing a virtual router, click the down arrow next to the virtual router name and click **Create New Virtual Router**.

Step 3 Configure the properties of the virtual router:

- **Name**—The name for the virtual router.
- **Description**—An optional description of the virtual router.
- **Interfaces**—Click **+** to select each interface that should be part of the virtual router. To remove an interface, hover over the interface and click **X** on the right side of the interface card. You can assign physical interfaces, subinterfaces, bridge groups, and EtherChannels to a virtual router, but not VLANs.

  The routing table will be constrained to these interfaces, unless you intentionally leak routes to other interfaces into the virtual routing table.

  You can assign the diagnostic (Management X/Y) interface to the global virtual router only.

Step 4 Click **OK** or **Save**.

You are taken to the view for this virtual router, where you can configure static routes or routing processes.

Configure Static Routes and Routing Processes in a Virtual Router

Each virtual router has its own static routes and routing processes, which operate separately from the routes and routing processes defined for any other virtual router.

When you configure static routes, you can select destination interfaces that are outside of the virtual router. This leaks the route into the virtual router that contains the destination interface. Ensure that you leak only those routes that you need to leak, to ensure that you do not send more traffic than you want to the other virtual router. For example, if you have one path to the Internet, then it makes sense to leak routes from each virtual router to the Internet-facing virtual router for traffic destined for the Internet.

Procedure

Step 1 Choose **Device** > **Routing**.

Step 2 Click the view icon ( ) in the Action cell for the virtual router to open it.

Step 3 Do any of the following:

- To configure static routes, click the **Static Routing** tab, then create or edit the routes. For detailed information, see Configuring Static Routes, on page 312.
- To configure Equal-Cost Multi-Path (ECMP) traffic zones, click the **ECMP Traffic Zones** tab, then create the zones. For detailed information, see Configuring ECMP Traffic Zones, on page 315.
- To configure the BGP routing process, click the **BGP** tab, then create the Smart CLI object that is needed to define the process. For detailed information, see Border Gateway Protocol (BGP), on page 401.

  There are also global settings for BGP that apply to all virtual routers. You must return to the virtual router listing page to click the **BGP Global Settings** button to configure these properties.
- To configure the OSPF routing process, click the **OSPF** tab, then create the Smart CLI objects that are needed to define up to 2 processes, and their associated interface configurations. For detailed information, see Open Shortest Path First (OSPF), on page 363.

- (Global virtual router only.) To configure the EIGRP routing process, click the **EIGRP** tab, then create the Smart CLI object that is needed to define a single process. For detailed information, see Enhanced Interior Gateway Routing Protocol (EIGRP), on page 383.

Delete a Virtual Router

If you no longer need a virtual router, you can delete it. You cannot delete the global virtual router.

When you delete a virtual router, you are also deleting all of the static routes, and routing processes, that are configured within the virtual router.

All interfaces assigned to the virtual router are reassigned to the global router.

Procedure

Step 1 Choose **Device** > **Routing**.

Step 2 Do one of the following:

- In the list of virtual routers, click the delete icon ( ) in the Action column for the virtual router.
- When viewing the virtual router you want to delete, click the delete icon ( ) next to the router's name.

You are prompted to confirm that you want to delete the virtual router.

Step 3 Click **OK** to confirm the deletion.

Examples for Virtual Routers

The following topics provide examples of implementing virtual routers.

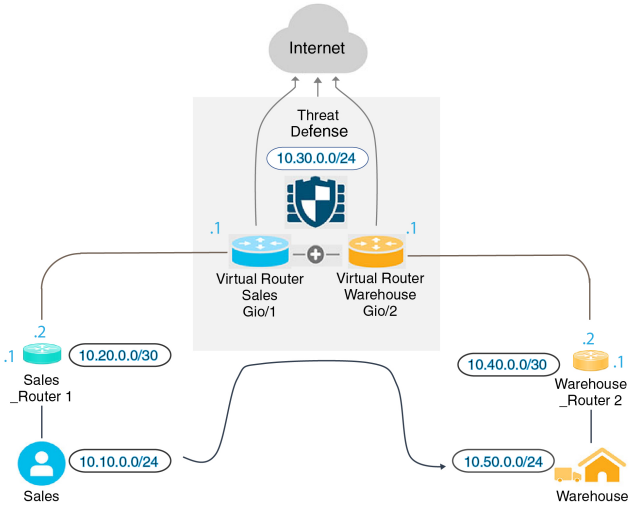
Related Topics

How to Secure Traffic from Networks in Multiple Virtual Routers over a Site-to-Site VPN, on page 665
How to Allow RA VPN Access to Internal Networks in Different Virtual Routers, on page 737

How to Route to a Distant Server through Multiple Virtual Routers

When you use virtual routers, you might have a situation where users in one virtual router need access to a server that is reachable only through a separate virtual router.

Consider the following illustration. The sales team's workstations are connected to the Sales virtual router. The warehouse servers are connected through the Warehouse virtual router. If the sales team needs to look up information on the warehouse server whose IP address is 10.50.0.5/24, you need to leak a route from the Sales virtual router into the Warehouse virtual router. The Warehouse virtual router must also have a route to the warehouse server, which is multiple hops away behind Warehouse Router 2.



Before you begin

This example assumes that you have already configured:

- Both the Sales and Warehouse virtual routers on the Firewall Threat Defense device, with GigabitEthernet 0/1 assigned to Sales, and GigabitEthernet 0/2 assigned to Warehouse.
- The Sales Router 1 has either a static or dynamic route that will send traffic to 10.50.0.5/24 out the 10.20.0.1/30 interface.

Procedure

Step 1 Create the network object for 10.50.0.5/24 or 10.50.0.0/24. Also, create the object for the gateway, 10.40.0.2/30.

If you want to limit the route to the single IP address of the warehouse server, use a host object to define 10.50.0.5. Alternatively, if the sales team should have access to other systems in the warehouse, create a network object for the 10.50.0.0/24 network. In this example, we will create a route for the host IP address.

a) Choose **Objects**, then **Networks** from the table of contents.
b) Click +, then fill in the object properties for the warehouse server:



c) Click **OK**.

d) Click +, then fill in the object properties for the router gateway to the warehouse network:



e) Click **OK**.

Step 2 Define the route leak in Sales that points to the Gi0/2 interface in the Warehouse virtual router.

In this example, Gi0/1 is named inside, and Gi0/2 is named inside-2.

a) Choose **Device**, then click **View Configuration** in the **Routing** summary.

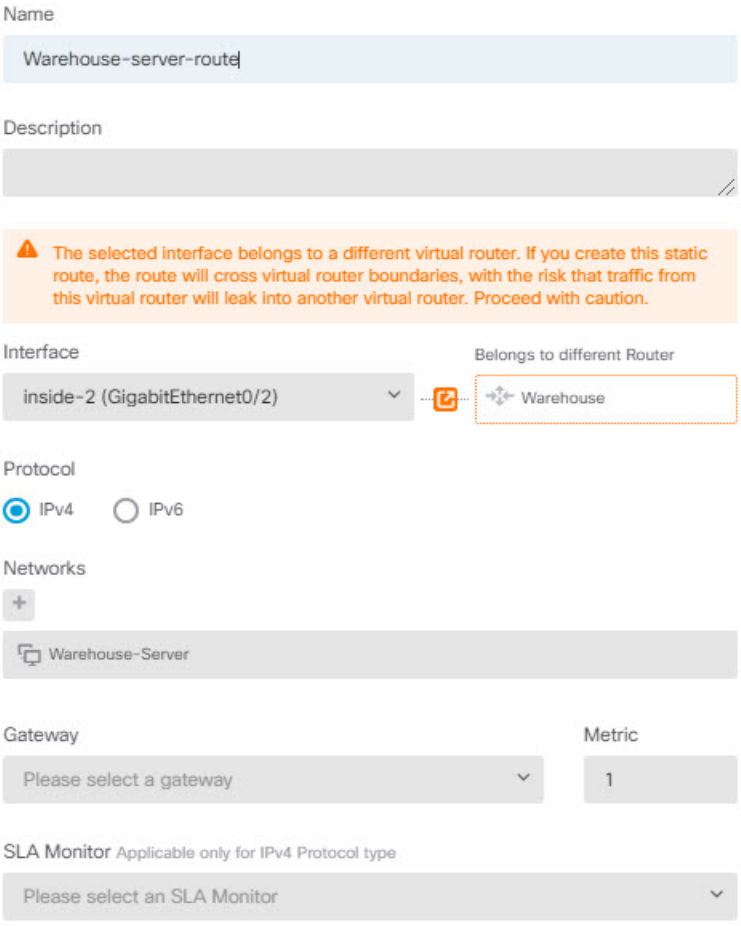
b) In the list of virtual routers, click the view icon ( ) in the action column for the Sales virtual router.

c) On the **Static Routing** tab, click + and configure the route:

- **Name**—Any name will do, such as Warehouse-server-route.
- **Interface**—Select **inside-2**. You will see a warning saying that the interface is in a different router and you are creating a route leak. This is what you want to do.
- **Protocol**—For this example, use **IPv4**. You can alternatively use IPv6 addresses to implement this example.
- **Networks**—Select the Warehouse-Server object.

- **Gateway**—Leave this item blank. When leaking a route into another virtual router, you do not select the gateway address.

The dialog box should look similar to the following:

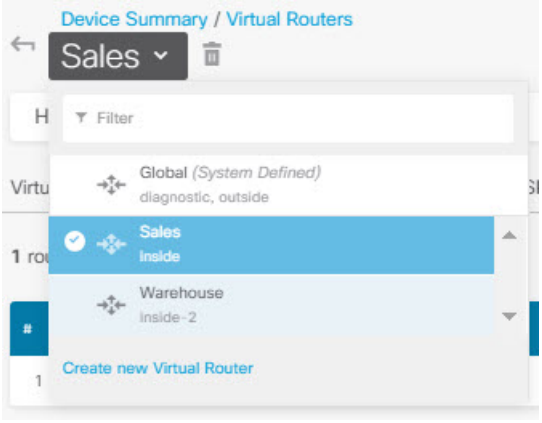


d) Click **OK**.

Step 3 In the Warehouse virtual router, define the route that points to the Warehouse Router 2 gateway.

Alternatively, this can be done by configuring a routing protocol that would dynamically discover the route from Warehouse Router 2. For this example, we will define the static route.

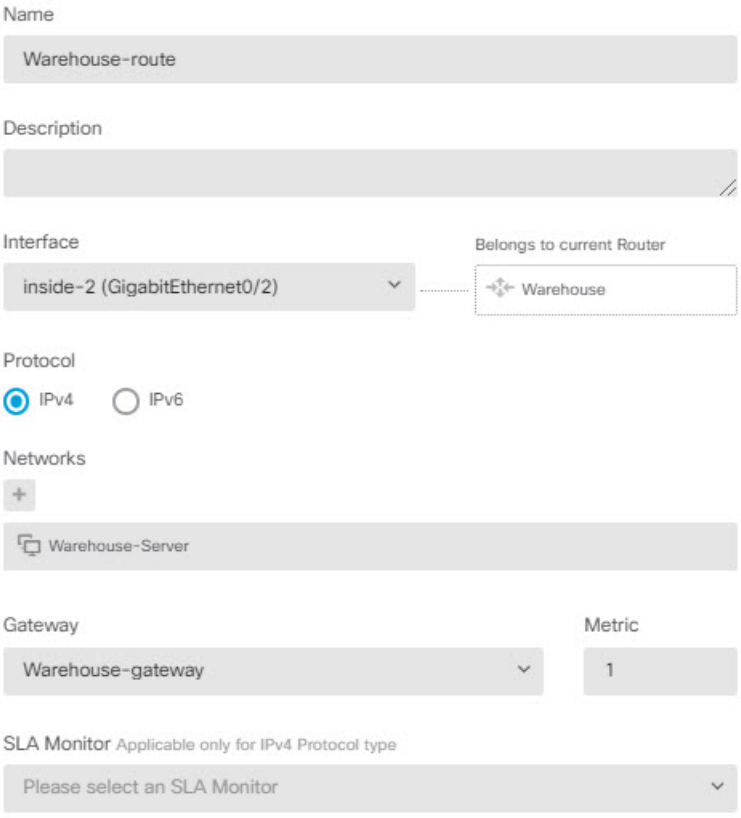
a) From the virtual router drop-down that currently says Sales, select the Warehouse virtual router to switch routers.



b) On the **Static Routing** tab, click + and configure the route:

- **Name**—Any name will do, such as Warehouse-route.
- **Interface**—Select **inside-2**.
- **Protocol**—Select **IPv4**.
- **Networks**—Select the Warehouse-Server object.
- **Gateway**—Select the Warehouse-gateway object.

The dialog box should look similar to the following:

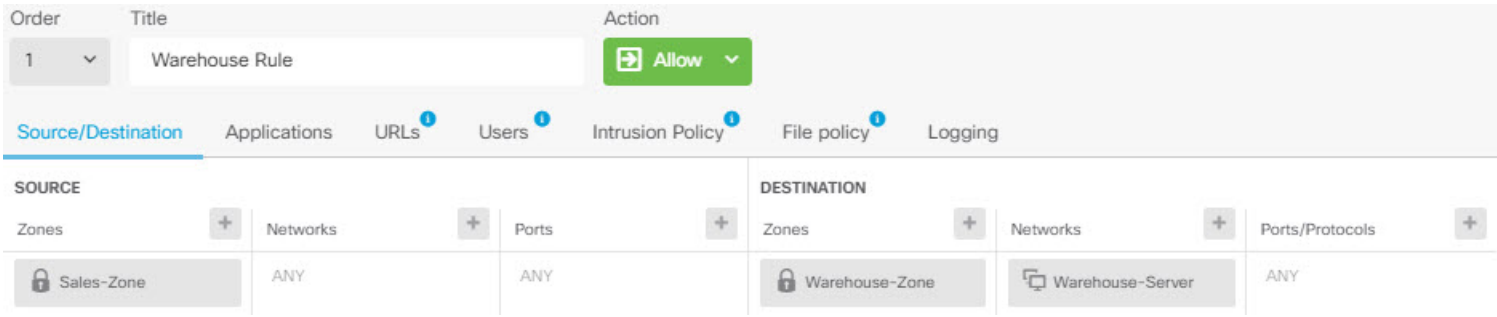


c) Click **OK**.

Step 4 Ensure that there is an access control rule that allows access to the warehouse server.

The simplest rule would allow traffic from the source interfaces in the Sales virtual router to the destination interfaces in the Warehouse virtual router for the destination Warehouse-Server network object. You can apply intrusion inspection to the traffic as you see fit.

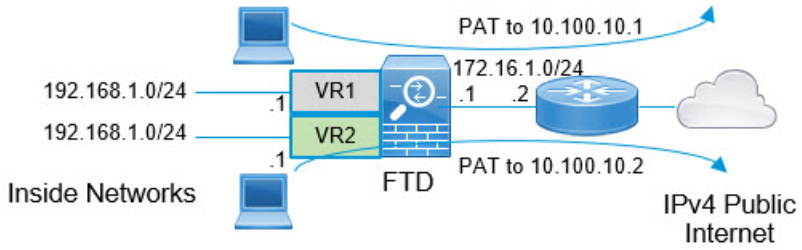
For example, if the interfaces in Sales are in the Sales-Zone security zone, and those in Warehouse are in the Warehouse-Zone security zone, the access control rule would look similar to the following:



How to Provide Internet Access to Multiple Virtual Routers with Overlapping Address Spaces

When using virtual routers, you can have the same network address for interfaces that reside in separate routers. For example, you could define interfaces inside and inside-2 to both use the IP address 192.168.1.1/24, managing endpoints on their segment in the 192.168.1.0/24 network. However, because the IP addresses routed in these separate virtual routers are the same, you need to carefully handle traffic that leaves the virtual routers, to ensure that return traffic goes to the correct destination.

For example, to allow Internet access from two virtual routers that use the same address space, you need to apply NAT rules separately to the interfaces within each virtual router, ideally using separate NAT or PAT pools. You could use PAT to translate the source addresses in virtual router 1 to 10.100.10.1, and for those in virtual router 2, to 10.100.10.2. The following illustration shows this setup, where the Internet-facing outside interface is part of the global router. Note that you must define the NAT/PAT rules with the source interface explicitly selected, because using “any” as the source interface will make it impossible for the system to identify the correct source because the same IP address could exist on 2 different interfaces.



Note This example is simplified, where each virtual router contains a single interface. If an “inside” virtual router has more than one interface, you need to create the NAT rules for each “inside” interface. Even if you have some interfaces within virtual routers that do not use overlapping address spaces, explicitly identifying the source interface in NAT rules might make troubleshooting easier, and ensure a cleaner separation between traffic from the virtual routers that is Internet-bound.

Procedure

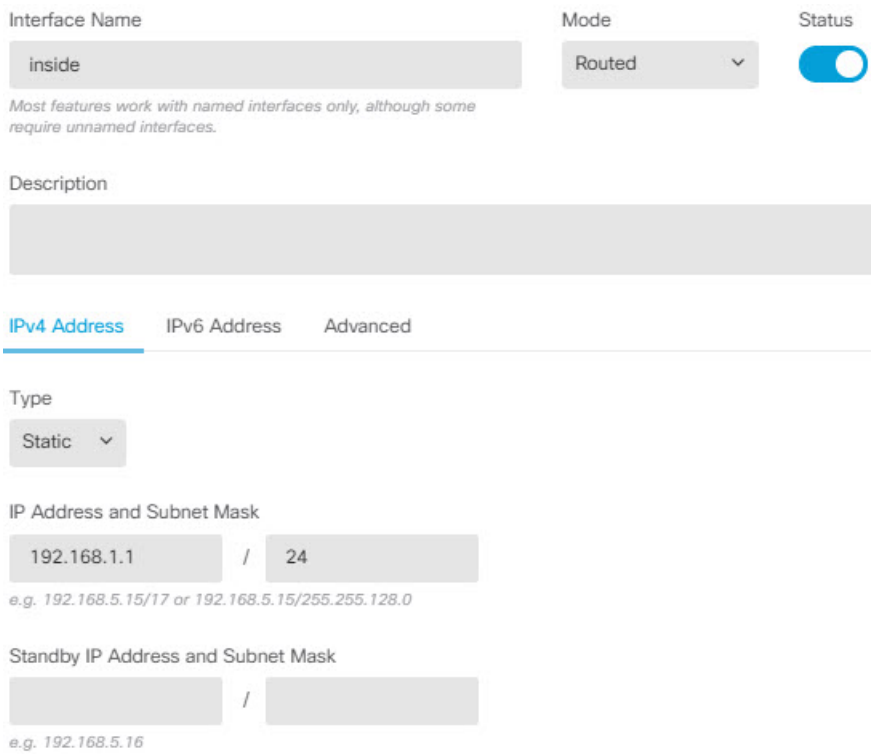
Step 1 Configure the inside interface for virtual router 1 (VR1).

a) Click **Device**, then **View All Interfaces** in the **Interface** summary.

b) Click the edit icon ( ) in the Action column for the interface you will assign to VR1.

c) Configure at least the following properties:

- **Name**—For this example, **inside**.
- **Mode**—Select **Routed**.
- **Status**—Enable the interface.
- **IPv4 Address** > **Type**—Select **Static**.
- **IPv4 Address and Subnet Mask**—Enter 192.168.1.1/24.



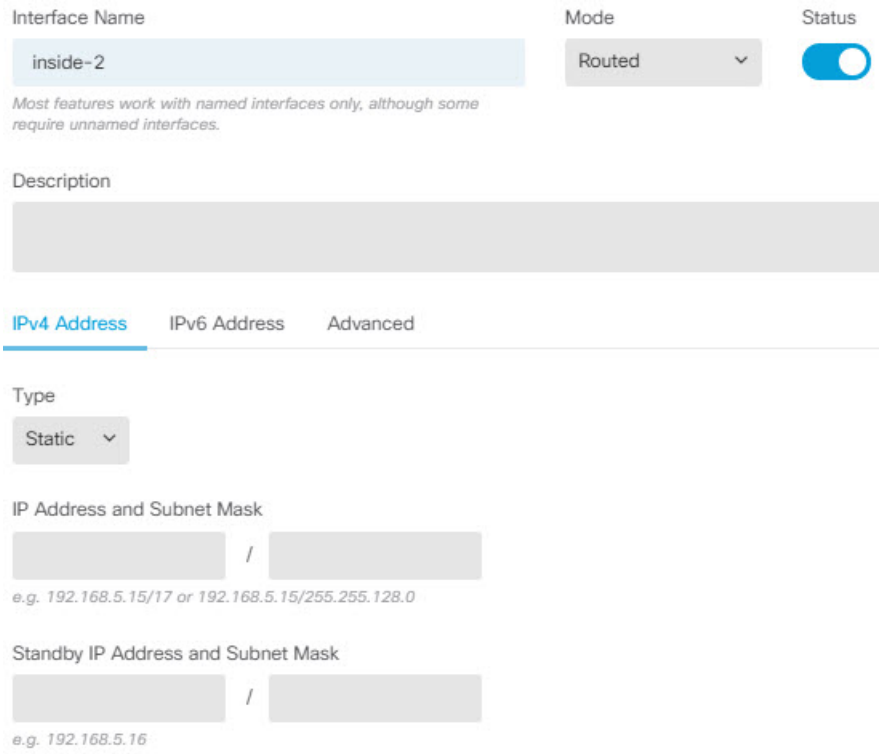
d) Click **OK**.

Step 2 Configure the inside-2 interface for virtual router 2 (VR2), but do not specify the IP address.

a) On the Interfaces listing page, click the edit icon ( ) in the Action column for the interface you will assign to VR2.

b) Configure at least the following properties:

- **Name**—For this example, **inside-2**.
- **Mode**—Select **Routed**.
- **Status**—Enable the interface.
- **IPv4 Address** > **Type**—Select **Static**.
- **IPv4 Address and Subnet Mask**—Leave these fields empty. If you try to configure the same address as the inside interface at this point, the system will show an error message and prevent you from creating a non-functional configuration. You cannot route to the same address space through different interfaces within the same router.



c) Click **OK**.

Step 3 Configure virtual router VR1, including the static default route leak to the outside interface.

a) Choose **Device**, then click **View Configuration** in the **Routing** summary.

b) Click **Add Multiple Virtual Routers** at the top of the Routing page.

c) At the bottom right of the explanatory panel, click **Create First Custom Virtual Router**.

d) Fill in the properties for virtual router VR1.

- **Name**—Enter VR1 or another name of your choosing.

- **Interfaces**—Click +, select **inside**, and click **OK**.



e) Click **OK**.

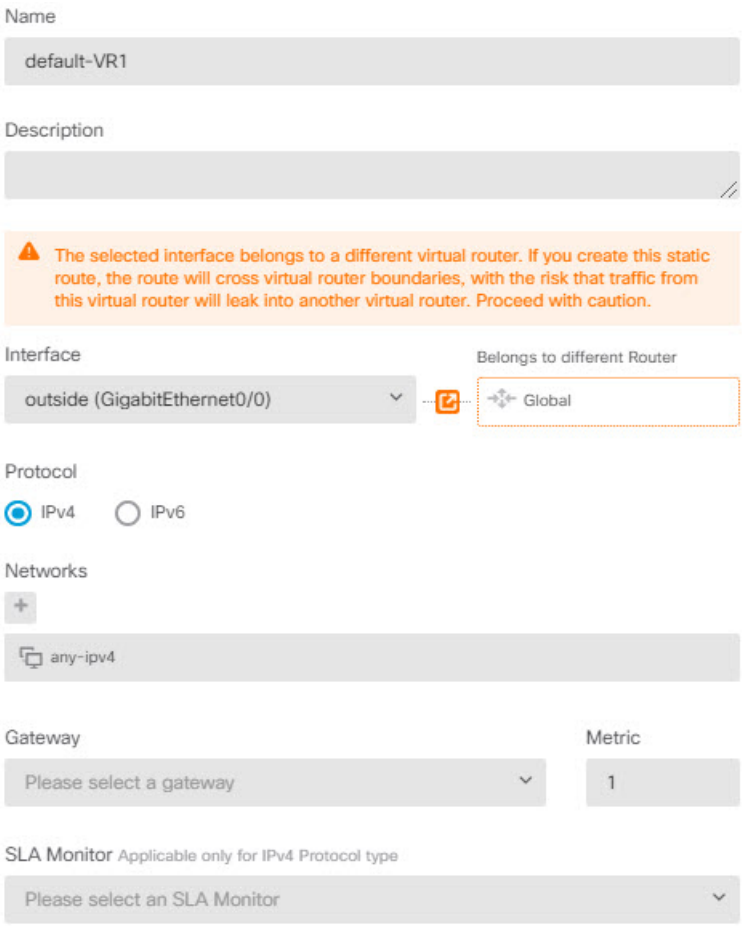
The dialog box closes, and you are shown the list of virtual routers.

f) In the list of virtual routers, click the view icon ( ) in the action column for the VR1 virtual router.

g) On the **Static Routing** tab, click + and configure the route:

- **Name**—Any name will do, such as **default-VR1**.
- **Interface**—Select **outside**. You will see a warning saying that the interface is in a different router and you are creating a route leak. This is what you want to do.
- **Protocol**—For this example, use **IPv4**.
- **Networks**—Select the **any-ipv4** object. This will be the default route for any traffic that cannot be routed within VR1.
- **Gateway**—Leave this item blank. When leaking a route into another virtual router, you do not select the gateway address.

The dialog box should look similar to the following:



h) Click **OK**.

Step 4 Configure virtual router VR2, including the static default route leak to the outside interface.

a) When viewing VR1, click the back button (⇦) to return to the list of virtual routers.

b) Click + at the top of the list.

c) Fill in the properties for virtual router VR2.

- **Name**—Enter VR2 or another name of your choosing.
- **Interfaces**—Click +, select **inside-2**, and click **OK**.



d) Click **OK**.

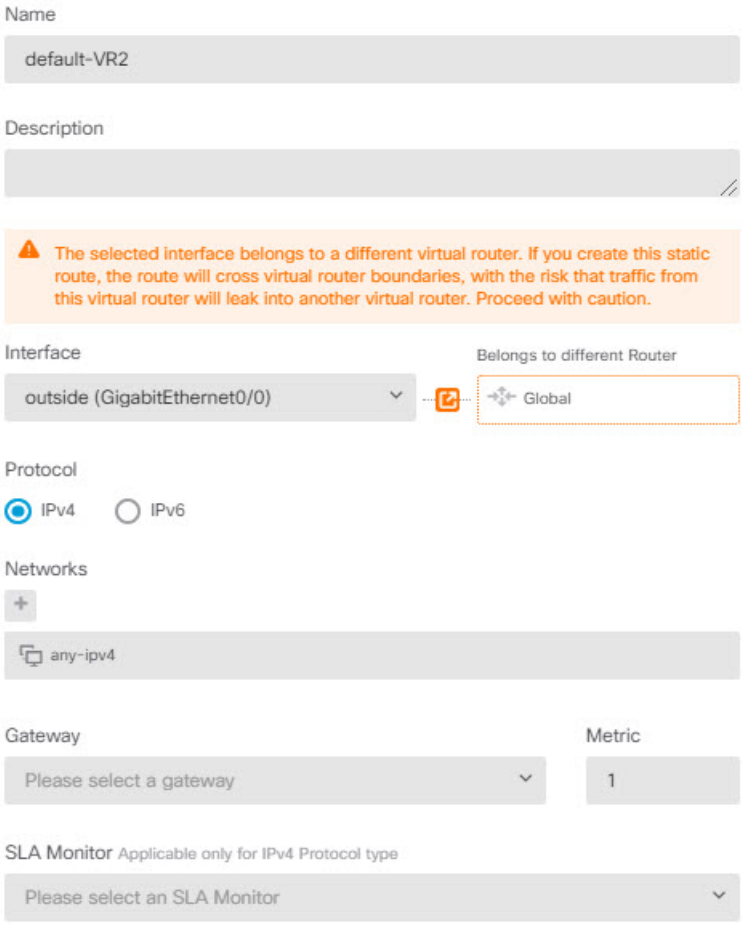
The dialog box closes, and you are shown the list of virtual routers.

e) In the list of virtual routers, click the view icon ( ) in the action column for the VR2 virtual router.

f) On the **Static Routing** tab, click + and configure the route:

- **Name**—Any name will do, such as **default-VR2**.
- **Interface**—Select **outside**. You will see a warning saying that the interface is in a different router and you are creating a route leak. This is what you want to do.
- **Protocol**—For this example, use **IPv4**.
- **Networks**—Select the **any-ipv4** object. This will be the default route for any traffic that cannot be routed within VR2.
- **Gateway**—Leave this item blank. When leaking a route into another virtual router, you do not select the gateway address.

The dialog box should look similar to the following:

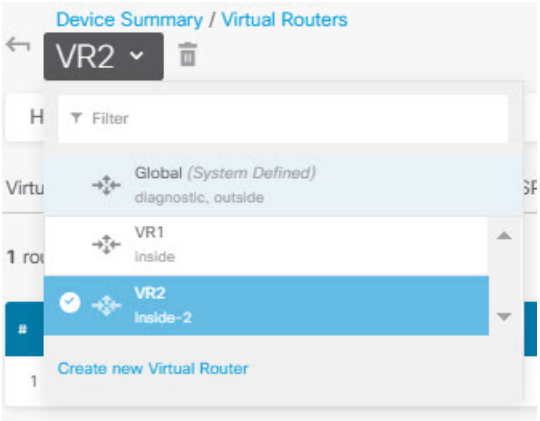


g) Click **OK**.

Step 5 Create the default route in global router to the outside interface.

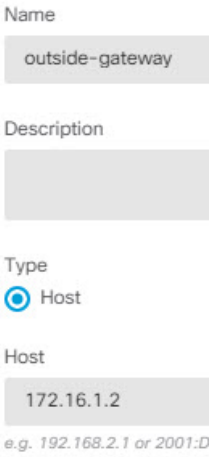
The purpose of this route is to assign the right gateway to traffic leaked from the two virtual routers to the outside interface in the global router.

a) When viewing VR2, click the VR2 name at the top of the page to open the list of virtual routers, and select the Global router.

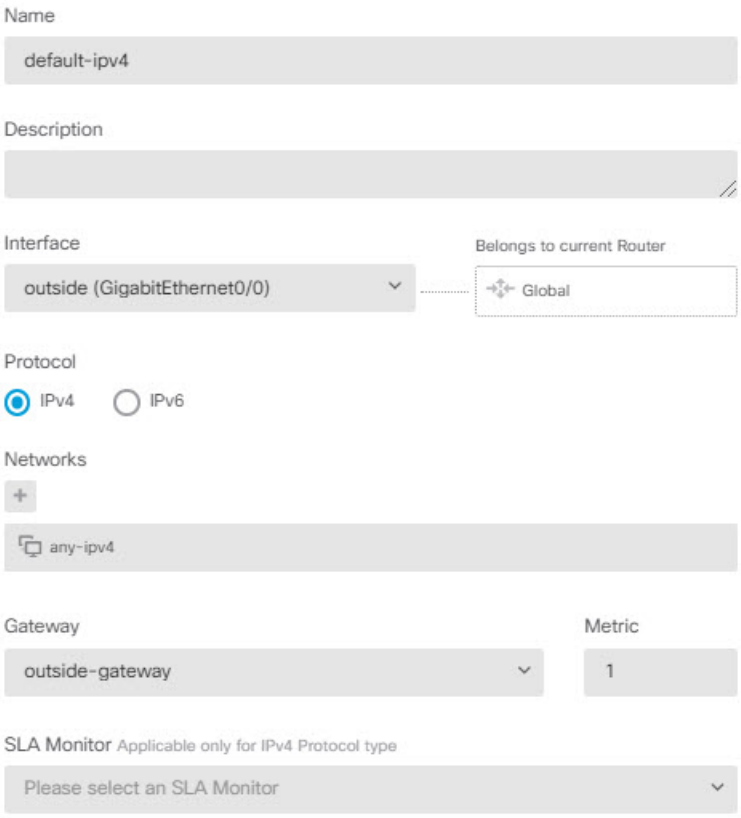


b) On the Static Routing tab for the Global router, click + and configure the route:

- **Name**—Any name will do, such as default-ipv4.
- **Interface**—Select **outside**.
- **Protocol**—For this example, use **IPv4**.
- **Networks**—Select the **any-ipv4** object. This will be the default route for any IPv4 traffic.
- **Gateway**—Assuming the object does not already exist, click **Create New Network Object**, then define a host object for the IP address of the gateway at the other end of the network link on the outside interface, in this case, 172.16.1.2. After you create the object, select it in the Gateway field of the static route.



The dialog box should look similar to the following:



c) Click **OK**.

Step 6 Go back to the **Interfaces** page and add the IP address to inside-2.

a) Click **Device**, then **View All Interfaces** in the **Interface** summary.

b) Click the edit icon ( ) in the Action column for the inside-2 interface that you assigned to VR2.

c) On the **IPv4 Address** tab, enter 192.168.1.1/24 as the IP address and subnet mask.

d) Click **OK**.

You do not get an error for a duplicate IP address this time because the inside and inside-2 interfaces are now in separate virtual routers.

Step 7 Create the NAT rule to PAT inside to outside traffic to 10.100.10.1.

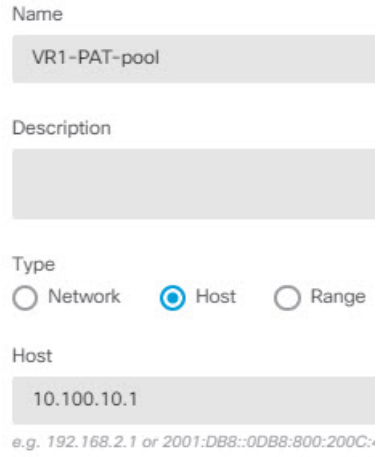
a) Choose **Policies**, then click **NAT**.

b) If there is already a manual NAT rule named InsideOutsideNatRule for the inside to outside interface, applying interface PAT, click the edit icon ( ) for the rule. Otherwise, click + to create a new rule.

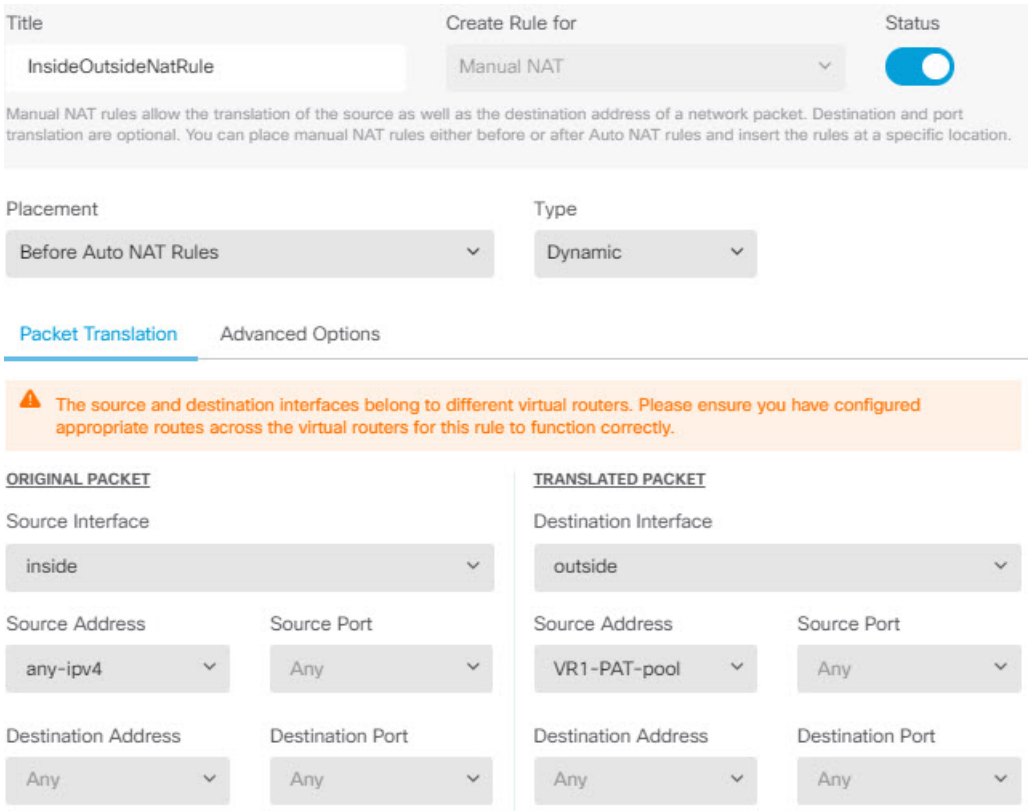
If you edit an existing rule, note that there is now a warning that the source and destination interfaces are in different virtual routers and you need to define routes. This is what you did earlier in the procedure.

c) Assuming you are editing an existing rule, click the drop-down arrow in **Translated Packet** > **Source Address**, and click **Create New Network** (assuming you do not already have a host object defining 10.100.10.1).

d) Configure the host network object for the PAT address. The object should be similar to the following:



e) Select the new object as the **Translated Packet** > **Source Address**. The NAT rule should look similar to the following:



f) Click **OK**.

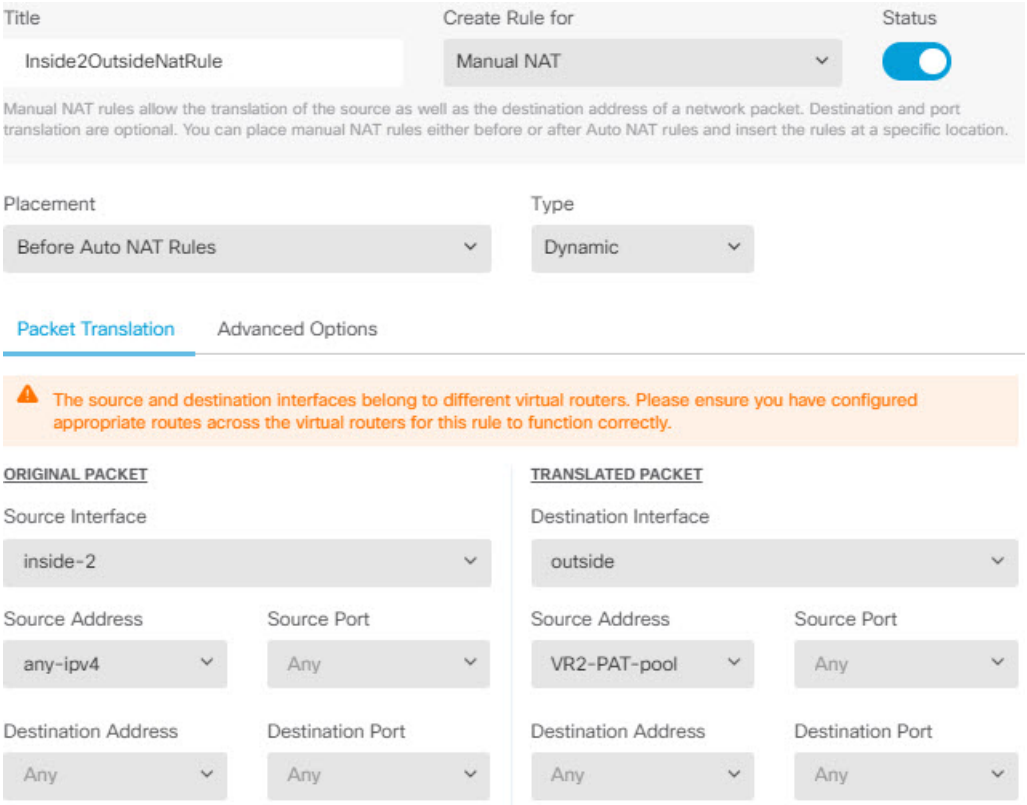
Step 8 Create the NAT rule to PAT inside-2 to outside traffic to 10.100.10.2.

This rule will look exactly like the rule for VR1, with these exceptions:

- **Name**—This must be unique, for example, Inside2OutsideNatRule.
- **Original Packet** > **Source Interface**—Select inside-2.

- **Translated Packet** > **Source Address**—Create a new host network object for 10.100.10.2.

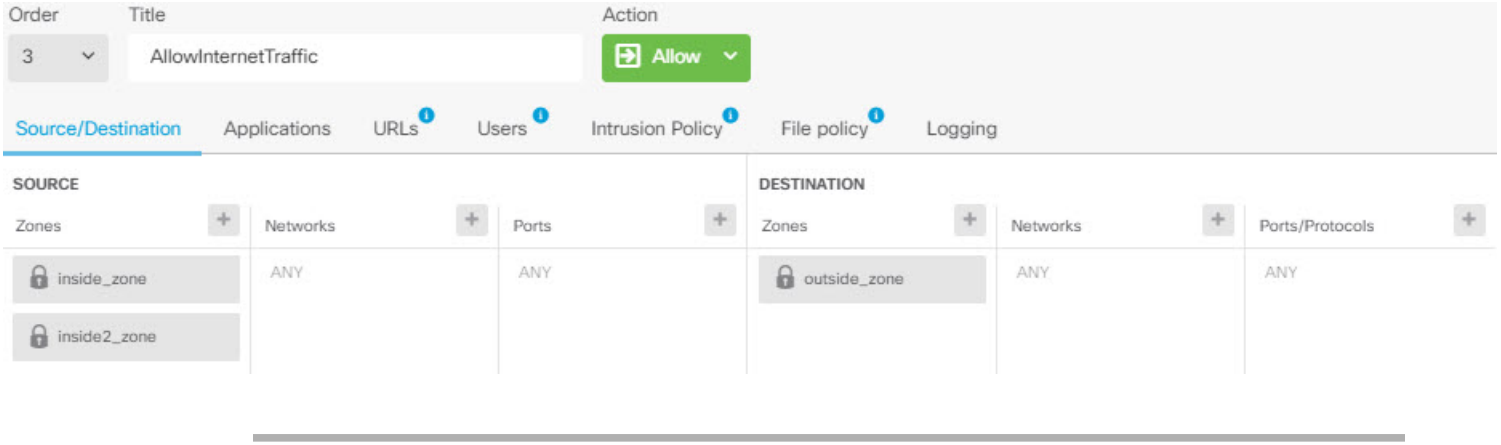
The rule should look similar to the following:



Step 9 Choose **Policies** > **Access Control**, and configure an access control rule to allow traffic from inside_zone and inside2_zone to outside_zone.

Finally, you need to configure the access control policy to allow traffic from the inside and inside-2 interfaces to the outside interface. Because the access control rule requires the use of security zones, you need to create zones for each of these interfaces. Alternatively, you could create a single zone to hold both inside and inside-2, but it is likely that you will want to create additional rules, here or in other policies, that differentiate how traffic is treated in these routers.

Assuming that you create zones named after the interfaces, a basic rule that allows all traffic to flow to the Internet will look like the following. You can apply an intrusion policy to this rule as you see fit. You can define additional rules to block unwanted traffic, for example, to implement URL filtering.



Monitoring Virtual Routers

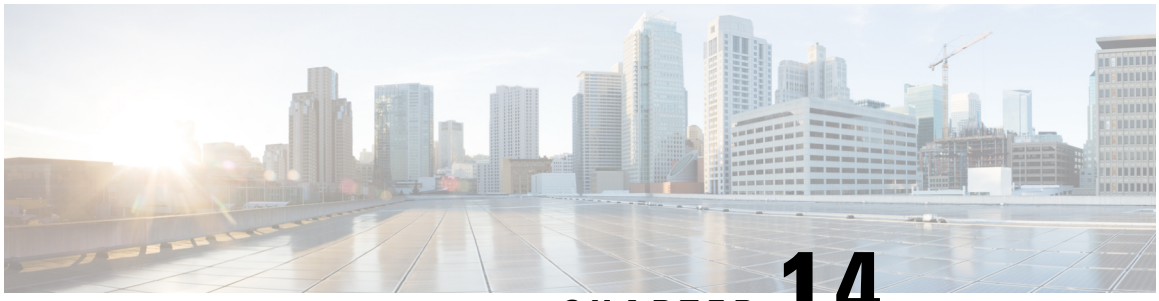
To monitor and troubleshoot virtual routers, open the CLI console or log into the device CLI and use the following commands. You can also select some of these commands from the **Commands** menu on the Routing page.

- **show vrf** displays information about the virtual routers defined on the system.
- **show ospf** [**vrf** *name* | **all**]

  Displays information about the OSPF process in a virtual router. You can specify a virtual router to see information about the process in that virtual router only, or omit the option so see information about VRF across all virtual routers. Use **show ospf ?** to see additional options.
- **show bgp** [**vrf** *name* | **all**]

  Displays information about the BGP process in a virtual router. You can specify a virtual router to see information about the process in that virtual router only, or omit the option so see information about VRF across all virtual routers. Use **show bgp ?** to see additional options.
- **show eigrp** *option*

  Displays information about the EIGRP process. Use **show eigrp ?** to see the available options.



CHAPTER 14

Route Maps and Other Objects for Route Tuning

The various routing protocols let you fine-tune activities such as route distribution and aggregation. For some tuning features, you use route maps or other objects to identify the routes that should be subject to your tuning policy. Route maps have the additional ability to set options on matching routes, so that you can make changes to the route that the next hop router can use to apply custom behavior.

Whether you need to create any of these objects is based on what your needs are for fine-tuning the behavior of the routing protocols that you implement. By first evaluating your requirements, you will determine which types of object you need for the tuning command you want to configure.

- Configure Route Maps, on page 345
- Configure Access Lists, on page 350
- Configure AS Path Access Lists, on page 354
- Configure Community Lists, on page 355
- Configure Policy Lists, on page 357
- Configure Prefix Lists, on page 358

Configure Route Maps

You can use route maps for a variety of purposes, with some routing protocols supporting more uses than others. The most typical use is to fine-tune route redistribution into another routing protocol.

Route Map Permit and Deny Clauses

A route map consists of one or more **permit** or **deny** clauses. The sequence of these clauses matter: routes are evaluated against the map top-down, first match wins. If a route does not match any clause, it is considered to not match the route map.

Each permit clause can contain zero or more **match** and **set** statements. The **match** statement determines which routes match the clause, whereas the **set** statements modify some characteristic of the routes, such as the route metric. You do not need any set statements: you can match a route for redistribution (or another service) without changing the routes in any way.

Each deny clause can contain zero or more match statements. But, because “denied” routes simply do not match the route map, it is pointless to include set clauses, because the set action cannot be applied.

Route Map Match and Set Statements

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 some attribute of the routes.

For example, for each route that is being redistributed, the router first evaluates the match criteria of a clause in the route map. If the route matches the criteria, then the route is redistributed or rejected as dictated by the permit or deny clause. For matches to permit clauses, some of the route's attributes might be modified by the values from the set commands. If the route does not match the criteria, then this clause is not applicable to the route, and the system 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. If there are no matches, the route is considered to not match the route map (equivalent to a deny action).

For the match and set statements in a single clause:

- Multiple match statements are ANDed. That is, a route must satisfy each statement to match the clause.
- Multiple values within a single match statement are ORed. That is, if a route matches any value within that match statement, it is considered to match the statement as a whole.
- If there are no match statements, all routes match the clause.
- If there are no set statements in a route map permit clause, then the feature (such as redistribution) is applied to the route without modification of the route's current attributes.
- Any set statements in a deny clause are ignored. “Denied” routes simply do not match the route map, so it is pointless to include set clauses, because the set action cannot be applied.
- An empty clause, one with no match or set statements, matches any routes that have not been matched by earlier clauses. For example:
  - An empty permit clause allows the redistribution of the remaining routes without modification.
  - An empty deny clause does not allow the redistribution of the remaining routes. This is the default action if a route map is completely scanned, but no explicit match is found.



Note Route-map set statements behave differently depending on where you apply the route map. When you apply a route map to BGP, **set ip next-hop** or **set ipv6 next-hop** changes the BGP NEXT_HOP attribute for matching routes. When you apply a route map to an interface for policy-based routing (PBR), **set ip next-hop** or **set ipv6 next-hop** sets the forwarding next hop for matching packets.

Configure a Route Map

You can use route maps for a variety of purposes, with some routing protocols supporting more uses than others. The most typical use is to fine-tune route redistribution into another routing protocol.

A route map consists of one or more **permit** or **deny** clauses. The sequence of these clauses matter: routes are evaluated against the map top-down, first match wins. If a route does not match any clause, it is considered to not match the route map.

Each permit clause can contain zero or more **match** and **set** statements. The **match** statement determines which routes match the clause, whereas the **set** statements modify some characteristic of the routes, such as the route metric. You do not need any set statements: you can match a route for redistribution (or another service) without changing the routes in any way.

Each deny clause can contain zero or more match statements. But, because "denied" routes simply do not match the route map, it is pointless to include set clauses, because the set action cannot be applied.

For a detailed explanation of how match and set statements are evaluated, carefully read Route Map Match and Set Statements, on page 346.

Before you begin

You can use a variety of other objects in a route map to define match criteria, such as access lists, AS path access lists, community lists, policy lists, and prefix lists. You must create these objects before you create the route map.

For ACL matching, you can use standard or extended ACLs for IPv4 addresses, but extended ACLs only for IPv6. Because the match clauses are based on IPv4 or IPv6 only, ensure your ACLs have the correct address scheme for the match statements.

Also note that the match and set criteria are different for BGP compared to the other routing protocols. Ensure you select the correct match/set criteria for the routing process that will use the route map.

Procedure

Step 1 Click **View Configuration** in **Device** > **Advanced Configuration**.

Step 2 Select **Smart CLI** > **Objects** in the table of contents.

Step 3 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 4 Select **Route Map** as the **CLI Template**.

Step 5 Enter a **Name** for the Smart CLI object. Note that this name is also entered as the route map name in the first line of the CLI template (in the **route-map** command).

Step 6 Create the first clause:

a) Click the *redistribution* variable and select one of the following:

- **permit**—Match. Connections that match this rule are selected for the feature you are configuring.
- **deny**—Do not match. Connections that match this rule are excluded from the feature. Note that "denied" traffic is not dropped, it simply does not get the service applied to it. For example, if you use this route map to define which routes get redistributed, "denied" address spaces are simply not redistributed.

b) Click the **sequence-number** variable and enter the number for the clause, from 1-65535.

This number is relative to the other numbered clauses within the route map. A typical practice is to skip count by 10, that is, 10, 20, 30, to leave room to insert new clauses in the future.

Step 7 Click **Show Disabled** and configure the **match** statements for the clause.

a) Click the + next to the **configure** *clause* command to enable it.

b) Click *clause* and either pick **bgp-match-clause** for BGP route maps, or **match-clause** for all other routing protocols.

c) (BGP route maps.) Configure any combination of the following **match** statements to identify the specific routes you are targeting in this clause. Be sure to click the - icon to disable any commands that you do not configure.

- **match as-path**. Click the variable and select the AS Path objects that define the autonomous system numbers to match.
- **match community**. Click the variable and select the community list objects that define the communities to match.
- **match policy-list**. Click the variable and select the policy list objects that define the match criteria for the clause.
- **match tag**. Click the variable and enter the route tag value to match, from 0-4294967295.

d) (All other routing protocols.) Configure any combination of the following **match** statements to identify the specific routes you are targeting in this clause. Be sure to click the - icon to disable any commands that you do not configure. You might need to click + to enable some of these commands.

- **match interface**. Click the variable and select all of the interfaces in the routes to match.
- **configure match ipv4/ipv6 ip address** *list-type*. Enable the correct command for your IP version. Then, click the *list-type* variable and select whether you want to match the IP address in the route based on **access-list** or **prefix-list**. This will add a **match ipv4/ipv6 address** command, where you can click the variable and select either the access lists or prefix lists that define the IP addresses to match.
- **configure match ipv4/ipv6 ip next-hop** *list-type*. Click the *list-type* variable and select whether you want to match the IP address of the next hop router in the route based on **access-list** or **prefix-list**. This will add a **match ipv4/ipv6 next-hop** command, where you can click the variable and select either the access lists or prefix lists that define the IP addresses to match.
- **configure match ipv4/ipv6 ip route-source** *list-type*. Click the *list-type* variable and select whether you want to match the IP address of the route source in the route based on **access-list** or **prefix-list**. This will add a **match ipv4/ipv6 route-source** command, where you can click the variable and select either the access lists or prefix lists that define the IP addresses to match.
- **match metric**. Click the variable and enter the routing metric to match, from 1-4294967295.
- **match route-type**. (OSPF, EIGRP.) Click the variable and select the route type:
  - **external-1**, **external-2**. OSPF or EIGRP external type-1 or type-2 routes.
  - **internal**. OSPF intra-area and interarea routes or EIGRP internal routes.
  - **local**. Locally generated BGP routes.
  - **nssa-external-1**, **nssa-external-2**. External Not So Stubby Area (NSSA) type-1 or type-2 routes.

Step 8 (Optional, permit clauses only.) For permitted, that is, matched routes, you can configure **set** statements to modify the route attributes. You do not need to modify routes; for example, you can redistribute them unchanged.

a) Click **...** > **Duplicate** to the left of the configure match-clause or **configure bgp-match-clause** command within the permit clause. A new **configure** *clause* command is added at the end of the permit clause.

b) Click *clause* and either pick **bgp-set-clause** or **set-clause**, based on what you picked for the match clause.

c) (BGP route maps.) Configure any combination of the following **set** statements to modify the attributes of matching routes. Be sure to click the - icon to disable any commands that you do not configure.

- **configure set as-path** *options*. Click *options* and select **properties**, which adds the following commands you need to configure. By adding items to the paths, even duplicate AS numbers, you lengthen the path and make the route less likely to be selected as the best route.
  - **set as-path prepend** *as-path*. Click *as-path* and enter up to 10 autonomous system numbers to be added to be beginning of the AS_PATH attribute of the route. The change applies to outbound BGP route maps.
  - **set as-path prepend last-as** *value*. Click *value* and enter how many times the system should prepend the advertising neighbor's autonomous system number to the start of the AS_PATH variable. The change applies to inbound BGP route maps.
  - **set as-path tag**. Converts the tag of a route into an autonomous system path. Applies only when redistributing routes into BGP.
- **set community** *community-number properties*. Click *community-number* and enter the community for the route, from 1-4694967295. Optionally, you can click *properties* and add one of the following:
  - **internet**—Routes with this community are advertised to all peers (internal and external).
  - **no-advertise**—Routes with this community are not advertised to any peer (internal or external).
  - **no-export**—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.
- **set local-preference**. Click the variable and enter the a preference value for the autonomous system path, from 0-4294967295. Unless you change it in the global BGP options, the default preference for BGP routes is 100. The route with the highest preference number is preferred.
- **set weight**. Click the variable and enter the weight for the route, 0-65535. If the router learns about more than one route to the same destination, the route with the highest weight is preferred.
- **set origin** *options*. The origin of a BGP route is based on the path information of the route in the main IP routing table. You can change this by clicking *options* and select how you want to set the BGP origin code.
  - **igp**. Set the origin to the remote Interior Gateway Protocol (IGP) system.
  - **incomplete**. Set the origin as unknown heritage.
- **configure next-hop ipv4/ipv6** *options*. These are separate commands. Select the IPv4 or IPv6 next-hop option. In a BGP route map, this setting changes the BGP NEXT_HOP attribute for matching routes. In a PBR route map, this setting selects the forwarding next hop for matching packets.

- **specific-ip**. Select this option if you want to explicitly set the IP address of the next hop gateway for this route. The **set ip/ipv6 next-hop ip-address** command is added. Click the variable and enter the IP address of the next hop gateway. You can add multiple IP addresses, separated by a space. If the address of the first gateway is unreachable, the next address is tried, and so forth.
- **user-peer-address**. Select this option if you want to set the next hop gateway as the BGP peer's IP address. If you use this option in an outbound route map of a BGP peer, the next hop of the advertised matching routes will be set to be the peering address of the local router, thus disabling the next hop calculation. No additional configuration is required for this command.

- **set ipv4/ipv6 address** *prefix-list*. These are separate commands. Change the IP address of the route based on the contents of the prefix list that you select.
- **set automatic-tag**. Have the system automatically compute a tag value for the route.

d) (All other routing protocols.) Configure any combination of the following **set** statements to modify the attributes of matching routes. Be sure to click the - icon to disable any commands that you do not configure.

- **set metric**. Click the variable and enter the metric value, from 0-4294967295. This value is not used by EIGRP.
- **set metric-type**. Click the variable and select the type of metric:
  - **type-1**, **type-2**. The type of external route in OSPF. Type-2 is the default.
  - **internal**. Sets the Multi Exit Discriminator (MED) value on prefixes advertised to external BGP (eBGP) neighbors to match the Interior Gateway Protocol (IGP) metric of the next hop of the route. This applies to generated, internal BGP (iBGP)-, and eBGP-derived routes.

Step 9 Add permit/deny clauses to complete the route map.

To add a clause, click **...** > **Duplicate** to the left of a **permit** or **deny** line. A new *redistribution sequence-number* clause is added immediately after the clause for which you click the Duplicate command.

Although route map clauses are evaluated in the order of the sequence number rather than the order the appear in the object, it is easier to edit your object if you insert new clauses in sequential order. You cannot move the clauses within the object.

Note that duplicating a clause simply inserts a new, empty clause, with no pre-configured characteristics. After creating the "duplicate," proceed as explained above to configure it for your needs.

Step 10 Click **OK** to save the object.

You can now use the object in a routing process configuration, or in a FlexConfig object, for a feature that requires a route map.

Configure Access Lists

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. 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.
- 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 a "permit" 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 match 10.100.10.1, but exclude the rest of 10.100.10.0/24, the permit entry for 10.100.10.1 must come before the deny entry for 10.100.10.0/24. In general, place more specific rules at the top of an ACL.

Packets that do not match a permit entry are considered to be denied or excluded from the match.

The following topics explain how to configure ACL objects.

Configure Extended Access Lists

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.

Before you begin

Create any network or port objects you will need in the ACEs you create in the object.

Procedure

Step 1 Click **View Configuration** in **Device** > **Advanced Configuration**.

Step 2 Select **Smart CLI** > **Objects** in the table of contents.

Step 3 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 4 Select **Extended Access List** as the **CLI Template**.

Step 5 Enter a **Name** for the Smart CLI object. Note that this name is also entered as the ACL name in the first line of the CLI template (in the **access list** command).

Step 6 Create the ACE that should be the top rule in the ACL.

Each list of commands contained within a single **configure access list entry** command is essentially one ACE, although when deployed, the system might break out the command into a series of ACEs, especially if you include more than one network object.

a) In the configure access list entry command, click *action* and select one of the following:

- **permit**—Match. Connections that match this ACE are selected for the feature you are configuring.
- **deny**—Do not match. Connections that match this ACE are excluded from the feature. Note that "denied" traffic is not dropped, it simply does not get the service applied to it. For example, in a route map, if you use this ACL to define which routes get redistributed, "denied" address spaces are simply not redistributed.

b) In the **permit/deny network** command, click the variables to select the network objects that define the source IP address and the destination IP address of the connection. You can select multiple objects. To specify "any" address, select the any-ipv4 and any-ipv6 objects.

The object or object group cannot contain an FQDN: the objects must specify IP addresses only. FQDN objects are valid in access control rules only.

c) In the **configure permit/deny port** command, click *options* and select one of the following, which will add the associated permit/deny command to the template:

- **any**—If the port does not matter. That is, you are matching any type of IP traffic.
- **any-source**—If the source TCP/UDP port does not matter, but you want to specify the destination port. Click the *destination-port* variable in the **permit/deny port** command and select the port object.
- **any-destination**—If the destination TCP/UDP port does not matter, but you want to specify the source port. Click the *source-port* variable in the **permit/deny port** command and select the port object.
- **source-destination**—If both the source and destination TCP/UDP ports matter. Click the *source-port* and *destination-port* variables in the **permit/deny port** command and select the port objects.

d) In the **configure logging** command, select **disabled**. Logging applies to ACLs that are used for access control, and you cannot use these objects for access control. Thus, the logging options are ignored no matter what you select.

Step 7 Add ACEs to complete the ACL.

To add an ACE, click **...** > **Duplicate** to the left of the **configure access list entry** line. A new ACE group is added immediately after the ACE for which you click the Duplicate command.

Thus, when you have many ACEs in the object, choose wisely which ACE you "duplicate." You cannot move the ACEs within the object, so if you make a mistake, you need to manually recreate the ACE in the correct location.

Note that duplicating an ACE simply inserts a new, empty ACE, with no pre-configured characteristics. After creating the "duplicate," proceed as explained above to configure it for your needs.

Step 8 Click **OK** to save the object.

You can now use the object in a route map object, or in a FlexConfig object, for a feature that requires an extended ACL.

Configure Standard Access Lists

Use standard ACL objects when you want to match traffic based on destination IPv4 address only, and the feature you are configuring supports standard ACLs. Otherwise, use extended ACLs.

Before you begin

Create any network objects you will need in the ACEs you create in the object.

Procedure

Step 1 Click **View Configuration** in **Device** > **Advanced Configuration**.

Step 2 Select **Smart CLI** > **Objects** in the table of contents.

Step 3 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon () for the object.

To delete an unreferenced object, click the trash can icon () for the object.

Step 4 Select **Standard Access List** as the **CLI Template**.

Step 5 Enter a **Name** for the Smart CLI object. Note that this name is also entered as the ACL name in the first line of the CLI template (in the **access list** command).

Step 6 Create the ACE that should be the top rule in the ACL.

Each list of commands contained within a single **configure action** command is one ACE.

a) In the **configure action** command, click *action* and select one of the following:

- **permit**—Match. Connections that match this ACE are selected for the feature you are configuring.
- **deny**—Do not match. Connections that match this ACE are excluded from the feature. Note that "denied" traffic is not dropped, it simply does not get the service applied to it. For example, in a route map, if you use this ACL to define which routes get redistributed, "denied" address spaces are simply not redistributed.

b) In the **permit/deny host** command, click the variable to select the network object that defines the destination IP address of the connection. The object can specify a network or host address. You can select one object per **permit/deny host** command; click **...** > **Duplicate** on the command to specify additional addresses, which will become unique ACEs with the same action. To specify "any" address, select the any-ipv4 object.

Step 7 Add ACEs to complete the ACL.

To add an ACE, click **...** > **Duplicate** to the left of the **configure action** line. A new ACE group is added immediately after the ACE for which you click the Duplicate command.

Thus, when you have many ACEs in the object, choose wisely which ACE you "duplicate." You cannot move the ACEs within the object, so if you make a mistake, you need to manually recreate the ACE in the correct location.

Note that duplicating an ACE simply inserts a new, empty ACE, with no pre-configured characteristics. After creating the "duplicate," proceed as explained above to configure it for your needs.

Step 8 Click **OK** to save the object.

You can now use the object in a route map object, or in a FlexConfig object, for a feature that requires a standard ACL.

Configure AS Path Access Lists

You can use an AS Path access list to filter BGP neighbor updates based on the autonomous system numbers in the updates. Permitted AS numbers have their updates accepted, whereas denied AS numbers have their updates rejected, that is, they are not added to the routing table.

You can also apply AS path filtering in the outbound direction and filter the updates you send to neighbors.

In addition, you can use AS path objects in route maps for BGP address aggregation,

Procedure

Step 1 Click **View Configuration** in **Device** > **Advanced Configuration**.

Step 2 Select **Smart CLI** > **Objects** in the table of contents.

Step 3 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 4 Select **ASPath** as the **CLI Template**.

Step 5 Enter a **Name** for the Smart CLI object. The name must be a number in the 1-500 range. Note that this name is also entered as the AS path access list name in the first line of the CLI template (in the **as-path** command).

Step 6 Configure an AS path entry.

Each entry is contained on a single line starting with the *action* option.

a) Click *action* and select one of the following:

- **permit**—Match. Connections that match this rule are selected for the feature you are configuring.
- **deny**—Do not match. Connections that match this rule are excluded from the feature. Note that "denied" traffic is not dropped, it simply does not get the service applied to it. For example, in a route map, if you use this object to define which routes get redistributed, "denied" address spaces are simply not redistributed.

b) Click *regex* and enter the regular expression that defines the AS numbers that should match this entry.

In its simplest form, the regular expression is just a complete AS path number, and you are either permitting or denying route updates from a single autonomous system.

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. The system supports asplain and asdot notation

as defined in RFC 5396. Which notation you need to use depends on whether you enable the **bgp asnotation dot** command in the BGP global settings.

Step 7 Add entries to complete the AS path access list.

To add an entry, click **...** > **Duplicate** to the left of the *action* line. A new entry is added immediately after the entry for which you click the Duplicate command.

Thus, when you have many entries in the object, choose wisely which entry you "duplicate." You cannot move the entries within the object, so if you make a mistake, you need to manually recreate the entry in the correct location. The rules are evaluated top down, first match wins.

Note that duplicating an entry simply inserts a new, empty entry, with no pre-configured characteristics. After creating the "duplicate," proceed as explained above to configure it for your needs.

Step 8 Click **OK** to save the object.

You can now use the object in a BGP object, route map object, or in a FlexConfig object, for a feature that requires an AS path access list.

Configure Community Lists

If you enable your BGP process to send community information, you can use community lists as a match clause in route maps to set attributes on matching routes. For example, you could change the route preference for certain communities.

A community is an optional attribute or label that a service provider would attach to advertised routes for a group of destinations that share some common attribute. The specific community numbers would be something your ISP would advertise: you would need to obtain the numbers and their meaning from your ISP, and then choose how you want to handle them using a route map.

Community lists are ordered, and matches are determined in a top-down, first match wins method similar to access and prefix lists.

There are two types of community list:

- Standard—Use a standard list when you want to target specific well-known communities, such as the ones obtained from your service provider.
- Expanded—Use an expanded list when you want to match a set of communities based on regular expression matching.

Procedure

Step 1 Click **View Configuration** in **Device** > **Advanced Configuration**.

Step 2 Select **Smart CLI** > **Objects** in the table of contents.

Step 3 Do one of the following:

- To create an object, click the + button.

- To edit an object, click the edit icon () for the object.

To delete an unreferenced object, click the trash can icon () for the object.

Step 4 Select **Standard Community List** or **Expanded Community List** as the **CLI Template**.

Step 5 Enter a **Name** for the Smart CLI object. Note that this name is also entered as the Community List name in the first line of the CLI template (in the **community-list** command).

Step 6 (Standard List.) Configure a community list entry.

Each entry is contained on a single line starting with the *action* option.

a) Click *action* and select one of the following:

- **permit**—Match. Connections that match this rule are selected for the feature you are configuring.
- **deny**—Do not match. Connections that match this rule are excluded from the feature. Note that "denied" traffic is not dropped, it simply does not get the service applied to it. For example, in a route map, if you use this rule to define which routes get redistributed, "denied" address spaces are simply not redistributed.

b) Click *community-number* and enter up to 10 space-separated communities. Multiple communities on a single rule are ANDed, so a match exists only if all communities are matched in the route.

Enter the community in decimal format (1-4294967295) or AA:NN format (each value is 1-66535) based on which numbering method is enabled for your BGP process. Get these numbers from your ISP or other BGP neighbors.

c) (Optional.) Click *properties* and add other well-known communities to the rule.

- **internet**—Routes with this community are advertised to all peers (internal and external).
- **no-advertise**—Routes with this community are not advertised to any peer (internal or external).
- **no-export**—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 7 (Expanded List.) Configure a community list entry.

a) Click *action* and select **permit** or **deny**. These actions are explained above.

b) Click *regex* and enter the regular expression that defines the communities that should match this entry.

The order for matching using the * or + character is longest construct first. Nested constructs are matched from the outside in. Concatenated constructs are matched beginning at the left side. If a regular expression can match two different parts of an input string, it will match the earliest part first. For more information about writing regular expressions, see the "Regular Expressions" appendix of the Cisco IOS Terminal Services Configuration Guide.

Step 8 Add entries to complete the community list.

To add an entry, click **...** > **Duplicate** to the left of the *action* line. A new entry is added immediately after the entry for which you click the Duplicate command.

Thus, when you have many entries in the object, choose wisely which entry you "duplicate." You cannot move the entries within the object, so if you make a mistake, you need to manually recreate the entry in the correct location.

Note that duplicating an entry simply inserts a new, empty entry, with no pre-configured characteristics. After creating the "duplicate," proceed as explained above to configure it for your needs.

Step 9 Click **OK** to save the object.

You can now use the object in a route map or routing process, or in a FlexConfig object, for a feature that requires a community list.

Configure Policy Lists

You can use policy lists in route maps as a replacement for one or more match clauses. Thus, if you have a set of match clauses that you want to reuse, a policy map simplifies your configuration, so you do not need to repeat the match clauses in each route map. You can use route maps that refer to policy lists in BGP.

Within a route map, you can include other match clauses in addition to policy lists. The policy list match clauses match on the incoming attributes only.

Policy lists support matching IPv4 addresses only; you cannot match IPv6 addresses.

For the match clauses in the policy map:

- Multiple match clauses are ANDed. That is, a route must satisfy each clause to match the policy list.
- Multiple values within a single match clause are ORed. That is, if a route matches any value within that match statement, it is considered to match the statement as a whole.

Before you begin

If you will configure match clauses for access lists, prefix lists, or AS path access lists, you must create those objects before you create the policy list.

Procedure

Step 1 Click **View Configuration** in **Device** > **Advanced Configuration**.

Step 2 Select **Smart CLI** > **Objects** in the table of contents.

Step 3 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon () for the object.

To delete an unreferenced object, click the trash can icon () for the object.

Step 4 Select **Policy List** as the **CLI Template**.

Step 5 Enter a **Name** for the Smart CLI object. Note that this name is also entered as the policy list name in the first line of the CLI template (in the **policy-list** command).

Step 6 Click *action* in the **policy-list** command select one of the following:

- **permit**—Match. Connections that match this list are selected for the feature you are configuring.

- **deny**—Do not match. Connections that match this list are excluded from the feature. Note that "denied" traffic is not dropped, it simply does not get the service applied to it. For example, in a route map, if you use this object to define which routes get redistributed, "denied" address spaces are simply not redistributed.

Step 7 Click **Show Disabled** above the template to expose the match commands. You must click the **+** icon to the left of the match statements you want to enable. Configure any combination of the following match statements to define the routes you are targeting.

- **match as-path**. Click the variable and select the AS Path objects that define the autonomous system numbers to match.
- **configure match ip address** *list-type*. Click the *list-type* variable and select whether you want to match the IP address in the route based on **access-list** or **prefix-list**. This will add a **match ip address** command, where you can click the variable and select either the standard access lists or IPv4 prefix lists that define the IP addresses to match.
- **configure match ip next-hop** *list-type*. Click the *list-type* variable and select whether you want to match the IP address of the next hop router in the route based on **access-list** or **prefix-list**. This will add a **match ip next-hop** command, where you can click the variable and select either the standard access lists or IPv4 prefix lists that define the IP addresses to match.
- **configure match ip route-source** *list-type*. Click the *list-type* variable and select whether you want to match the IP address of the route source in the route based on **access-list** or **prefix-list**. This will add a **match ip route-source** command, where you can click the variable and select either the standard access lists or IPv4 prefix lists that define the IP addresses to match.
- **match community** *community-list options*. Click the *community-list* variable and select the community list objects that define the communities to match. If you want routes to match the community list only if all communities in the list are matched, click *options* and select **exact-match**.
- **match interface**. Click the variable and select all of the interfaces in the routes to match.
- **match metric**. Click the variable and enter the routing multi-exit discriminator (MED) metric to match, from 1-4294967295.
- **match tag**. Click the variable and enter the route tag value to match, from 0-4294967295.

Step 8 Click **OK** to save the object.

You can now use the object in a route map object for use in BGP routing.

Configure Prefix Lists

Prefix lists are similar to access control lists. A prefix list is an ordered list of permit/deny rules, where permit indicates address prefixes that should match the list, and deny indicates address prefixes that should not match the list. The system evaluates matches top-down, and assigns the action based on the first matched rule, not necessarily the best matched rule. So, you need to specify sequence numbers carefully to ensure you get the matches you require.

You can use prefix lists for OSPF filtering or BGP, OSPF, or EIGRP route maps for route redistribution or injection, or for BGP neighbor filtering.

There are separate prefix lists for IPv4 and IPv6 addresses, but the structure of the lists are the same.

Procedure

Step 1 Click **View Configuration** in **Device** > **Advanced Configuration**.

Step 2 Select **Smart CLI** > **Objects** in the table of contents.

Step 3 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon () for the object.

To delete an unreferenced object, click the trash can icon () for the object.

Step 4 Select **IPv4 Prefix List** or **IPv6 Prefix List** as the **CLI Template**.

Step 5 Enter a **Name** for the Smart CLI object. Note that this name is also entered as the Prefix List name in the first line of the CLI template (in the **prefix-list** command).

Step 6 Configure a prefix list entry, the **seq** command line.

Each entry is contained on a single line starting with the **seq** option.

a) In **seq**, click *sequence-number* and enter the number for this rule, from 1-4294967294. The number is relative to the sequence numbers for the other rules, with 1 being the first rule evaluated. Common practice is to skip count by 5, that is, 5, 10, 15, etc. This leaves you room to insert new rules without needing to change the sequence numbers of other rules.

b) Click *action* and select one of the following:

- **permit**—Match. Connections that match this rule are selected for the feature you are configuring.
- **deny**—Do not match. Connections that match this rule are excluded from the feature. Note that "denied" traffic is not dropped, it simply does not get the service applied to it. For example, in a route map, if you use this rule to define which routes get redistributed, "denied" address spaces are simply not redistributed.

c) Click *ip-address-mask* and enter the network address and mask (in CIDR format for IPv4) or prefix-length for IPv6. For example, 10.100.10.0/24 (IPv4) or 2001:DB8:0:CD30::/60 (IPv6).

The system uses an exact match for this address/mask unless you also include one of the **ge** or **le** options. For example, 10.100.10.10/8 does not match 10.100.10.0/24 unless you include **ge 9** in the rule.

The mask or prefix length can be:

- IPv4 = 0-32
- IPv6 = 0-128

d) (Optional.) You can use the **ge** and **le** keywords to specify the range of the prefix length to be matched for prefixes that are more specific than the IP address and mask/prefix length. Without these keywords, only exact matches are considered to match the rule.

ge *min-prefix-length* specifies the minimum prefix length to be matched. The minimum must be greater than the mask/prefix length and less than or equal to the maximum defined in the **le** option, if present.

le *max-prefix-length* specifies the maximum prefix length to be matched. The maximum must be greater than or equal to the minimum, if present, or greater than the mask/prefix length if the minimum value is not defined.

Besides the relative length limits mentioned above, the lengths in these options have the following outside limits:

- IPv4 = 1-32
- IPv6 = 0-128

Step 7 Add entries to complete the prefix list.

To add an entry, click **...** > **Duplicate** to the left of a **seq** line. A new entry is added immediately after the entry for which you click the Duplicate command.

For your convenience, it is best to try to keep the entries in sequence order. However, when deployed, the prefix list will be rewritten in the sequential order, even if you have them mixed up in the object.

Note that duplicating an entry simply inserts a new, empty entry, with no pre-configured characteristics. After creating the "duplicate," proceed as explained above to configure it for your needs.

Step 8 Click **OK** to save the object.

You can now use the object in a route map or routing process, or in a FlexConfig object, for a feature that requires a prefix list.

Examples

Following are some examples for how to match prefixes using a prefix list. The sequence number is left out of the examples for simplicity. The actual behavior of each rule is modified by any sequentially earlier rule that matches a subset of the covered address spaces.

- Deny the default route 0.0.0.0/0:

```
deny 0.0.0.0/0
```

- Permit the prefix10.0.0.0/8:

```
permit 10.0.0.0/8
```

- Accept a mask length of up to 24 bits in routes with the prefix 192/8:

```
permit 192.168.0.0/8 le 24
```

- Deny mask lengths greater than 25 bits in routes with a prefix of 192/8:

```
deny 192.168.0.0/8 ge 25
```

- Permit mask lengths from 8 to 24 bits in all address spaces:

```
permit 0.0.0.0/0 ge 8 le 24
```

- Deny mask lengths greater than 25 bits in all address spaces:

```
deny 0.0.0.0/0 ge 25
```

- Deny all routes with a prefix of 10/8:

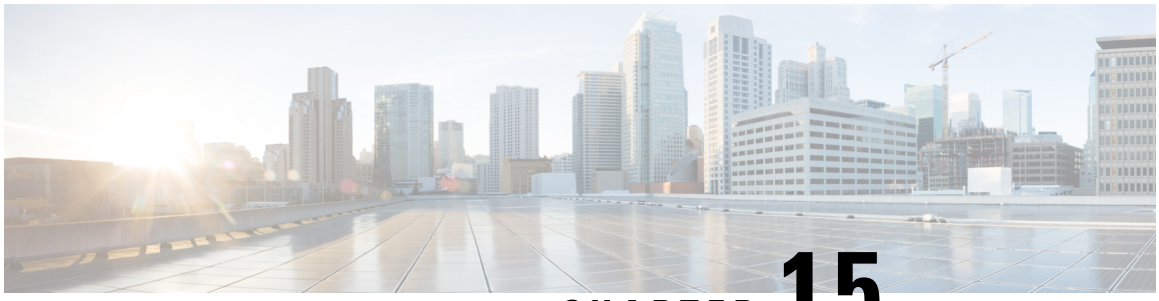
```
deny 10.0.0.0/8 le 32
```

- Deny all masks with a length greater than 25 bits for routes with a prefix of 192.168.1/24:

```
deny 192.168.1.0/24 ge 25
```

- Permit all routes with a prefix of 0/0:

```
permit 0.0.0.0/0 le 32
```



CHAPTER 15

Open Shortest Path First (OSPF)

Open Shortest Path First (OSPF) is a link-state interior gateway protocol. OSPF routers flood link-state information to neighboring routers so that all routers in an OSPF area have a complete view of the network topology.

There are separate OSPF versions based on the IP version: OSPFv2 for IPv4 networks, and OSPFv3 for IPv6 networks. These versions are independent; that is, OSPFv3 is not a replacement for OSPFv2.

You can configure OSPFv2 using Smart CLI objects to integrate your device into the OSPFv2 network topology. You cannot configure OSPFv3.

- Configure the OSPFv2 Process and Areas, on page 363
- Customizing OSPF Process and Area Characteristics, on page 365
- Configure OSPFv2 Interface Settings and OSPF Authentication, on page 377
- Monitoring OSPF, on page 381

Configure the OSPFv2 Process and Areas

You can configure up to 2 OSPFv2 processes. The process numbers are purely internal indicators; they do not need to match any process numbers used on other devices, although you can make the numbers consistent for your own tracking purposes.

If you use private network numbering, such as 192.168.1.0/24, for any internal networks, you might need to segregate private addresses from public addresses, using one OSPFv2 process for these internal networks, and a second process for the external, publicly-addressable networks. Even if you do not use private numbering, 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. If you use NAT, and 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.

Area numbers, in contrast, do exist in the network and you must use the same numbers used by other adjacent routers. If you are configuring a single-area network, use Area 0, also known as the backbone area. For multiple-area networks, where you have a hierarchical network design, you must understand the areas defined in the network, and know which areas this device is supposed to participate in.

If you are using virtual routers, you can configure 2 OSPFv2 processes per virtual router.

The following procedure explains how to create a single OSPFv2 process. Repeat the procedure to create a second process.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring OSPF.

Step 3 Click the **OSPF** tab.

Step 4 Do one of the following:

- To create a new process, click **+** > **OSPF** or click the **Create OSPF Object** > **OSPF** button.
- Click the edit icon ( ) for the object you want to edit. Note that when you edit an object, you might see lines that you did not directly configure. These lines are exposed to show you the default values that are being configured.

If you no longer need a process, click the trash can icon for the object to delete it.

Step 5 Enter a name for the object, and optionally, a description.

Step 6 Configure the basic process properties:

- **router ospf** *process-id*. Click *process-id* and enter a number from 1-65535. This number has meaning within this device only and does not need to match any process numbers configured on other routers. The number must be unique within a virtual router.
- **log-adj-changes** *log-state*. Click *log-state* and select one of the following options:
  - **enable** (recommended)—The system generates a syslog message when an OSPFv2 neighbor goes up or down. If you select this option, an additional **log-adj-changes** *log-type* line is added to the object. Click *log-type* and select **detail** if you want to generate a syslog message for each state change, not just when a neighbor goes up or down.

    If you do not want detailed messages, simply leave *log-type* as the option. Do not delete this line from the object.
  - **disable**—No syslog message is generated. The **no log-adj-changes** line is added to the object: do not delete this line.

Step 7 Click the **Show Disabled** link above the object body to add all other possible configuration lines.

Step 8 Configure the area number.

a) Click the **+** to the left of the **area** *area-id* line to enable the command. You cannot configure a command until you enable it.

b) Click *area-id* and enter the number of the area. This area number needs to be the same as the one used by the other routers that define the OSPFv2 area. You can specify the area ID as either a decimal number or an IP address. Valid decimal values range from 0 to 4294967295.

Step 9 Configure the networks and interfaces that should be routed within the area.

a) Click the **+** to the left of the **configure area** *area-id options* line.

b) Click *area-id* and enter the same area number from the **area** command.

c) Click *options* and select **properties**. This action adds several lines, including one that is enabled by default, the **network** command.

d) In the **network** command, click *network-object* and select the object that defines a network that should be included in this area. Typically, this would be a directly-connected network. For example, if the IP address of the inside interface is 192.168.1.1/24, the associated network object for this command would contain 192.168.1.0/24. If the object does not already exist, click **Create New Network** and create it now.

e) (Optional.) In the **network** command, click *tag-interface* and select the interface that hosts or routes to the network. If you select the interface, the system can prevent you from changing the address on the interface because it is used in the routing process. This helps remind you that any change to interface addressing can impact your routing configuration.

If you select an interface here, before you can change the address on an interface, you must first remove it from the routing process. Then, after the changing the IP address, remember to return here and select the new networks and interface to ensure a correctly-configured routing process.

f) All other new area lines are optional and disabled by default. Configure them only if you need these services. For more information, see Customizing OSPF Process and Area Characteristics, on page 365.

Step 10 If you are configuring the process for a multiple-area network, mouse over the area to the left of the circled **-** on the **area** and **configure area** lines and click **...** > **duplicate**. Then, configure the new area and its networks as explained above. Repeat this process until you have defined all areas in which this routing process should participate.

Step 11 Click **OK**.

Customizing OSPF Process and Area Characteristics

OSPF includes many options that have default values. These values work well for many networks. However, you might need to adjust one or more settings to get the precise behavior you need. The following topics explain the various ways you can customize your OSPFv2 routing process.

Configure Advanced Settings for an OSPF Process

You can configure several settings that control the overall behavior of an OSPFv2 process, including distance metrics, timers, graceful restart, and the router ID used for sending link state advertisements and other routing updates. Many of these settings have defaults that are appropriate for most networks.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring OSPF.

Step 3 Click the **OSPF** tab.

Step 4 Add or edit an OSPF process object.

Step 5 Look for the **setup ospf** line.

When adding an object, you must click the **Show Disabled** link to see the line. Then, click the **+** for the command to enable it, and click *configuration* and select **advanced**. The commands that will be enabled by default are already enabled with their default values.

When editing an object, the line will already be enabled.

The remainder of this procedure assumes you have clicked **Show Disabled**. If you cannot see a command, make sure you expose the disabled commands.

Step 6 (Optional.) Configure the router ID.

Click **+** to enable the **router-id** command, then click the variable and enter the IPv4 address that should be used when sending any router updates from this device. No two routers in an OSPF system can have the same router ID, so ensure that it is unique in the area.

If you do not explicitly specify a router ID for the process, the system uses the highest IP address assigned to an active interface. Thus, the router ID can change if you disable the selected interface, or you change its addresses. By assigning a router ID explicitly, you ensure consistency for your process.

Step 7 (Optional.) Configure RFC 1583 compatibility when calculating summary route costs.

Click **+** to enable the **configure summary-route-cost** command, then click the variable and select either **any**, which turns off RFC 1583 compatibility, or **rfc1583**, which turns it on.

Even though this command is not enabled by default in the OSPF object, in fact RFC 1583 compatibility is the default method used when calculating summary route costs. If you examine the configuration defined in the CLI, only the disabled setting is shown.

Routing loops can occur with RFC 1583 compatibility enabled. Disable it to prevent routing loops. Ensure that you set RFC 1583 compatibility the same on all OSPF routers in an OSPF routing domain.

Step 8 (Optional.) Suppress syslog messages for multicast OSPF (MOSPF) link state advertisements (LSA).

Click **+** to enable the **ignore lsa mospf** command.

The system does not support LSA Type 6 MOSPF packets. You can enable this command so the system does not send syslog messages when it receives these packets, to reduce the noise in your syslog server.

Step 9 Configure the distance metrics.

The following **distance** commands are enabled by default. You can change OSPF route administrative distances based on route type. The distances are 1 to 255, with the higher numbers being less trusted than lower numbers. These metrics are used to judge the relative value of a learned route when comparing similar routes from different processes.

- **distance ospf inter-area 110**. Click the number and set the distance for all routes from one area to another area.
- **distance ospf intra-area 110**. Click the number and set the distance for all routes within an area.
- **distance ospf external 110**. Click the number and set the distance for routes from other routing domains that are learned by redistribution.

Step 10 Configure the route calculation timers for the OSPF process.

The following timer commands are enabled with these default values.

- **timers lsa arrival 1000**. Click the number and set the minimum interval at which the system accepts the same link-state advertisement (LSA) from OSPF neighbors, from 0 to 600000 milliseconds. Use this command to indicate the minimum interval that must pass between acceptance of the same LSA that is arriving from neighbors. LSAs arriving before this minimum time are ignored.

- **timers pacing flood 33**. Click the number and set the time at which LSAs in the flooding queue are paced in-between updates, from 5 to 100 milliseconds.
- **timers pacing lsp-group 240**. Click the number and set the interval at which OSPF link-state advertisements (LSAs) are collected into a group and refreshed, checksummed, or aged, from 10 to 1800 seconds.
- **timers pacing retransmission 66**. Click the number and set the time interval at which LSAs in the retransmission queue are paced, 5 milliseconds to 200 milliseconds. We recommend that you do not change the packet retransmission pacing timer unless all other options to meet OSPF packet flooding requirements have been exhausted. Specifically, configure summarization, stub area usage, queue tuning, and buffer tuning before changing the default flooding timers.
- **timers throttle lsa 0 5000 5000**. Click the numbers and set rate-limiting values for Open Shortest Path First (OSPF) link-state advertisement (LSA) generation. LSA and SPF throttling provide a dynamic mechanism to slow down LSA updates in OSPF during times of network instability and allow faster OSPF convergence. The values are:
  - **Start Interval** (first number)—The minimum delay to generate the first occurrence of an LSA, from 1 to 600000 milliseconds. The first instance of LSA is generated immediately after a local OSPF topology change. The next LSA is generated only after this start interval. Specify 0 to have LSAs generated without delay.
  - **Hold Time** (second number)—The minimum delay to generate the LSA again, from 1 to 600000 milliseconds. This value is used to calculate the subsequent rate limiting times for LSA generation.
  - **Maximum Interval** (third number)—The maximum delay to generate the LSA again, from 1 to 600000 milliseconds.
- **timers throttle spf 5000 10000 10000**. Click the numbers and set rate-limiting values for the shortest path first (SPF) generation. The values are:
  - **Start Interval** (first number)—The delay to receive a change to the SPF calculation, from 1 to 600000 milliseconds.
  - **Hold Time** (second number)—The delay between the first and second SPF calculations, from 1 to 600000 milliseconds.
  - **Maximum Interval** (third number)—The maximum wait time for SPF calculations, from 1 to 600000 milliseconds.

Step 11 (Optional.) Generate a default external route into an OSPF routing domain.

Click + to enable the **default-information originate** command. You can optionally enable and configure the following commands to fine-tune the feature:

- **default-information originate always**. Always advertise a default route even if there is no default route.
- **default-information originate metric 1 metric-type** *metric-type-value*. The metric type and value for generating the default route.
  - Click the **metric** number and enter the OSPF default metric value, from 0 to 16777214. Unless you know you need a different value, enter 10.
  - Click the **metric-type** number and select 1 or 2 as the external link type associated with the default route advertised into the OSPF routing domain. The default is 2.

- **default-information originate route-map** *route-map*. Select a route map that specifies the routing process that generates the default route if the route map is satisfied.

Step 12 (Optional.) Configure Non-Stop Forwarding (NSF) graceful restart if the device is configured for High Availability (HA).

The system can 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 component failure (for example, in HA, the active unit fails over to the standby unit, or in a cluster, the primary unit fails with a secondary unit being elected as new primary), or 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).

You can configure 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.

- You can configure a device as NSF-aware irrespective of the mode in which it operates.
- A device has to be in either High Availability (failover) or Spanned Etherchannel (L2) cluster mode for you to configure it as NSF-capable.

Note
You must not configure the OSPF process to use fast hello packets if you also configure graceful restart. Graceful restart cannot occur with fast hello packets, because the time taken for the role change between the active and standby units is more than the configured dead interval.

To configure graceful restart:

a) Click **+** to enable the **configure nsf graceful-restart** command.
b) Click the *mechanism* variable and select one of the following:
   - **cisco** to configure an NSF-capable device according to Cisco RFC 4811 and RFC 4812.
   - **ietf** to configure an NSF-capable device according to IETF RFC 3623.
   - **both** to configure the device as an NSF-aware helper rather than an NSF-capable device.
   - **none** to disable graceful restart if you have configured it previously.
c) Your selection in the previous step adds the commands required to implement graceful restart according to your specifications. Do not disable these commands. There is only one command that optionally needs further configuration. Following is an explanation of the added commands; a **no** form of the command turns off the related feature.
   - **nsf cisco helper**. Enable Cisco nonstop forwarding (NSF) helper mode. When the NSF-capable Firewall Threat Defense device is performing graceful restart, the helper Firewall Threat Defense devices assist in the nonstop forwarding recovery process.
   - **nsf ietf helper** *mode-option*. Enable IETF nonstop forwarding (NSF) helper mode. When the NSF-capable Firewall Threat Defense device is performing graceful restart, the helper Firewall Threat Defense devices assist in the nonstop forwarding recovery process. Optionally, you can click *mode-option* and enable strict link-state advertisement (LSA) checking. With strict LSA checking enabled, the helper system will terminate the helping process of the restarting system if it detects

that there is a change to an LSA that would be flooded to the restarting system or if there is a changed LSA on the retransmission list of the restarting system when the graceful restart process is initiated.

- **capability lls**. Enables Link Local Signaling (LLS), which is needed for Cisco graceful restart.
- **capability opaque**. Enables opaque Link State Advertisements (LSAs), which is needed for IETF graceful restart.

Step 13 Click **OK**.

Configure OSPF Area Properties

You can configure several OSPF area parameters. You can define the networks to advertise in the area, plus filtering and virtual links. In addition, these area parameters include setting authentication, defining stub areas, and assigning specific costs to the default summary route. Authentication provides password-based protection against unauthorized access to an area.

When you configure area parameters, you need to know how the system functions within the area.

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 link-state advertisements (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 system 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.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring OSPF.

Step 3 Click the **OSPF** tab.

Step 4 Add or edit an OSPF process object.

Step 5 Configure the area number.

a) Click the + to the left of the **area** *area-id* line to enable the command. You cannot configure a command until you enable it.

b) Click *area-id* and enter the number of the area. This area number needs to be the same as the one used by the other routers that define the OSPFv2 area. You can specify the area ID as either a decimal number or an IP address. Valid decimal values range from 0 to 4294967295.

Step 6 Configure the networks and interfaces that should be routed within the area.

a) Click the + to the left of the **configure area** *area-id options* line.

b) Click *area-id* and enter the same area number from the **area** command.

c) Click *options* and select **properties**. This action adds several lines, including one that is enabled by default, the **network** command.

d) In the **network** command, click *network-object* and select the object that defines a network that should be included in this area. Typically, this would be a directly-connected network. For example, if the IP address of the inside interface is 192.168.1.1/24, the associated network object for this command would contain 192.168.1.0/24. If the object does not already exist, click **Create New Network** and create it now.

e) (Optional.) In the **network** command, click *tag-interface* and select the interface that hosts or routes to the network. If you select the interface, the system can prevent you from changing the address on the interface because it is used in the routing process. This helps remind you that any change to interface addressing can impact your routing configuration.

   If you select an interface here, before you can change the address on an interface, you must first remove it from the routing process. Then, after the changing the IP address, remember to return here and select the new networks and interface to ensure a correctly-configured routing process.

Step 7 (Optional.) Configure the cost for the default summary route sent to a stub area or a not-so-stubby area (NSSA).

This option is meaningful only if you configure the area to be a stub or NSSA, as explained below. Click + to enable to following command in the area properties:

area *area-id* **default-cost 1**

If necessary, enter the correct area ID. Then, click the number and enter the relative cost of the route, from 0 to 16777214. The default is 1. The higher the number, the less likely the route will be used over another route that applies for the destination.

Step 8 (Optional.) Configure prefix filtering for the area.

You can filter prefixes advertised in Type 3 link-state advertisements (LSAs) between OSPFv2 areas of an area border router (ABR). Prefix filtering improves your control of route distribution between OSPF areas. With prefix filtering, you can 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.

Before configuring this command, you must create the prefix lists, which are Smart CLI objects, on the **Device** > **Advanced Configuration** page. You can configure separate prefix lists for inbound or outbound advertisements: select the direction for the filter-direction parameter.

area *area-id* **filter-list prefix** *prefix-list filter-direction*

Step 9 (Optional.) Configure the area as a stub 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. You must use default routing in the stub area for it to work properly. To further reduce the number of LSAs sent into a stub area, you can use the **no-summary** keyword of the **area stub** command on the ABR to prevent it from sending a summary link advertisement (LSA Type 3) into the stub area.

To configure the area as a stub:

a) Click the + to the left of the setup area-id as type line.
b) Click *type* and select **stub**. This will add the area stub command after the setup line.
c) Optionally, in the **area stub** command, click *stub-parameters* and select **no-summary**.

Step 10 (Optional.) Configure the area as a not-so-stubby area (NSSA).

A not-so-stubby area (NSSA) is similar to a stub area. NSSA does not flood Type 5 external LSAs from the core into the area, but it can import autonomous system external routes in a limited way within the area.

NSSA imports Type 7 autonomous system external routes within an NSSA area by redistribution. These Type 7 LSAs are translated into Type 5 LSAs by NSSA area border routers (ABRs), which are flooded throughout the whole routing domain. Summarization and filtering are supported during the translation.

You can simplify administration if you are an ISP or a network administrator that must connect a central site using OSPFv2 to a remote site that is using a different routing protocol by running the connecting area as an NSSA. The connection between the corporate site border router and the remote router cannot be run as an OSPFv2 stub area because routes for the remote site cannot be redistributed into the stub area, which means that two routing protocols would have to be maintained. A simple protocol such as RIP would usually be run to handle the redistribution. With NSSA, you can extend OSPFv2 to cover the remote connection by defining the area between the corporate router and the remote router as an NSSA.

Before you use this feature, consider these guidelines:

- You can set a Type 7 default route that can be used to reach external destinations. When configured, the router generates a Type 7 default into the NSSA or the NSSA area boundary router.
- Every router within the same area must agree that the area is NSSA; otherwise, the routers cannot communicate with each other.

To configure the area as an NSSA:

a) Click the **+** to the left of the **setup** *area-id* **as** *type* line.
b) Click *type* and select **nssa**. This will add several commands after the setup line, including the **area nssa** command, which you must leave enabled.
c) (Optional.) To generate a Type 7 default route into the NSSA, click the + to enable the following command:

   area *area-id* **nssa default-information-originate metric 1 metric-type 2**

   You can optionally adjust the following values:

   - Click the **metric** number and enter the OSPF default metric value, from 0 to 16777214. Unless you know you need a different value, enter 10.
   - Click the **metric-type** number and select 1 or 2 as the external link type associated with the default route advertised into the OSPF routing domain. The default is 2.

d) (Optional.) If the system is an ABR, and you want redistribution from other routing protocols to import routes into normal areas only and not into the NSSA, click the **+** to enable the following command:

   area *area-id* **nssa no-redistribution**

e) (Optional.) If you do not want to inject summary routes into the NSSA, click **+** to enable the following command:

   area *area-id* **nssa no-summary**

Step 11 (Optional.) Configure virtual links for the area.

In OSPF, all areas must be connected to a backbone area. If the connection to the backbone is lost, it can be repaired by establishing a virtual link. You can configure virtual links to routers that are connected to the backbone area.

a) Click the **+** to the left of the **configure area** *area-id* **virtual-link** *ip_address option* line.
b) Click *ip_address* and enter the router ID of the router to which you are establishing the virtual link.

c) (Optional.) Click *option* and select **properties** to adjust the following attributes, which all have default values appropriate for most networks. The first part of these commands is omitted, because these are just parameters on the same command:

- **authentication** *auth-type*. Click **+** to enable the command, then click *auth-type* and select **none**, **password**, or **message-digest**. Configure the key options if you select something other than none. The options are the same as those you would configure on an OSPF interface, as explained in Configure OSPFv2 Interface Settings and OSPF Authentication, on page 377. Configure authentication only if the other router uses authentication.
- **hello-interval 10**. Click the number and enter interval between hello packets sent on the interface, from 1 to 65535 seconds.
- **retransmit-interval 5**. Click the number and enter the time between LSA retransmissions for the virtual link, from 1 to 65535 seconds.
- **transmit-delay 1**. Click the number and enter the delay time between when OSPF receives a topology change and when it starts a shortest path first (SPF) calculation, from 0 to 65535 seconds.

d) You can click **...** > **Duplicate** next to the **configure area virtual-link** command to define another virtual link. Define as many as you require.

Step 12 (Optional.) If the system is an area border router (ABR), configure ranges to consolidate or summarize routes for the area.

When you configure the **area range** command, the result is that a single summary route is advertised to other areas by the ABR. Routing information is condensed at area boundaries. External to the area, a single route is advertised for each address range. This behavior is called route summarization. You can configure multiple **area range** commands for an area. In this way, OSPF can summarize addresses for many different sets of address ranges.

To configure route summarization:

a) Click the **+** to the left of the **area** *area-id* **range** *network-object range-parameters* line.
b) Click *network-object* and select the network object that defines the address range whose routes you want summarized.
c) (Optional.) Click *range-parameters* and select one of the following attributes:

- **advertise**. Sets the address range status to advertise and generates Type 3 summary link-state advertisements (LSAs). This is the default if you select no option.
- **not-advertise**. Sets the address range status to DoNotAdvertise. The Type 3 summary LSA is suppressed, and the component networks remain hidden from other networks.

d) You can click **...** > **Duplicate** next to the **area range** command to define another route summarization. Define as many as you require.

Step 13 If you are configuring the process for a multi-area network, mouse over the area to the left of the circled - on the **area** and **configure area** lines and click **...** > **Duplicate**. Then, configure the new area and its networks and other settings as explained above. Repeat this process until you have defined all areas in which this routing process should participate.

Step 14 Click **OK**.

Configure Static OSPF Neighbors

You need to define static OSPF neighbors to advertise OSPF routes over a point-to-point, non-broadcast network, that is, a VPN tunnel.

You do not need to define static neighbors that are on regular broadcast networks, as these routers can form adjacencies themselves.

Before you begin

Determine the interface through which the system should reach the neighbor. You must configure the OSPF settings of this interface before you can define the neighbor router.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring OSPF.

Step 3 Click the **OSPF** tab.

Step 4 Add or edit the OSPF interface object, and enable the **ospf network point-to-point non-broadcast** command for the selected interface. Save your changes.

Step 5 Add or edit an OSPF process object.

Step 6 Click **Show Disabled** to expose all commands, then click **+** to enable the **neighbor** command.

Step 7 Configure the neighbor address.

neighbor *ip-address* **interface** *interface*

- Click *ip-address* and enter the IP address of the neighbor router.
- Click *interface* and select the interface through which the system can reach the router.

Step 8 If necessary, configure a static route for the neighbor router.

If the IP address of the router is on the same network as the selected interface, a static route is not necessary. For example, if you select an interface whose IP address is 10.100.10.1/24, and the neighbor address is 10.100.10.2/24, you do not need a static route.

Step 9 You can click **...** > **Duplicate** next to the **neighbor** command to define another static neighbor. Define as many as you require.

Step 10 Click **OK**.

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 system 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. You can suppress routes that match the specified IP address mask pair. You can use the tag value as a match value for controlling redistribution through route maps.

Route summarization is the consolidation of advertised addresses. You can summarize routes learned from other routing protocols. 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.

Before you begin

Create network objects for all addresses that you want to summarize.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring OSPF.

Step 3 Click the **OSPF** tab.

Step 4 Add or edit an OSPF process object.

Step 5 Click **Show Disabled** to expose all commands, then click **+** to enable the **configure** *network-object* **as** *option* **summary-address** command.

Step 6 Click *network-object* and select the object that defines the address space that you want to summarize.

Step 7 Click *option* and select one of the following:

- **advertising**. Advertise routes that match the address.
- **non-advertising**. Suppress routes that match the address.

Step 8 (Optional.) To add a tag value to the summarized route, click **+** to enable the **summary-address tag** command, click the *tag-number* variable, and enter the tag number, from 0 to 4294967295.

This value is not used by OSPF itself. It may be used to communicate information between autonomous system boundary routers (ASBR). If none is specified, then the remote autonomous system number is used for routes from BGP and EGP; for other protocols, zero (0) is used.

The primary reason to use tag values is for controlling redistribution based on the tag number. If you do not use it in your redistribution route maps, there is no need to configure it here.

Step 9 You can click **...** > **Duplicate** next to the **configure summary-address** command to define another route summarization. Define as many as you require.

Step 10 Click **OK**.

Configure OSPF Filter Rules

Create the Smart CLI standard access list objects you need for each filter rule. Use deny access control entries (ACEs) to filter out routes that match the entry, and permit ACEs for the routes that should be updated.

Before you begin

You can configure area border rouer (ABR) Type 3 LSA filters to allow only specified prefixes to be sent from one area to another area and to 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 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon () for the router in which you are configuring OSPF.

Step 3 Click the **OSPF** tab.

Step 4 Add or edit an OSPF process object.

Step 5 Click **Show Disabled** to expose all commands, then click **+** to enable the **configure filter-rules** *direction* command.

Step 6 Click *direction* and select **in**, for filtering incoming updates, or **out**, for filtering outbound updates.

Step 7 For inbound filters, you can optionally specify the interface on which to filter updates. If you do not specify an interface, the filter applies to all updates received on any interface.

a) Click **+** to enable the **distribute-list** *acl-name* **in interface** *interface* command.

b) Click the *interface* variable and select the interface.

Step 8 For outbound filters, you can optionally specify the protocol, to limit the filter to routes advertised to that routing process.

There are two forms of the **distribute-list out** command, one with an *identifier* variable after the *protocol* variable, and one without the identifier. You can select the following protocols, but they are divided between these command versions based on whether you must provide the additional identifier information.

- **connected**. For routes established for networks that are directly connected to the system's interfaces.
- **static**. For static routes you manually created.
- **rip**. For routes advertised to RIP.
- **bgp** *automonous-system*. For routes advertised to BGP. Click *identifier* and enter the autonomous system number for the BGP process defined on the system.
- **eigrp** *automonous-system*. For routes advertised to EIGRP. Click *identifier* and enter the autonomous system number for the EIGRP process defined on the system.
- **ospf** *process-id*. For routes advertised to OSPF. Click *identifier* and enter the process ID for the other OSPF process defined on the system.

Step 9 You can click **...** > **Duplicate** next to the **configure filter-rules** command to define another filter rule. Define as many as you require.

Step 10 Click **OK**.

Configure OSPF Redistribution

You can control the redistribution of routes into an OSPF process from other routing protocols, connected routes, and static routes.

Before you begin

It is best practice to configure the routing process from which you will redistribute routes, and deploy your changes, before you configure redistribution into OSPF.

If you want to apply a route map to fine-tune which routes are redistributed, create the Smart CLI route map object. Routes that match the route map are redistributed, and all non-matching routes are not redistributed.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring OSPF.

Step 3 Click the **OSPF** tab.

Step 4 Add or edit an OSPF process object.

Step 5 Click **Show Disabled** to expose all commands, then click + to enable the **configure redistribution** command.

Step 6 Click the *protocol* variable and select the source process from which you are redistributing routes. You can redistribute **connected** and **static** routes, or routes generated by **bgp**, **eigrp**, **isis**, **ospf**, or **rip**.

Step 7 If you select a routing process, click the *identifier* variable and enter the required value:

- **bgp**, **eigrp**. Enter the autonomous system number.
- **ospf**. Enter the process ID number.
- **connected**, **static**, **isis**, **rip**. Enter **none**. Even if you enter a different value, it will be ignored.

Step 8 (Optional; IS-IS only.) On the **redistribute isis level-2** command, click **level-2** and select whether you are redistributing routes learned only within an IS-IS area (**level-1**), between IS-IS areas (**level-2**) or both (**level-1-2**).

Step 9 (Optional; all protocols.) If you apply tags to routes in order to control redistribution, click + to enable the **redistribute tag** *tag-number* command, then click the variable and enter the tag associated with routes you want to redistribute. The tag number is from 0 to 4294967295.

Step 10 (Optional; all protocols.) If you want to redistribute routes for all subnets, not just those that abide by the standard class, click + to enable the **redistribute subnets** command.

For example, if you do not enable this command, a specific route for 10.100.10.0/24 would not be redistributed; instead, only a route for 10.0.0.0/8 would be redistributed.

Step 11 (Optional; all protocols.) To fine-tune which routes are redistributed based on a route map, click + to enable the **redistribute** *route-map* command, click the variable, and select the route map that defines your restrictions.

If you do not apply a route map, all routes for the process (that fit the other commands configured for redistribution), are redistributed.

Step 12 (Optional; all protocols.) To fine-tune the metrics for redistributed routes, click + to enable the following command and configure the options:

redistribute *protocol* **metric** *metric-value* **metric-type** *metric-type-value*

Click the variables and configure the following:

- **metric**. The metric value for the routes being distributed, 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 you do not specify a metric value. When redistributing other processes to an OSPF process, the default metric is 20.
- **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. The default is 2.

Step 13 (Optional; OSPF only.) The following commands are enabled by default when you redistribute routes from another OSPF process. You can click - to disable unwanted commands.

These commands specify the criteria by which OSPF routes are redistributed into other routing domains.

- **redistribute ospf match external 1**. Routes that are external to the autonomous system, but are imported into OSPF as Type 1 external routes.
- **redistribute ospf match external 2**. Routes that are external to the autonomous system, but are imported into OSPF as Type 2 external routes.
- **redistribute ospf match internal**. Routes that are internal to a specific autonomous system.
- **redistribute ospf match nssa-external 1**. Routes that are external to the autonomous system, but are imported into OSPF as Type 1 external routes and marked as Not-So-Stubby-Area (NSSA) only.
- **redistribute ospf match nssa-external 2**. Routes that are external to the autonomous system, but are imported into OSPF as Type 2external routes and marked as Not-So-Stubby-Area (NSSA) only.

Step 14 You can click **...** > **Duplicate** next to the **configure redistribution** command to configure redistribution for another protocol. Configure redistribution for each protocol that makes sense for your network.

Step 15 Click **OK**.

Configure OSPFv2 Interface Settings and OSPF Authentication

Any interface that faces a neighbor OSPF router communicates with the router using hello packets and other methods to verify the health of the neighbor and to share routing updates. While some of these characteristics have default settings, it is best practice to set the options explicitly using an OSPF Interface Settings object. Create an object for each interface that is adjacent to an OSPF neighbor router.



Note The routers on a given network must have the same values for authentication and for lost neighbor detection hello and dead intervals.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring OSPF.

Step 3 Click the **OSPF** tab.

Step 4 Do one of the following:

- To create a new object, click **+** > **OSPF Interface Settings** or click the **Create OSPF Object** > **OSPF Interface Settings** button.
- Click the edit icon ( ) for the object you want to edit. Note that when you edit an object, you might see lines that you did not directly configure. These lines are exposed to show you the default values that are being configured.

If you no longer need an interface settings object, click the trash can icon for the object to delete it.

Step 5 Enter a name for the object, and optionally, a description.

Step 6 Configure authentication for the interface.

configure authentication *auth-type*

To configure OSPF authentication, you must configure the password or authentication key on each of the OSPF interfaces, then enable authentication on the area itself. You must choose the same authentication method on the interfaces and the area.

You can select the following options by clicking *auth-type*.

- **none**—Do not use OSPF authentication. Any OSPF router operating on the link can establish an adjacency with this router. The following command is added to the object: **ospf authentication null**.
- **password**—Authenticate the OSPF connection using a shared password. You can configure a separate password to each network on a per-interface basis. However, all neighboring routers on the same network must have the same password to be able to exchange OSPF information.

  When you select this option, two commands are added: **ospf authentication** and **ospf authentication-key** *key*. Click the variable to configure the following:

  - *key*—Select the secret key object that contains the password. The password can be up to 8 characters. You can include blank space between two characters. Spaces at the beginning or end of the password are ignored. If the object does not yet exist, click **Create New Secret Key** at the bottom of the list and create it now.

- **message-digest**—Authenticate the OSPF connection using message digest (MD5). MD5 authentication verifies the integrity of the communication, authenticates the origin, and checks for timeliness. Both routers must be configured to use the same MD5 key.

  When you select this option, two commands are added: **ospf authentication message-digest** and **ospf message-digest-key** *key-id* **md5** *key*. Click the variables to configure the following:

  - *key-id*—The authentication key ID number, from 1 to 255. You must configure the neighbor router with the same key ID and associated MD5 key.

- *key*—Select the secret key object that contains the MD5 key. The key is an alphanumeric password up to 16 characters. You can include spaces between characters. Spaces at the beginning or end of the key are ignored. If the object does not yet exist, click **Create New Secret Key** at the bottom of the list and create it now.

Step 7 (Optional.) Configure link-state advertisement (LSA) timers.

These timers have default values, so you need to change them only if your network requires different settings. Configure the following commands:

- **ospf retransmit interval 5**—The number of seconds between LSA retransmissions for adjacencies belonging to an OSPF interface. The number of seconds must be greater than the expected round-trip delay between any two routers on the attached network. The range is from 1 to 8192 seconds. The default value is 5 seconds. Click the 5 and type in a new number to change the value.
- **ospf transmit-delay 1**—The estimated number of seconds required to send a link-state update packet on an OSPF interface, from 1 to 8192 seconds. The default value is 1 second. Click the 1 and type in a new number to change the value.

Step 8 (Optional.) All other settings have default values or they are optional. Change them, or enable them, only if you need different behavior. Click the **Show Disabled** link to expose the options.

Following are the additional interface settings. To enable a setting, click the **+** to the left of the command, then configure the command (if necessary).

- **ospf cost** *value*—The cost (a link-state metric) of sending a packet on an OSPF interface, from 1 to 65535. The value 1 represents a network that is directly connected to the interface. Click the variable and enter the cost that represents the capability of the interface based on the numbers you are using in your network.

  When deciding on a value, the higher the interface bandwidth, the lower the associated cost to send packets across that interface should be. In other words, a large cost value represents a low bandwidth interface and a small cost value represents a high bandwidth interface. The specific number you select has no inherent meaning: the value is relative to the other values you configure for interface across the OSPF area. These values then affect the calculation of the best route for a destination.

  The OSPF interface default cost on the Firewall Threat Defense device is 10. This default differs from Cisco IOS software, where the default cost is 1 for Fast Ethernet and Gigabit Ethernet and 10 for 10BaseT. This is important to take into account if you are using ECMP in your network.
- **ospf database-filter all out**—Filters out all outgoing LSAs to an OSPF interface during synchronization and flooding.
- **ospf mtu-ignore**—Disables OSPF maximum transmission unit (MTU) mismatch detection on receiving database packets. OSPF checks whether neighbors are using the same MTU on a common interface. This check is performed when neighbors exchange Database Descriptor (DBD) packets. If the receiving MTU in the DBD packet is higher than the MTU configured on the incoming interface, OSPF adjacency will not be established. If you cannot fix the MTU values on the interfaces to be the same, you can disable MTU checking.
- **ospf network point-to-point non-broadcast**—Configures the OSPF interface as a point-to-point, non-broadcast network. This lets you transmit OSPF routes over VPN tunnels. If you configure this option, dynamic discover of neighbors is not possible. You must also:

- Update the OSPF process object to define a single static neighbor for this interface. Also, update the OSPF process of the neighbor router to define this device as its static neighbor.
- Create static routes (on each router) that point to the neighbor router.

- **ospf priority** *value*—The priority of the router relative to the other routers in the network, from 0 to 255. The default priority is 1. When two routers attached to a network both attempt to become the designated router, the one with the higher router priority takes precedence. If there is a tie, the router with the higher router ID takes precedence. A router with a router priority set to zero is ineligible to become the designated router or backup designated router. Click the variable and select the priority based on the relative numbering system you use in your network.
- **ospf lost-neighbor-detection** *detection-mechanism*—Defines how the system determines if a neighbor router is down. OSPF must recalculate routes whenever an OSPF router is declared down. For detailed information on configuring lost neighbor detection, see Configure OSPFv2 Lost Neighbor Detection and Fast Hello Packets (OSPF Interface Settings), on page 380.

Step 9 Click **OK**.

Configure OSPFv2 Lost Neighbor Detection and Fast Hello Packets (OSPF Interface Settings)

The OSPF process regularly sends hello packets to each neighbor router to verify that the neighbor can still respond. Continued failure to respond indicates that the neighbor router (either entirely or just the adjacent interface) is not available for routing, and OSPF must recalculate routes and the OSPF system must converge on the updated routing table.

You can adjust the following values to fine-tune your network. Ideally, you want to minimize how often neighbors are declared down and routes recalculated. On the other hand, you also want to minimize how long it takes for the network to reconverge on a good routing table when an OSPF router (or interface) truly is down.

- Hello interval—This is the time between sending hello packets. The default is every 10 seconds. If desired, you can configure fast hello packets, where hellos are sent at sub-second intervals. Fast hello packets provide the quickest detection of a down neighbor and reconvergence of the routing table.
- Dead interval—The length of time during which, if no hello packets are seen from a neighbor, the neighbor is declared dead. The default is 40 seconds (4 times the default hello interval), unless you are using fast hello packets, in which case the dead interval is always 1 second. Specifying a smaller dead interval will give faster detection of a neighbor being down and improve convergence, but might cause more routing instability. In any case, you must configure the dead interval to be larger than the hello interval. You must set the same dead interval on all OSPF routers in the network.

You configure lost neighbor detection in the OSPF Interface Settings object.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring OSPF.

Step 3 Click the **OSPF** tab.

Step 4 Do one of the following:

- To create a new object, click **+** > **OSPF Interface Settings** or click the **Create OSPF Object** > **OSPF Interface Settings** button.
- Click the edit icon ( ) for the object you want to edit. Note that when you edit an object, you might see lines that you did not directly configure. These lines are exposed to show you the default values that are being configured.

Step 5 If the **ospf lost-neighbor-detection** *detection-mechanism* command is not displayed, click the **Show Disabled** link.

Step 6 Click the **+** to the left of the command to enable it.

Step 7 Click *detection-mechanism* and select the mechanism that you want to implement:

- **dead-interval**—To configure a standard hello interval in seconds. The following commands are added; adjust their values as needed:
  - **ospf hello-interval 10**—The hello interval, from 1 to 8192 seconds. The default is 10. This value must be less than the dead interval. Click the value to enter the desired number.
  - **ospf dead-interval 40**—The dead interval, from 1 to 8192 seconds. The recommended value is 4 times the hello interval, but you can configure a shorter time for faster convergence.
- **hello-multiplier**—To configure sub-second fast hello packets. The following command is added, you must configure the value.
  - **ospf dead-interval minimal hello-multiplier** *value*—Click the variable and enter the number of hello packets that should be sent each second, from 3 to 20. The dead interval is set to 1 second by the **minimal** keyword.

Step 8 Click **OK**.

Monitoring OSPF

To monitor and troubleshoot OSPF, open the CLI console or log into the device CLI and use the following commands. You can also select some of these commands from the **Commands** menu on the Routing page.

Use **show ospf ?** to get lists of additional options. For example, you can specify process ID, area ID, and virtual router to limit the information you see, as well as other options to target just the information you are looking for. The following list is a summary only.

- **show ospf**

  Displays general information about OSPFv2 routing processes.

- **show ospf border-routers**

  Displays the internal OSPFv2 routing table entries to the ABR and ASBR.

- **show ospf database**

  Displays lists of information related to the OSPFv2 database for a specific router.

- **show ospf events**

  Displays OSPF internal event information.

- **show ospf flood-list**

  Displays a list of LSAs waiting to be flooded over an interface, to observe OSPF v2packet pacing. OSPFv2 update packets are automatically paced so they are not sent less than 33 milliseconds apart. Without pacing, some update packets could get lost in situations where the link is slow, a neighbor could not receive the updates quickly enough, or the router could run out of buffer space.

  Pacing is also used between resends to increase efficiency and minimize lost retransmissions. You also can display the LSAs waiting to be sent out of an interface. Pacing enables OSPFv2 update and retransmission packets to be sent more efficiently.

- **show ospf interface**

  Displays OSPFv2-related interface information.

- **show ospf neighbor**

  Displays OSPFv2 neighbor information on a per-interface basis.

- **show ospf nsf**

  Displays the OSPFv2 related Non-Stop Forwarding (NSF) information.

- **show ospf request-list**

  Displays a list of all LSAs requested by a router.

- **show ospf retransmission-list**

  Displays a list of all LSAs waiting to be resent.

- **show ospf rib**

  Displays the OSPF Router Information Base (RIB).

- **show ospf statistics**

  Displays various OSPF statistics, such as the number of times SPF was executed, the reasons, and the duration.

- **show ospf summary-addresses**

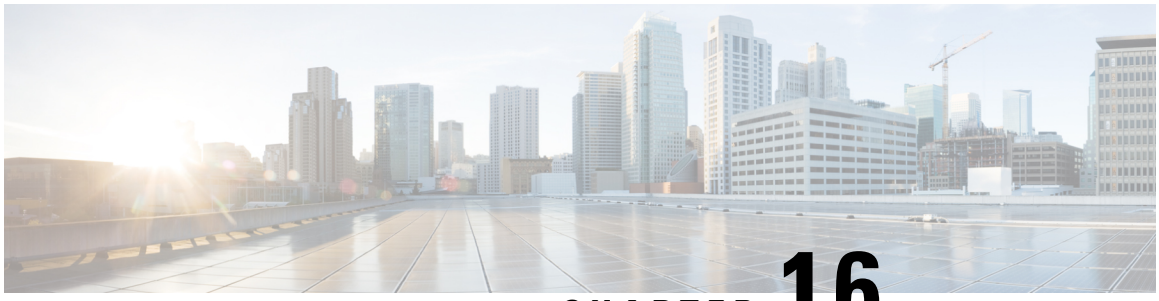
  Displays a list of all summary address redistribution information configured under an OSPFv2 process.

- **show ospf traffic**

  Displays a list of different types of packets being sent or received by a specific OSPFv2 instance.

- **show ospf virtual-links**

  Displays OSPFv2-related virtual links information.



CHAPTER 16

Enhanced Interior Gateway Routing Protocol (EIGRP)

The Enhanced Interior Gateway Routing Protocol (EIGRP) is a hybrid dynamic distance-vector and link-state interior gateway routing protocol. Originally a proprietary protocol developed by Cisco, it is now an open standard defined in RFC 7868. You can configure EIGRP to manage interior routes within an autonomous system.

- Best Practices for EIGRP, on page 383
- About EIGRP, on page 384
- Guidelines for EIGRP, on page 385
- Configure the Core EIGRP Process, on page 386
- Customize the EIGRP Process, on page 389
- Monitoring EIGRP, on page 399

Best Practices for EIGRP

Following are some tips for configuring EIGRP:

- If you are inserting the device into an existing EIGRP autonomous system, examine the configuration of the other routers in the autonomous system to determine the system number and any other customizations. Ensure that you implement the same, or at least consistent, customizations on the Firewall Threat Defense device that you are adding.
- Determine whether to configure a full EIGRP process or a stub process:
  - If the Firewall Threat Defense device is in the middle of the autonomous system, where it is connected to more than one other EIGRP router, you probably need the full EIGRP process. See Configure the EIGRP Process for Full Routing, on page 386.
  - If the Firewall Threat Defense device is at the edge of the autonomous system, where it is connected to just one other EIGRP router, and otherwise hosts connected networks only, then configuring it as a stub router might make the most sense. You can configure the stub so that the Firewall Threat Defense device sends information to the EIGRP neighbor about connected routes, so that other EIGRP routers in the autonomous system can get routes to the Firewall Threat Defense device's connected networks. See Configure the EIGRP Process for Stub Routing, on page 387.

- The default settings work for most networks, so adjust them only if you have adjusted them on other EIGRP routers in the autonomous system. You can have a fully-functional EIGRP process by simply configuring the autonomous system number and specifying the networks to route.
- Configure the router ID to ensure a stable address is used to identify the router. This can help make troubleshooting routing problems easier. See Configure EIGRP Advanced Settings, on page 389.
- Do not enable automatic route summarization (the **auto-summary** command) unless you have determined that it will not create routing loops and it will provide some benefit for your network. How to determine whether auto summarization will work for your network is outside the scope of this document.

About EIGRP

The Enhanced Interior Gateway Routing Protocol (EIGRP) is a hybrid dynamic distance-vector and link-state interior gateway routing protocol. EIGRP sends routing updates to routers within the same autonomous system. Normally, EIGRP discovers neighbor routers using multicast updates, but you can configure static neighbors that are outside of a multicast boundary, and these static neighbors get unicast updates.

The convergence technology of EIGRP is based on an algorithm called the Diffusing Update Algorithm (DUAL). The algorithm guarantees loop-free operation at every instant throughout a route computation and allows all devices involved in a topology change to synchronize. Devices that are not affected by topology changes are not involved in recomputations.

To a limited extent, you can adjust routing metrics to control how routes are chosen. The following topics provide some background on these advanced concepts.



Note If you adjust these metrics, you must make the same adjustment to all routers within the autonomous system or you might get routing loops.

DUAL Finite State Machine

The DUAL finite state machine embodies the decision process for all route computations. It tracks all routes advertised by all neighbors. DUAL uses the distance information (known as the metric) to select efficient, loop-free paths.

DUAL selects routes to be inserted into a routing table based on feasible successors. A successor is a neighboring device (used for packet forwarding) that has the least-cost path to a destination that is guaranteed not to be part of a routing loop.

When a topology change occurs, DUAL will test for feasible successors. If there are feasible successors, DUAL will use any feasible successors it finds to avoid unnecessary recomputation.

When there are no feasible successors but only neighbors advertising the destination, a recomputation must occur to determine a new successor. The time required to recompute the route affects the convergence time.

EIGRP Metric Weights

EIGRP uses metric weights, known as K values, in routing and metric computations. EIGRP metric defaults have been carefully selected to provide optimal performance in most networks.

Unlike IOS routers, you cannot adjust these default K values for EIGRP running on the Firewall Threat Defense device. Because you need to use the same K values on all systems within an autonomous network, you should not change them on any routers within an autonomous system that includes Firewall Threat Defense devices.

For an explanation of how the K values are used, see EIGRP Cost Metrics, on page 385.

EIGRP Cost Metrics

EIGRP uses metric weights (K values) in addition to link characteristics to compute a composite cost metric. To avoid churn in the network as a result of a change in the link characteristics, you can tune some of the values used in this computation.

The actual computation is quite complex, using five K values (as multipliers) and five vector attributes. However, three K values are 0 by default, and because you cannot change the K value defaults, the actual computation is greatly simplified:

cost metric = 256 * (bandwidth + delay)

What you can change is the bandwidth and delay values for routes redistributed from or to the EIGRP process. Specifically, you can adjust these values in the **default-metric** command (to set a default for all types of redistributed routes) or in the **redistribute metric** command (to set the metrics for a specific type of route). Note the following:

- Bandwidth is the minimum bandwidth of the route, in kilobits per second. It can be 1 to 4294967295 Kilobytes per second. The bandwidth for the formula is scaled and inverted by the following formula:

  $(10^7$/minimum bandwidth in kilobits per second)

- Delay is the route delay, in tens of microseconds.

The other characteristics, which are not used by Firewall Threat Defense, are the delay reliability, the effective load on the route, and the minimum MTU (maximum transmission unit) on the route. Even though the values are not used, you must configure them if you adjust these commands.

For more detailed information on how EIGRP calculates the cost metric, see *IP Routing: EIGRP Configuration Guide*. For example, https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/xe-16-7/ire-xe-16-8-book/ire-enhanced-igrp.html.

Guidelines for EIGRP

IPv6 Guidelines

Does not support IPv6.

Additional Guidelines

- A maximum of one EIGRP process is supported.
- You cannot change the autonomous system number for an EIGRP process. Instead, delete the process, deploy your changes, then configure a new process using the new autonomous system number.
- You cannot include a network in the EIGRP process that belongs to a Bridge Virtual Interface (BVI).

- EIGRP adjacency flap occurs whenever a configuration change is applied which results in modifying the routing information (sent or received) from neighbors especially in distribute lists, offset lists, and changes to summarization. After the routers are synchronized, EIGRP reestablishes the adjacency between neighbors. When an adjacency is torn down and reestablished, all learned routes between the neighbors are erased and the entire synchronization between the neighbors is performed newly with the new distribute list.

Configure the Core EIGRP Process

The following topics explain how to get EIGRP up and running on the device. You can either configure the full routing process, or configure it as a stub process for systems that should not fully participate in the autonomous network as an EIGRP router.

Configure the EIGRP Process for Full Routing

You can configure one EIGRP process. If you configure multiple virtual routers, EIGRP is supported in the global virtual router only.

The following procedure sets up basic EIGRP routing for a set of networks using all default values for EIGRP routing. Completing this procedure is sufficient to enable EIGRP on the device. You can complete the other procedures to fine-tune the EIGRP process as needed.

Before you begin

Determine what autonomous system number you are using for EIGRP in your network.

Create network objects that define each of the networks that you want to route within the EIGRP autonomous system. For example, if you want to use EIGRP for the 192.168.1.0/24 and 192.168.2.0/24 networks, create two network objects, one for each network.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the global virtual router.

Step 3 Click the **EIGRP** tab.

Step 4 Do one of the following:

- To create a new process, click **+** or click the **Create EIGRP Object** button.
- Click the edit icon ( ) for the object you want to edit. Note that when you edit an object, you might see lines that you did not directly configure. These lines are exposed to show you the default values that are being configured.

If you no longer need a process, click the trash can icon for the object to delete it.

Step 5 Enter a **Name** for the Smart CLI object and optionally, a description.

Step 6 Configure the basic process properties:

router eigrp *autonomous-system*

Click the variable and enter a number from 1-65535. Use the same autonomous system number used on the other routers in the network that should operate within the same routing domain as this device.

Step 7 Configure the networks and interfaces that should be routed within the EIGRP autonomous system.

a) Click the **Show Disabled** link above the object body to add all other possible configuration lines.

b) Click the + to the left of the **network** *network-object* line.

c) In the **network** command, click the variable and select the object that defines a network that should be included in this autonomous system.

Typically, this would be a directly-connected network. For example, if the IP address of the inside interface is 192.168.1.1/24, the associated network object for this command would contain 192.168.1.0/24. If the object does not already exist, click **Create New Network** and create it now.

Directly connected and static networks that fall within the defined network are advertised by the process. Additionally, only interfaces with an IP address that fall within the defined network participate in the EIGRP routing process.

If you have an interface that you do not want to have participate in EIGRP routing, but that is attached to a network that you want advertised, see Configure EIGRP Passive Routing Interfaces, on page 392.

d) If you have additional networks to route, click **...** > **Duplicate** to the left of the **network** command to add a new one. Keep adding **network** lines until you have configured all the networks to route.

Step 8 (Optional.) If necessary, adjust settings for the other initially disabled commands. See Customize the EIGRP Process, on page 389.

Step 9 Click **OK**.

Configure the EIGRP Process for Stub Routing

You can configure the device to be an EIGRP stub router. Stub routing decreases memory and processing requirements on the system. As a stub router, the system does not need to maintain a complete EIGRP routing table because it forwards all non-local traffic to a distribution router. Generally, the distribution router need not send anything more than a default route to the stub router.

Only specified routes are propagated from the stub router to the distribution router. As a stub router, the system responds to all queries for summaries, connected routes, redistributed static routes, external routes, and internal routes with the message “inaccessible.” The system sends a special peer information packet to all neighboring routers to report its status as a stub router. Any neighbor that receives a packet informing it of the stub status will not query the stub router for any routes, and a router that has a stub peer will not query that peer. The stub router depends on the distribution router to send the correct updates to all peers.

Before you begin

Determine what autonomous system number you are using for EIGRP in your network.

Create network objects that define each of the networks that you want to route within the EIGRP autonomous system. For example, if you want to use EIGRP for the 192.168.1.0/24 and 192.168.2.0/24 networks, create two network objects, one for each network.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the global virtual router.

Step 3 Click the **EIGRP** tab.

Step 4 Do one of the following:

- To create a new process, click + or click the **Create EIGRP Object** button.
- Click the edit icon ( ) for the object you want to edit. Note that when you edit an object, you might see lines that you did not directly configure. These lines are exposed to show you the default values that are being configured.

If you no longer need a process, click the trash can icon for the object to delete it.

Step 5 Enter a **Name** for the Smart CLI object and optionally, a description.

Step 6 Configure the basic process properties:

router eigrp *autonomous-system*

Click the variable and enter a number from 1-65535. Use the same autonomous system number used on the other routers in the network that should operate within the same routing domain as this device.

Step 7 Configure the networks and interfaces that should be routed within the EIGRP autonomous system.

a) Click the **Show Disabled** link above the object body to add all other possible configuration lines.

b) Click the + to the left of the **network** *network-object* line.

c) In the **network** command, click the variable and select the object that defines a network that should be included in this autonomous system.

Typically, this would be a directly-connected network. For example, if the IP address of the inside interface is 192.168.1.1/24, the associated network object for this command would contain 192.168.1.0/24. If the object does not already exist, click **Create New Network** and create it now.

Directly connected and static networks that fall within the defined network are advertised by the process. Additionally, only interfaces with an IP address that fall within the defined network participate in the EIGRP routing process.

If you have an interface that you do not want to have participate in EIGRP routing, but that is attached to a network that you want advertised, see Configure EIGRP Passive Routing Interfaces, on page 392.

d) If you have additional networks to route, click **...** > **Duplicate** to the left of the **network** command to add a new one. Keep adding **network** lines until you have configured all the networks to route.

Step 8 Configure the stub settings.

a) Click the + to the left of the **setup eigrp** *configuration* line.

b) Click the variable and select **advanced**.

c) Click the + to the left of the **setup eigrp stub** *stub-options* command.

d) To restrict the device from sharing any of its routes with any other router in the autonomous system, so that it only receives updates from the EIGRP neighbor router, click *stub-options* and select **receive**. Then, configure the following command:

eigrp stub *stub-parameters*

Click the variable and select **receive-only**.

e) To allow the device to advertise routes to EIGRP neighbor routers, click *stub-options* and select **other**. Then, configure the following command to select the types of route that should be advertised.

eigrp stub *connected-parameter redistributed-parameter static-parameter summary-parameter*

Click the variables to make your selection. You must select at least one route type, but you can select all or any combination.

- *connected-parameter*. Select **connected** to advertise connected routes. If the connected routes are not covered by a **network** statement, it might be necessary to configure redistribution for the connected routes in the EIGRP process.
- *redistributed-parameter*. Select **redistributed** to advertise routes redistributed into the EIGRP routing process from other routing protocols.
- *static-parameter*. Select **static** to advertise static routes. You must also enable the **configure redistribution** command and configure the redistribution of static routes.
- *summary-parameter*. Select **summary** to advertise summary routes.

Step 9 (Optional.) If necessary, adjust settings for the other initially disabled commands. See Customize the EIGRP Process, on page 389.

Step 10 Click **OK**.

Customize the EIGRP Process

EIGRP includes many options that have default values. These values work well for many networks. However, you might need to adjust one or more settings to get the precise behavior you need. The following topics explain the various ways you can customize your EIGRP routing process.

Configure EIGRP Advanced Settings

You can configure several settings that control the overall behavior of an EIGRP process, including automatic route summarization, distance metrics, logging. and the router ID used for sending link state advertisements and other routing updates. Many of these settings have defaults that are appropriate for most networks.

Before you begin

This procedure assumes you have already configured the EIGRP process; see Configure the Core EIGRP Process, on page 386.

When you create the process, certain advanced options are enabled by default. You will see these enabled options when you edit the EIGRP object.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the global virtual router.

Step 3 Click the **EIGRP** tab.

Step 4 Click the edit icon ( ) for the EIGRP object.

Note that when you edit an object, you might see lines that you did not directly configure. These lines are exposed to show you the default values that are being configured.

Step 5 Click **Show Disabled** above the object body to add all other possible configuration lines.

Step 6 The **setup eigrp** *configuration* line should already be enabled as **setup eigrp advanced**. If it is not, click the + to the left of the line to enable it, then click the variable and select **advanced**.

Step 7 (Optional; not recommended.) To automatically summarize routes on network number boundaries, click + next to the **auto-summary** command.

This can cause routing problems if you have noncontiguous networks.

For example, if you have a router with the networks 172.16.1.0, 172.16.2.0, and 172.16.3.0 connected to it, and those networks all participate in EIGRP, the EIGRP routing process creates the summary address 172.16.0.0 for those routes. If an additional router is added to the network with the networks 172.16.10.0 and 172.16.11.0, and those networks participate in EIGRP, they will also be summarized as 172.16.0.0. Thus, automatically summarizing the routes leads to traffic being routed to the wrong router.

Step 8 (Optional; recommended.) Configure the router ID.

Click + to enable the **router-id** command, then click the variable and enter the IPv4 address that should be used when sending any router updates from this device. No two routers in an EIGRP autonomous system can have the same router ID, so ensure that it is unique in the system.

If you do not explicitly specify a router ID for the process, the system uses the highest IP address assigned to an active interface. Thus, the router ID can change if you disable the selected interface, or you change its addresses. By assigning a router ID explicitly, you ensure consistency for your process.

Step 9 (Optional.) Configure the administrative distances for internal and external EIGRP routes.

The following command is enabled by default when you configure the process. If you are configuring a new object, you might need to click + to enable the command.

distance eigrp 90 170

Because every routing protocol has metrics based on algorithms that are different from the other routing 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. Administrative distance is a route parameter that the system uses to select the best path when there are two or more different routes to the same destination from two different routing protocols.

Administrative distances for EIGRP are from 1 to 255. These numbers are relative to the administrative values assigned to other routing processes when the system chooses the best routes. In general, the higher the value, the lower the trust rating. The default values should work for most networks. Adjust them it you want to either give preference to EIGRP routes, or if you want to reduce the likelihood of EIGRP routes being used.

The numbers mean the following:

- First value (90): **Internal Distance**. The administrative distance for EIGRP internal routes. Internal routes are those that are learned from another entity within the same autonomous system.
- Second value (170): **External Distance**. The administrative distance for EIGRP external routes. External routes are those for which the best path is learned from a neighbor external to the autonomous system

Step 10 The **default-metric** command is used when redistributing routes from other routing processes. Configure it only if you also configure redistribution. For details, see Configure EIGRP Route Redistribution, on page 397.

Step 11 Configure neighbor logging.

The following commands are enabled by default when you configure the process. If you are configuring a new object, you might need to click + to enable the command. If you want to disable logging, click - to disable the commands.

- **eigrp log-neighbor-changes** enables the logging of EIGRP neighbor adjacency changes.
- **eigrp log-neighbor-warnings 10** enables the logging of EIGRP neighbor warning messages. The number is the time interval between repeated neighbor warning messages, from 1 to 65535 seconds. Repeated warnings are not logged if they occur during this interval.

Step 12 If you want to configure the **setup stub** command, see Configure the EIGRP Process for Stub Routing, on page 387.

Step 13 Click **OK**.

Configure the Networks for EIGRP to Advertise

Use the **network** command to identify the networks, and by implication, interfaces that should be included in EIGRP routing. For an interface to participate in EIGRP routing, it must fall within the range of addresses defined by the network entries. For directly connected and static networks to be advertised, they must also fall within the range of the network entries.

Before you begin

This procedure assumes you have already configured the EIGRP process; see Configure the Core EIGRP Process, on page 386.

Create the network objects that define the networks to advertise.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon () for the global virtual router.

Step 3 Click the **EIGRP** tab.

Step 4 Click the edit icon () for the EIGRP object.

Note that when you edit an object, you might see lines that you did not directly configure. These lines are exposed to show you the default values that are being configured.

Step 5 Click **Show Disabled** above the object body to add all other possible configuration lines.

Step 6 Assuming you already configured a network, click **...** > **Duplicate** next to a **network** line to create a new empty command.

If you have defined no networks yet, click + next to the empty **network** *network-object* line.

Step 7 In the **network** command, click the variable and select the object that defines a network that should be included in this autonomous system.

Typically, this would be a directly-connected network. For example, if the IP address of the inside interface is 192.168.1.1/24, the associated network object for this command would contain 192.168.1.0/24. If the object does not already exist, click **Create New Network** and create it now.

Directly connected and static networks that fall within the defined network are advertised by the process. Additionally, only interfaces with an IP address that fall within the defined network participate in the EIGRP routing process.

If you have an interface that you do not want to have participate in EIGRP routing, but that is attached to a network that you want advertised, see Configure EIGRP Passive Routing Interfaces, on page 392.

Step 8 If you have additional networks to route, click **...** > **Duplicate** to the left of the **network** command to add a new one. Keep adding **network** lines until you have configured all the networks to route.

Step 9 Click **OK**.

Configure EIGRP Passive Routing Interfaces

If you have an interface that you do not want to have participate in EIGRP routing, but that is attached to a network that you want advertised, you can configure a **network** command that includes the network to which the interface is attached, and use the **passive-interface** command to prevent that interface from sending or receiving EIGRP updates.

By default, the system enables the **no passive-interface default** command, which sets all interfaces to be active, sending and receiving EIGRP updates.

The following procedure explains how to change an interface to be passive.

Before you begin

This procedure assumes you have already configured the EIGRP process; see Configure the Core EIGRP Process, on page 386.

When you create the process, you add the **network** commands to indicate which networks should be routed using EIGRP. To configure additional networks to route, see Configure the Networks for EIGRP to Advertise, on page 391.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the global virtual router.

Step 3 Click the **EIGRP** tab.

Step 4 Click the edit icon ( ) for the EIGRP object.

Note that when you edit an object, you might see lines that you did not directly configure. These lines are exposed to show you the default values that are being configured.

Step 5 Click **Show Disabled** above the object body to add all other possible configuration lines.

Step 6 If you are editing an object, the **configure interface passive** command and its child, **no passive-interface default**, are enabled.

For a new object, click **+** to enable the **configure routing-interface** *parameters* command.

Step 7 To configure interfaces to be active by default, and then selectively make interfaces passive:

a) In the **configure routing-interface** command, click the variable and select **passive**.

This action enables the **no passive-interface default** command, which makes EIGRP interfaces active by default.

b) Click **+** next to the **passive-interface** *interface* command, click the variable, and select the interface that should be passive and not participate in EIGRP routing updates.

c) Click **...** > **Duplicate** next to the **passive-interface** *interface* command if you need to configure additional passive interfaces. Continue until you have a **passive-interface** command for each interface that should be passive.

Step 8 To configure interfaces to be passive by default, and then selectively make interfaces active:

a) In the **configure routing-interface** command, click the variable and select **active**.

This action enables the **passive-interface default** command, which makes EIGRP interfaces passive by default.

b) Click **+** next to the **no passive-interface** *interface* command, click the variable, and select the interface that should actively participate in EIGRP routing updates.

c) Click **...** > **Duplicate** next to the **no passive-interface** *interface* command if you need to configure additional active interfaces. Continue until you have a **no passive-interface** command for each interface that should be active.

Step 9 To switch an interface back to the default behavior (either passive or active), click **-** next to the command that makes that specific interface passive or active. This deletes the exception and makes the interface behave according to the default action you set.

Step 10 Click **OK**.

Configure Static EIGRP Neighbors

EIGRP hello packets are sent as multicast packets. If an EIGRP neighbor is located across a non broadcast network, such as a VPN tunnel, you must manually define that neighbor. When you manually define an EIGRP neighbor, hello packets are sent to that neighbor as unicast messages.

You do not need to define static neighbors that are on regular broadcast networks, as these routers can form adjacencies themselves.

Before you begin

This procedure assumes you have already configured the EIGRP process; see Configure the Core EIGRP Process, on page 386.

Determine the interface through which the system should reach the neighbor.

You can also configure logging settings for neighbors, as explained in Configure EIGRP Advanced Settings, on page 389.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the global virtual router.

Step 3 Click the **EIGRP** tab.

Step 4 Click the edit icon ( ) for the EIGRP object.

Note that when you edit an object, you might see lines that you did not directly configure. These lines are exposed to show you the default values that are being configured.

Step 5 Click **Show Disabled** to expose all commands, then click **+** to enable the **neighbor** command.

Step 6 Configure the neighbor address.

neighbor *ip-address* **interface** *interface*

- Click *ip-address* and enter the IP address of the neighbor router.
- Click *interface* and select the interface through which the system can reach the router.

Step 7 If necessary, configure a static route for the neighbor router.

If the IP address of the router is on the same network as the selected interface, a static route is not necessary. For example, if you select an interface whose IP address is 10.100.10.1/24, and the neighbor address is 10.100.10.2/24, you do not need a static route.

Step 8 You can click **...** > **Duplicate** next to the **neighbor** command to define another static neighbor. Define as many as you require.

Step 9 Click **OK**.

Control EIGRP Candidate Default Route Dissemination

You can control the sending or receiving of default route candidates from the EIGRP process. By default, all candidate routes are advertised or accepted contingent on your route filtering and redistribution settings.

You cannot directly turn off the sending or receiving of a default route. If you want to prevent default route dissemination from EIGRP, configure these commands with a standard ACL that denies the any-ipv4 network.

Before you begin

This procedure assumes you have already configured the EIGRP process; see Configure the Core EIGRP Process, on page 386.

Create the Smart CLI standard access list objects you need for each filter rule. Use deny access control entries (ACEs) to filter out routes that match the entry, and permit ACEs for the routes that should be updated.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the global virtual router.

Step 3 Click the **EIGRP** tab.

Step 4 Click the edit icon ( ) for the EIGRP object.

Note that when you edit an object, you might see lines that you did not directly configure. These lines are exposed to show you the default values that are being configured.

Step 5 Click **Show Disabled** above the object body to add all other possible configuration lines.

Step 6 Click + to enable one or both of the following commands:

- **default-information in** *acl* for controlling the receiving of candidate default routes.
- **default-information out** *acl* for controlling the sending of candidate default routes.

Step 7 Click the variable and select the standard ACL that applies your filter.

Step 8 Click **OK**.

Configure EIGRP Filter Rules

You can filter incoming or outgoing routing updates based on network prefixes as defined in a standard access control list. Filtering improves your control of route distribution into an EIGRP autonomous system or outbound to other routing processes.

Before you begin

This procedure assumes you have already configured the EIGRP process; see Configure the Core EIGRP Process, on page 386.

Create the Smart CLI standard access list objects you need for each filter rule. Use deny access control entries (ACEs) to filter out routes that match the entry, and permit ACEs for the routes that should be updated.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the global virtual router.

Step 3 Click the **EIGRP** tab.

Step 4 Click the edit icon ( ) for the EIGRP object.

Note that when you edit an object, you might see lines that you did not directly configure. These lines are exposed to show you the default values that are being configured.

Step 5 Click **Show Disabled** to expose all commands, then click + to enable the **configure filter-rules** *direction* command.

Step 6 Click *direction* and select **in**, for filtering incoming updates, or **out**, for filtering outbound updates.

This action adds the **distribute-list** commands to the object.

Step 7 For inbound filters, you can optionally specify the interface on which to filter updates. If you do not specify an interface, the filter applies to all updates received on any interface. Click + to enable one of the following options:

- **distribute-list** *acl-name* **in**

  Select your standard ACL object.

- **distribute-list** *acl-name* **in interface** *interface*

  Select your standard ACL object and the interface on which to filter incoming updates.

Step 8 For outbound filters, you can optionally specify the protocol, to limit the filter to routes generated by that routing process, and the interface on which to filter updates. Click + to enable one of the following options:

- **distribute-list** *acl-name* **out**

  Select your standard ACL object.

- **distribute-list** *acl-name* **out interface** *interface*

  Select your standard ACL object and the interface on which to filter outgoing updates.

- **distribute-list** *acl-name* **out** *protocol*

  Select your standard ACL object and one of the following route types:

  - **connected**. For routes established for networks that are directly connected to the system's interfaces.
  - **static**. For static routes you manually created.
  - **rip**. For routes generated by RIP.

- **distribute-list** *acl-name* **out** *protocol identifier*

  Select your standard ACL object and one of the following route types:

  - **ospf** *process-id*. For routes generated by OSPF. Click identifier and enter the process ID for the OSPF process defined on the system.
  - **bgp** *automonous-system*. For routes generated by BGP. Click identifier and enter the autonomous system number for the BGP process defined on the system.

Step 9 You can click **...** > **Duplicate** next to the **configure filter-rules** command to define another filter rule. Define as many as you require.

Step 10 Click **OK**.

Configure EIGRP Route Redistribution

You can control the redistribution of routes into an EIGRP process from other routing protocols, connected routes, and static routes.

Before you begin

It is best practice to configure the routing process from which you will redistribute routes, and deploy your changes, before you configure redistribution into EIGRP.

If you want to apply a route map to fine-tune which routes are redistributed, create the Smart CLI route map object. Routes that match the route map are redistributed, and all non-matching routes are not redistributed.

This procedure assumes you have already configured the EIGRP process; see Configure the Core EIGRP Process, on page 386.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the global virtual router.

Step 3 Click the **EIGRP** tab.

Step 4 Click the edit icon ( ) for the EIGRP object.

Note that when you edit an object, you might see lines that you did not directly configure. These lines are exposed to show you the default values that are being configured.

Step 5 Click **Show Disabled** to expose all commands.

Step 6 (Optional.) Click + to enable the **default-metric** command, which is in the **setup eigrp advanced** group of commands.

The **default-metric** command sets the metrics to use for redistributed routes if you do not configure a specific **redistribute metric** command for a route type.

default-metric *bandwidth-metric delay-metric reliability-metric effective-bandwidth path-MTU*

Click the variables and configure the following. You must configure all metric variables.

- *bandwidth-metric*. Click the variable and enter the bandwidth for connections on this route from 1 to 4294967295 Kilobytes per second.
- *delay-metric*. Click the variable and enter the delay for connections on the route, in 10 microsecond units, from 0 to 4294967295.
- *reliability-metric*. Click the variable and enter the EIGRP reliability metric for the route, from 0 to 255, where 255 indicates 100% reliability. This metric is ignored but you still must configure it.
- *effective-bandwidth*. Click the variable and enter the EIGRP effective bandwidth for the route, from 1 to 255, where 255 indicates 100% loaded. This metric is ignored but you still must configure it.

- *path-MTU*. Click the variable and enter the mean transmission unit (MTU) of the path, from 1 to 65535. This metric is ignored but you still must configure it.

Step 7 Click **+** to enable the **configure redistribution** command.

Step 8 Click the *protocol* variable and select the source process from which you are redistributing routes. You can redistribute **connected** and **static** routes, or routes generated by **bgp**, **isis**, **ospf**, or **rip**.

Step 9 If you select a routing process, click the *identifier* variable and enter the required value:

- **bgp**. Enter the autonomous system number.
- **ospf**. Enter the process ID number.
- **connected**, **static**, **isis**, **rip**. Enter **none**. Even if you enter a different value, it will be ignored.

Step 10 (Optional; IS-IS only.) On the **redistribute isis route-level** *route-level* command, click the variable and select whether you are redistributing routes learned only within an IS-IS area (**level-1**), between IS-IS areas (**level-2**) or both (**level-1-2**).

Step 11 (Optional; all protocols.) To fine-tune which routes are redistributed based on a route map, click **+** to enable the **redistribute** *route-map* command, click the variable, and select the route map that defines your restrictions.

If you do not apply a route map, all routes for the process (that fit the other commands configured for redistribution), are redistributed.

Step 12 (Optional; all protocols.) To fine-tune the metrics for redistributed routes, click **+** to enable the following command and configure the options:

redistribute *protocol* **metric** *bandwidth-metric delay-metric reliability-metric effective-bandwidth path-MTU*

Click the variables and configure the values, which are explained under the **default-metric** command above. You must configure all metric variables.

Step 13 (Optional; OSPF only.) The following commands are enabled by default when you redistribute routes from an OSPF process. You can click **-** to disable unwanted commands.

These commands specify the criteria by which OSPF routes are redistributed into other routing domains.

- **redistribute ospf match external 1**. Routes that are external to the autonomous system, but are imported into OSPF as Type 1 external routes.
- **redistribute ospf match external 2**. Routes that are external to the autonomous system, but are imported into OSPF as Type 2 external routes.
- **redistribute ospf match internal**. Routes that are internal to a specific autonomous system.
- **redistribute ospf match nssa-external 1**. Routes that are external to the autonomous system, but are imported into OSPF as Type 1 external routes and marked as Not-So-Stubby-Area (NSSA) only.
- **redistribute ospf match nssa-external 2**. Routes that are external to the autonomous system, but are imported into OSPF as Type 2external routes and marked as Not-So-Stubby-Area (NSSA) only.

Step 14 You can click **...** > **Duplicate** next to the **configure redistribution** command to configure redistribution for another protocol. Configure redistribution for each protocol that makes sense for your network.

Step 15 Click **OK**.

Monitoring EIGRP

You can use the following commands to monitor the EIGRP routing process. For examples and descriptions of the command output, see the command reference.

- **show eigrp events** [{*start end*} | **type**]

  Displays the EIGRP event log.

- **show eigrp interfaces** [*if-name*] [**detail**]

  Displays the interfaces participating in EIGRP routing.

- **show eigrp neighbors** [**detail** | **static**] [*if-name*]

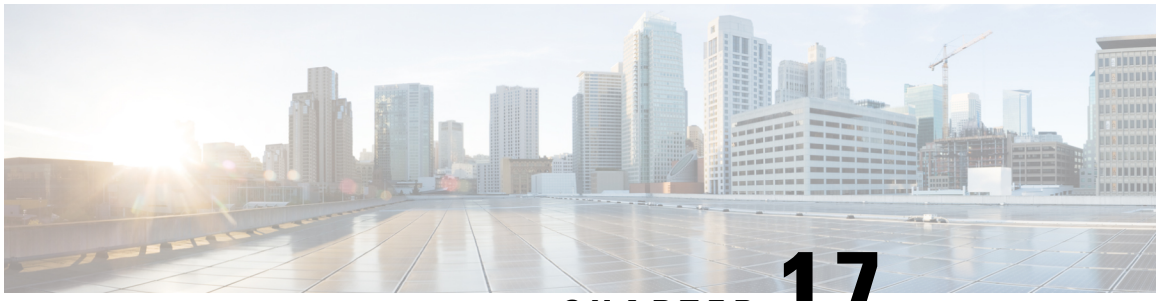
  Displays the EIGRP neighbor table.

- **show eigrp topology** [*ip-addr* [*mask*] | **active** | **all-links** | **pending** | **summary** | **zero-successors**]

  Displays the EIGRP topology table.

- **show eigrp traffic**

  Displays EIGRP traffic statistics.



CHAPTER 17

Border Gateway Protocol (BGP)

BGP is used to exchange routing information for the Internet and is the protocol used between Internet service providers (ISP). If your system is a gateway to the service provider network, you might need to implement BGP. You can configure one BGP process on the device, for a single autonomous system.

- BGP, on page 401
- Configure BGP, on page 405
- Monitoring BGP, on page 425

BGP

Border Gateway Protocol (BGP) is an inter and intra-autonomous system routing protocol that

- operates within and between autonomous systems (networks or groups of networks under common administration with common routing policies),
- exchanges routing information for the Internet, and
- serves as the protocol used between Internet service providers (ISPs).

Routing table changes

Routing table changes are network updates that

- occur when BGP neighbors exchange full routing information upon initial TCP connection establishment
- involve sending only changed routes to neighbors when routing table modifications are detected, and
- advertise only the optimal PATH to a destination network without periodic updates.

BGP attributes for route selection

RRoutes learned via BGP have properties used to determine the best route when multiple paths to a destination exist. These properties are called 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 selects an exit point from the local AS. Unlike the weight attribute, it is shared throughout the local AS. If multiple exit points exist, the one with the highest local preference attribute is chosen 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. However, when the next hop is in the same subnet AS the peering address of the EBGP peer, the next hop is not modified. This behavior is referred to AS the third party next hop.
- 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.



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 device in performing loop checks and to avoid delays in the existing outgoing update tasks.

BGP usage scenarios

BGP usage scenarios are network configurations that

- enable ISPs to exchange customer and ISP routes between autonomous systems using External BGP (EBGP),

- allow service providers to exchange routes within an autonomous system using Interior BGP (IBGP), and
- support IPv6 prefix routing over IPv6 networks.

BGP implementation contexts

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 path selection is a routing process that selects only one path AS the best path when BGP receives multiple advertisements for the same route from different sources, puts the selected path in the IP routing table, and propagates the path to its neighbors using specific criteria in order of precedence.

BGP path selection criteria

BGP uses these criteria, in the order presented, to select a path for a destination:

1. If the path specifies a next hop that is inaccessible, drop the update. Otherwise, prefer the path with the largest weight.
2. 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.
3. 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).
4. 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.
5. If the paths are still the same, prefer the path through the closest IGP neighbor. Then, determine if multiple paths require installation in the routing table for BGP multipath, on page 403.
6. Determine if multiple paths require installation in the routing table for BGP multipath, on page 403.
7. If both paths are external, prefer the path that was received first (the oldest one).
8. 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.
9. Prefer the path that comes from the lowest neighbor address.

BGP multipath

BGP multipath is a BGP feature that

- allows installation into the IP routing table of multiple equal-cost BGP paths to the same destination prefix,
- enables traffic to the destination prefix to be shared across all installed paths, and
- does not affect best-PATH selection algorithm.

BGP multipath PATH requirements

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.

Configure BGP

The following topics explain how to configure BGP.

Configure BGP Global Settings

If you configure BGP, the global settings apply across all virtual routers, if you use virtual routers. There are additional BGP settings you configure to define the BGP process. When using virtual routers, you can create a separate BGP process for each virtual router.

Before you begin

After you create the BGP Global Settings object, you can delete it if you no longer need it. Simply edit the object following this procedure, but click the **Delete BGP Global Settings Object** button at the bottom of the dialog box.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 On the main routing or virtual routers page, click the **BGP Global Settings** button.

If you are viewing a virtual router, you must return to the main list of virtual routers.

Step 3 If you have not yet configured the BGP Global Settings object, click **Create BGP Global Settings Object**.

Step 4 (Optional.) You can change the object name or enter a description for the object. The default object name is BgpGeneralSettings.

Step 5 Configure at least the following basic settings:

- **router bgp** *as-number*. Click *as-number* and enter the autonomous system (AS) number for the BGP process. 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. The system supports asplain and asdot notation as defined in RFC 5396.
- **log-neighbor-changes** *state*. Click *state* and select enable or disable. When enabled, which is recommended, BGP neighbor changes (up or down) and resets are logged. This helps in troubleshooting network connectivity problems and measuring network stability.
- **transport path-mtu-discovery** *state*. Click *state* and select enable or disable. When enabled, which is recommended, the system determines the maximum transmission unit (MTU) size on the network path between two IP hosts, and then takes advantage of the highest-MTU path. This avoids IP fragmentation.
- **fast-external-fallover** *state*. Click *state* and select enable or disable. When enabled, which is recommended, the system uses fast external fallover for BGP peering sessions with directly connected external peers. The session is immediately reset if the link goes down. If you disable BGP fast external fallover, the BGP routing process will wait until the default hold timer expires (3 keepalives) to reset the peering session.
- **enforce-first-as** *state*. Click *state* and select enable or disable. When enabled, which is recommended, the system denies incoming updates received from eBGP peers that do not list their autonomous system

number as the first segment in the AS_PATH attribute. Enabling this command prevents a misconfigured or unauthorized peer from misdirecting traffic (spoofing the local router) by advertising a route as if it was sourced from another autonomous system.

Step 6 (Optional.) Click the **Show Disabled** link above the object body to add all other possible configuration lines.

You can enable the following options by clicking the + to the left of the option.

- **bgp asnotation dot**. Changes the default display and regular expression match format of BGP 4-byte autonomous system numbers from asplain (decimal values) to asdot (dot notation). The system uses asplain as the default display format for autonomous system numbers, but you can configure 4-byte autonomous system numbers in both the asplain and asdot format even if you do not enable this command.

  In addition, the default format for matching 4-byte autonomous system numbers in regular expressions is asplain, so you must ensure that any regular expressions to match 4-byte autonomous system numbers are written in the asplain format if you do not enable this command.

- **bgp scan time 60**. Click the number and enter the scanning interval of BGP routers for next hop validation, from 5 to 60 seconds. The default is 60 seconds.

- **configure nexthop trigger** *state*. Click *state* and select either **enable** or **disable**. BGP next-hop address tracking is event driven. BGP prefixes are automatically tracked as peering sessions are established. Next-hop changes are rapidly reported to BGP as they are updated in the routing information base (RIB). This optimization improves overall BGP convergence by reducing the response time to next-hop changes for routes installed in the RIB. When a best-path calculation is run in between BGP scanner cycles, only the changes are processed and tracked. If you enable next hop address tracking, the following commands are added. Note that if you do not configure the general options in a new object, the default is to enable this feature.

  - **bgp nexthop trigger enable**. BGP next-hop address tracking improves BGP response time significantly. However, unstable Interior Gateway Protocol (IGP) peers can introduce instability to BGP. We recommend that you aggressively dampen unstable IGP peering sessions to mitigate the possible impact to BGP.

  - **bgp nexthop trigger delay 5**. Click the number to change the delay interval between routing table walks for BGP next-hop address tracking. You can increase the performance of BGP next-hop address tracking by tuning the delay interval between full routing table walks to match the tuning parameters for the IGP. The default delay interval is 5 seconds, which is an optimal value for a fast-tuned IGP. In the case of an IGP that converges more slowly, you can change the delay interval to 20 seconds or more, depending on the IGP convergence time. You can set the delay from 0 to 100 seconds.

- **bgp aggregate-timer 30**. Click the number to set the interval at which BGP routes will be aggregated, from 6 to 60 seconds. The default is 30 seconds.

- **bgp router-id** *router-id*. Click *router-id* and enter the IPv4 address that should be used as the global router ID. This ID is used for any BGP process in a virtual router that does not itself specify a router ID. If you do not enable this command, the router ID is set to the highest IP address on a physical interface that is assigned to the virtual router. Use this command to ensure that the router ID remains stable.

- **bgp maxas-limit** *value*. Click *value* and enter the maximum number of autonomous system numbers in the AS-path attribute of the BGP Update message, ranging from 1 to 254. 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. The system discards routes that have a number of autonomous systems in the AS-path that exceed the specified value. In addition to setting the limit on the number of autonomous system

numbers within the AS-path segment, the command limits the number of AS-path segments to ten. If you do not enable this command, no routes are discarded.

Step 7 (Optional.) Configure BGP advanced options.

Click the **Show Disabled** link if necessary to expose the following command. When editing the settings, both the **timers** and **bestpath** option sets are shown, as they have some defaults that are enabled even if you do not explicitly set them.

configure bgp advanced *advanced-option*

Click *advanced-option* and select one of the following. You can configure all of these options by clicking ... in the left column and selecting **Duplicate**.

- **timers**. Configures the timers used when communicating with BGP neighbor routers.

  timers bgp 60 180 0

  - First value (default 60): **Keepalive Interval**. Click the number and enter the frequency with which the system sends keepalive messages to its BGP neighbor, from 0 to 65535 seconds. We recommend that you do not specify 20 or less, or you might find routes flapping unnecessarily.
  - Second value (default 180): **Hold Time**. Click the number and enter how long the system should wait after not receiving a keepalive message before declaring a BGP neighbor dead, from 0 to 65535 seconds.
  - Third value (default 0): **Minimum Hold Time**. Click the number and specify the minimum acceptable hold-time configured on the BGP neighbor. The minimum acceptable hold-time must be less than, or equal to, the interval specified as the hold time for this system. The range is from 0 to 65535 seconds.

- **bestpath**. Configures options that are used in the BGP best path selection algorithm. The **bgp default local-preference** command is configured by default, but you can add the other commands by clicking the + for the command.
  - **bgp default local-preference 100**. Click the number and enter a value that indicates the preference of this system relative to other routers in the BGP AS, from 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. This attribute is exchanged only between iBGP peers and is used to determine local policy.
  - **bgp always-compare-med**. Allow the comparison of Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems. By default, the system does not compare the MED for paths from neighbors in different autonomous systems.
  - **bgp bestpath compare-routerid**. Use the router ID as the tie breaker for best path selection when two identical routes are received from two different peers (all the attributes are the same except for the router ID). When this command is enabled, the lowest router ID will be selected as the best path when all other attributes are equal. Otherwise, the first route received is used.
  - **bgp deterministic-med**. Select the best MED path advertised from the neighboring AS.
  - **bgp bestpath med missing-as-worst**. Set a path with a missing MED attribute as the least preferred path. By default, the system considers the route with a missing MED to be the best route.

- **graceful-restart**. Configure graceful restart for the systems in a high availability or cluster configuration.

- **bgp graceful-restart**. Enables graceful restart for non-stop forwarding. With graceful restart, the system can advertise the ability to maintain the forwarding state for an address group during restart.
- **bgp graceful-restart restart-time 120**. Click the number and enter the maximum time period that the system will wait for a graceful-restart-capable neighbor to return to normal operation after a restart event occurs, from 1 to 3600 seconds. The default is 120 seconds.
- **bgp graceful-restart stalepath-time 360**. Click the number and enter the maximum time period that the system will hold stale paths for a restarting peer, from 1 to 3600 seconds. All stale paths are deleted after this timer expires. The default value is 360 seconds.

Step 8 Click **OK**.

Configure the BGP Process

After you configure the BGP global settings, you can configure the BGP process. If you are using virtual routers, you can configure a separate process for each virtual router. You can configure at most one BGP process for the system or per virtual router.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring BGP.

Step 3 Click the **BGP** tab.

Step 4 Do one of the following:

- To create a new process, click + or click the **Create BGP Object** button.
- Click the edit icon ( ) for the object you want to edit. Note that when you edit an object, you might see lines that you did not directly configure. These lines are exposed to show you the default values that are being configured.

If you no longer need a process, click the trash can icon for the object to delete it.

Step 5 Enter a name, and optionally, a description for the object.

Step 6 Configure the minimum settings for the process:

- **router bgp** *as-number*. Click *as-number* and enter the same autonomous system (AS) number for the BGP process that you specified for the global settings. 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. The system supports asplain and asdot notation as defined in RFC 5396.
- **configure address-family** *ip-protocol*. Click *ip-protocol* and select IPv4 or IPv6. If you are using virtual routers, you can configure IPv6 for the global router only. You can configure IPv4 for any virtual router. Selecting an option adds the **address-family ipv4 unicast** or **address-family ipv6 unicast** commands, plus the following command which you must configure:

- **configure address-family {ipv4 | ipv6}** *settings*. Click settings and select either **general** or **advanced**. You must configure at least one command under these options to have a minimal process, but that will not be sufficient for a meaningful process.

Step 7 Click **Show Disabled** and customize the process to function correctly in your network.

So long as you configure a minimum set of commands, as explained above, you can save the object and customize the process settings later. The following topics explain the various sets of options. At minimum, you should configure the network settings to identify the networks for which the process will distribute routes. Both the general and advanced settings have command defaults that are appropriate in most cases.

- Configure BGP General Settings, on page 409
- Configure BGP Advanced Settings, on page 410
- Configure the Networks for BGP to Advertise, on page 412
- Configure BGP Route Injection, on page 413
- Configure BGP Aggregate Address Settings, on page 414
- Configure BGP Filter Settings for IPv4, on page 416
- Configure BGP Neighbors, on page 417
- Configure BGP Route Redistribution from Other Routing Protocols, on page 423

Step 8 (Optional.) Configure the router ID for this process.

You can configure the router ID to use for the BGP process in the BGP global settings. You can optionally configure it in the process object, too. Any router ID configured in the process object overrides the global router ID. This makes it easy to override the global value for specific virtual routers.

Click **Show Disabled** if the following command is not displayed, and click the + next to it to enable the command.

- **bgp router-id** *router-id*. Click *router-id* and enter the IPv4 address that should be used as the router ID for this process. If you do not enable this command, the router ID is set to the global router ID, or to the highest IP address on a physical interface that is assigned to the virtual router. Use this command to ensure that the router ID remains stable.

Step 9 Click **OK**.

Configure BGP General Settings

The general settings define administrative distances, timers, and for IPv4 only, next hop address tracking. These options have defaults appropriate for most networks.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring BGP.

Step 3 Click the **BGP** tab.

Step 4 Add or edit the BGP process object.

Step 5 Find the **configure address-family ipv4** or **ipv6** line. If the **general** option is already selected, move to the next step. However:

- If the *settings* variable is still displayed, click it and select **general**.
- If you have already configured advanced options, click the **...** button to the left of the command and select **Duplicate**. Then, click *settings* and select **general**.

Step 6 Configure the following commands:

- **distance bgp 20 200 200**. Configures the administrative distances for BGP, from 1 to 255. These numbers are relative to the administrative values assigned to other routing processes when the system chooses the best routes. In general, the higher the value, the lower the trust rating. Use this command if another protocol is known to be able to provide a better route to a node than was actually learned via external BGP (eBGP), or if some internal routes should be preferred by BGP. Routes with a distance of 255 are not installed in the routing table. The numbers mean the following:
  - First value (default 20): **External Distance**. Click the number and enter the administrative distance for external BGP routes. Routes are external when learned from an external autonomous system.
  - Second value (default 200): **Internal Distance**. Click the number and enter the administrative distance for internal BGP routes. Routes are internal when learned from peers in the local autonomous system. Changing the administrative distance of internal BGP routes is considered dangerous and is not recommended. Improper configuration can introduce routing table inconsistencies and break routing.
  - Third value (default 200): **Local Distance**. Click the number and enter the administrative distance for local BGP routes. Local routes are for those networks listed with a **network** command in the BGP routing process, that is, the networks that the process is advertising, or for the networks that are being redistributed to BGP from another process.

Step 7 Click **OK**.

Configure BGP Advanced Settings

Use the advanced settings to configure a variety of options that are needed only under special circumstances. Most of these options are disabled by default.

Before you begin

If you intend to configure the **table-map** command, you must first go to the **Device** > **Advanced Configuration** page and create the Smart CLI route map object needed by the command.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring BGP.

Step 3 Click the **BGP** tab.

Step 4 Add or edit the BGP process object.

Step 5 Find the **configure address-family ipv4** or **ipv6** line. If the **advanced** option is already selected, move to the next step. However:

- If the *settings* variable is still displayed, click it and select **advanced**.
- If you have already configured general options, click the **...** button to the left of the command and select **Duplicate**. Then, click *settings* and select **advanced**.

Step 6 Configure the following commands. You need to click **Show Disabled** to see all but the first command when initially creating the object.

Click + for a command to enable it.

- **bgp redistribute-internal**. Configure iBGP redistribution into an interior gateway protocol (IGP), such as EIGRP or OSPF. Exercise caution when redistributing iBGP into an IGP. Use IP prefix-list and route-map statements to limit the number of prefixes that are redistributed. Redistributing an unfiltered BGP routing table into an IGP can have a detrimental effect on normal IGP network operation. This command is enabled by default, so you need to click the - button for it to turn it off.
- **bgp suppress-inactive**. Prevent routes that are not installed in the RIB (inactive routes) from being advertised to peers. By default, BGP advertises inactive routes. Note that BGP marks routes that are not installed into the RIB with a RIB-failure flag. This flag will also appear in the output of the **show bgp** command; for example, Rib-Failure (17). This flag does not indicate an error or problem with the route or the RIB.
- **auto-summary**. (IPv4 only.) Automatically summarize subnet routes into network-level routes. Route summarization reduces the amount of routing information in the routing tables. Disable automatic summarization if you must perform routing between disconnected subnets. When automatic summarization is disabled, subnets are advertised.
- **synchronization**. Enable synchronization between BGP and your Interior Gateway Protocol (IGP) system, such as OSPF. 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. Use this command if other routers in the autonomous system do not speak BGP.
- **table-map** *route-map options*. (IPv4 only.) Apply a route map that sets metrics, a tag value, or a traffic index for routes that are updated in the BGP routing table, or controls whether routes are downloaded to the RIB. Click *route-map* and select the Smart CLI object that defines the route map. In the route map, you can use match clauses for IP access list, autonomous system paths, communities, prefix lists, and next hop.

  You can determine how the route map is used by clicking *options* and choosing either blank or **filter**:

  - If you do not select **filter**, the system uses the route map to set certain properties of a route before the route is installed into the RIB. The route is always downloaded, regardless of whether it is permitted or denied by the route map.
  - If you select **filter**, the route map also controls whether the BGP route is downloaded to the RIB. Only routes permitted in the route map are downloaded; denied routes are not downloaded.

- **default-information originate**. Configure BGP to advertise a default route (network 0.0.0.0). The configuration of the **default-information originate** command is similar to the configuration of the **network** command. The **default-information originate** command, however, requires explicit redistribution of the route 0.0.0.0, which you must also configure in this object. The **network** command requires only that the route 0.0.0.0 be present in the Interior Gateway Protocol (IGP), such as OSPF, routing table. For this reason, the **network** command is preferred for distributing a default route.
- **maximum paths 1**. Control the maximum number of parallel BGP routes that can be installed in a routing table, from 1 to 8. Use this command to configure equal-cost or unequal-cost multipath load sharing for BGP peering sessions. In order for a route to be installed as a multipath in the BGP routing table, the route cannot have a next hop that is the same as another route that is already installed. The BGP routing process will still advertise a best path to BGP peers when BGP multipath load sharing is configured. For equal-cost routes, the path from the neighbor with the lowest router ID is advertised as the best path.

  To configure BGP equal-cost multipath load sharing, all path attributes must be the same. The path attributes include weight, local preference, autonomous system path (the entire attribute and not just the length), origin code, Multi Exit Discriminator (MED), and Interior Gateway Protocol (IGP) distance.
- **maximum paths ibgp 1**. Control the maximum number of internal BGP routes that can be installed to the routing table, from 1 to 8. The considerations for multipath iBGP are the same as described under the **maximum paths** command above.

Step 7 Click **OK**.

Configure the Networks for BGP to Advertise

You need to define the networks to be advertised by the BGP routing process.

Before you begin

Create the network objects that define the networks to advertise. You can define IPv4 or IPv6 networks, or both, depending on the address families you configure for BGP.

If the network objects specify large network spaces, you can also create route maps to apply against the network object to filter out subnets within the larger space that you do not want to advertise. Only routes that match the route map specifications are advertised. Use Smart CLI to create the route map object.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring BGP.

Step 3 Click the **BGP** tab.

Step 4 Add or edit the BGP process object.

The network commands are within the command sets beneath the **configure address family ipv4** or **ipv6** command. You must configure the address family to configure the networks to advertise.

The **network** command within each address group must specify addresses that match the address family you are configuring.

Step 5 Click **Show Disabled** to expose all commands, then click **+** to enable the **network** or **network route-map** command, and configure the options:

- *network-object*. Click the variable and select the network object that defines the networks to advertise: IPv4 network address and mask or IPv6 network address and prefix.
- **route-map** *map-tag*. Click the variable and select the route map that should be applied to the network object to filter which addresses within the range should be advertised.
- (Optional; IPv6 only.) *prefix-name*. Click the variable and enter the name of a DHCPv6 prefix to advertise the prefix. If you configure this option, the network object acts as a subnet for the prefix. To use this option, you must enable the DHCPv6 Prefix Delegation client, which requires that you use FlexConfig to add the **ipv6 dhcp client pd** command to an interface in interface configuration mode.

Step 6 You can click **...** > **Duplicate** next to the **network** or **network route-map** command to configure additional networks to advertise.

Step 7 Click **OK**.

Configure BGP Route Injection

You can configure conditional route injection to inject more specific routes into a BGP routing table. Conditional route injection allows you to originate a more specific prefix into a BGP routing table without a corresponding match. A valid parent route must exist for any injected prefixes. Only prefixes that are equal to or more specific than the aggregate route (existing prefix) can be injected.

Before you begin

You must create the route maps needed to define the prefixes. These route maps must meet the requirements explained in the procedure.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring BGP.

Step 3 Click the **BGP** tab.

Step 4 Add or edit the BGP process object.

The route injection commands are within the command sets beneath the **configure address family ipv4** or **ipv6** command. You must configure the address family to configure the networks to advertise.

Step 5 Click **Show Disabled** to expose all commands, then click **+** to enable the **bgp inject-map** command.

Step 6 Configure the command properties:

- **inject-map** *inject-map*. Click the variable and select the route map that defines the prefixes that will be created and installed into the routing table. Injected prefixes are installed in the local BGP RIB. A valid parent route must exist; Only prefixes that are equal to or more specific than the aggregate route (existing prefix) can be injected. The route map must use a prefix list to specify the routes to be injected.

- **exist-map** *exist-map*. Click the variable and select the route map that defines the prefix that the BGP speaker will track. This route map must use prefix lists to specify the aggregate prefix and also the route source. The route source would be a router, for example, 10.2.1.1/32, rather than a subnet.
- *options*. Optionally, click the variable and select **copy-attributes**. This option configures the injected prefix to inherit the same attributes as the aggregate route. If you do not select this keyword, the injected prefix will use the default attributes for locally originated routes.

Step 7 You can click **...** > **Duplicate** next to the **bgp inject-map** command to configure additional route injection rules.

Step 8 Click **OK**.

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.

If you configure an aggregate route with no keywords on the command, the system will create an aggregate entry in the BGP routing table if any more-specific BGP routes are available that fall within the specified range. (A longer prefix that matches the aggregate must exist in the Routing Information Base (RIB).) The aggregate route will be advertised as coming from your autonomous system and will have the atomic aggregate attribute set to show that information might be missing. The atomic aggregate attribute is set unless you specify the **as-set** keyword.

The following procedure explains how to configure the aggregation of specific routes into one route.

Before you begin

If you want to apply route maps to fine-tune which routes are aggregated or the attributes set on the aggregate route, create the Smart CLI route map objects.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring BGP.

Step 3 Click the **BGP** tab.

Step 4 Add or edit the BGP process object.

The aggregation commands are within the command sets beneath the **configure address family ipv4** or **ipv6** command. You must configure the address family to configure aggregation.

Step 5 Click **Show Disabled** to expose all commands, then click **+** to enable the **configure aggregate-address** command.

Step 6 Click the *map-type* variable and select which types of route map you want to apply to this particular aggregate route.

This option simply determines which parameters are included on the **aggregate-address** command that will be added to the object. You can apply up to 3 separate route maps, to suppress routes from the aggregation, to advertise routes, and to define attributes to apply to the aggregate route.

- Select **no-map** if you do not need to apply any route maps.
- Select **all** if you want to apply route maps for all three options.
- Select the appropriate keyword combinations if you want to apply one or two maps, but not all: **suppress-map**, **advertise-map**, **attribute-map**, **suppress-advertise**, **suppress-attribute**, **advertise-attribute**.

Step 7 Configure the properties of the route to be aggregated.

Following is a complete list of the properties. What you see depends on your map type selection.

- *network-object*. Click the variable and select the network object that defines the address space that you want to aggregate. The object must use IPv4 or IPv6 addressing that matches the address type you are configuring. For example, you could aggregate routes for all 10.0.0.0/8 subnets.
- **suppress-map** *suppress-route-map*. Click the variable and select the route map to suppresses the advertisement of specified routes. You can use the match clauses of the route map to selectively suppress some more-specific routes of the aggregate and leave others unsuppressed. The route map can match routes based on access lists and autonomous system paths.
- **advertise-map** *advertise-route-map*. Click the variable and select the route map that selects specific routes that will be used to build different components of the aggregate route, such as AS_SET or community. This is useful when the components of an aggregate are in separate autonomous systems and you want to create an aggregate with AS_SET, and advertise it back to some of the same autonomous systems. You must remember to omit the specific autonomous system numbers from the AS_SET to prevent the aggregate from being dropped by the BGP loop detection mechanism at the receiving router. The route map can match routes based on access lists and autonomous system paths.
- **attribute-map** *attribute-route-map*. Click the variable and select the route map that changes attributes of the aggregate route. This is useful when one of the routes forming the AS_SET is configured with an attribute such as the community no-export attribute, which would prevent the aggregate route from being exported.
- *options*. Click the variable and select one, all, or none of the following options:
  - **as-set**. Generate autonomous system set path information for the aggregate route. The path advertised for this route will be an AS_SET consisting of all elements contained in all paths that are being summarized. Do not use this keyword when aggregating many paths, because this route must be continually withdrawn and updated as autonomous system path reachability information for the summarized routes changes.
  - **summary-only**. Suppress advertisements of more-specific routes to all neighbors.

Step 8 You can click **...** > **Duplicate** next to the **configure aggregate-address** command to configure additional routes to aggregate.

Step 9 Click **OK**.

Configure BGP Filter Settings for IPv4

You can create filter rules to restrict the routing information the system learns from, or advertises to, other routing protocols.

The configuration explained here applies to all local processes and for filtering updates to all BGP neighbors. You can configure different filtering rules per neighbor in the neighbor settings.

Before you begin

Create the Smart CLI standard access list objects you need for each filter rule. Use deny access control entries (ACEs) to filter out routes that match the entry, and permit ACEs for the routes that should be updated.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring BGP.

Step 3 Click the **BGP** tab.

Step 4 Add or edit the BGP process object.

The filtering commands are within the command sets beneath the **configure address family ipv4** command. You must configure the address family to configure filtering. These rules are not available for IPv6.

Step 5 Click **Show Disabled** to expose all commands, then click + to enable the **configure filter-rules** *direction* command.

Step 6 Click *direction* and select **in**, for filtering incoming updates, or **out**, for filtering outbound updates.

Step 7 For inbound filters, you can optionally specify the interface on which to filter updates. If you do not specify an interface, the filter applies to all updates received on any interface.

a) Click + to enable the **distribute-list** *acl-name* **in interface** *interface* command.

b) Click the *interface* variable and select the interface.

Step 8 For outbound filters, you can optionally specify the protocol, to limit the filter to routes advertised to that routing process.

There are two forms of the **distribute-list out** command, one with an *identifier* variable after the *protocol* variable, and one without the identifier. You can select the following protocols, but they are divided between these command versions based on whether you must provide the additional identifier information.

- **connected**. For routes established for networks that are directly connected to the system's interfaces.
- **static**. For static routes you manually created.
- **rip**. For routes advertised to RIP.
- **bgp** *automonous-system*. For routes advertised to BGP. Click *identifier* and enter the autonomous system number for the BGP process defined on the system.
- **eigrp** *automonous-system*. For routes advertised to EIGRP. Click *identifier* and enter the autonomous system number for the EIGRP process defined on the system.
- **ospf** *process-id*. For routes advertised to OSPF. Click *identifier* and enter the process ID for the OSPF process defined on the system.

Step 9 You can click **...** > **Duplicate** next to the **configure filter-rules** command to define another filter rule. Define as many as you require.

Step 10 Click **OK**.

Configure BGP Neighbors

You need to define the neighbors with whom BGP will exchange routing updates.

Before you begin

Several optional commands require Smart CLI objects, for route maps, prefix lists, and so forth. Examine the options you need to configure to determine if you need objects. You must create the Smart CLI objects before configuring the associated BGP command.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring BGP.

Step 3 Click the **BGP** tab.

Step 4 Add or edit the BGP process object.

The neighbor commands are within the command sets beneath the **configure address family ipv4** or **ipv6** command. You must configure the neighbors separately for each address family.

Step 5 Click **Show Disabled** to expose all commands, then click + to enable the **configure neighbor** command.

Step 6 Configure the basic neighbor parameters on the neighbor command:

- **neighbor** *neighbor-address*. Click the variable and enter the IPv4 or IPv6 address of the BGP neighbor router, as appropriate for the address group you are configuring.
- **remote-as** *as-number*. Click the variable and enter the autonomous system number of the BGP neighbor router.
- *config-options*. Click the variable and select **properties**. The only property that is configured by default activates the neighbor. You can adjust other options as explained in this procedure.

Step 7 (Optional.) Configure the neighbor general settings.

a) Click + to enable the **configure neighbor** *neighbor-address* **remote-as** *settings* command. Click **Show Disabled** if you cannot see the command.
b) Click *settings* and select **general**.
c) On the **configure neighbor description** command, either click the variable and enter a description of the neighbor (such as its location or purpose), up to 80 characters, or click - to disable the command if you do not want a description. A description cannot include spaces or question marks.
d) (IPv4 only.) The **configure neighbor shutdown** command is initially enabled. This command disables communication with this BGP neighbor, terminating any active session and removing all associated routing information. If you want to actively communicate with this neighbor, click - to disable this command.
e) On the **configure neighbor fall-over bfd** command, either click *option* and select **single-hop** or **multi-hop** (based on your BFD configuration), or click - to disable the command.

This command registers BGP to receive forwarding path detection failure messages from Bidirectional Forwarding Detection (BFD). Whether you select single-hop or multi-hop depends on the type of BFD template you created and attached to the interface that faces this neighbor. Ensure that your selection here is consistent with the BFD template. You must use FlexConfig to build and apply BFD templates.

Step 8 (Optional.) Configure the neighbor advanced settings.

a) If you have already configured it, click **...** > **Duplicate** for the **configure neighbor** *neighbor-address* **remote-as** *settings* command, or simply click **+** to enable it if it is not yet in use. Click **Show Disabled** if you cannot see the command.

b) Click *settings* and select **advanced**.

c) On the **neighbor password** command, either click the *secret* variable and select the secret key object that contains the password to use when authenticating the neighbor, or click **-** to disable the command if you do not want use message digest 5 (MD5) authentication. You can create the key object while editing the BGP object.

The secret key object must contain a key case-sensitive password of up to 25 characters in length. The string can contain any alphanumeric characters, including spaces, and the special characters ` ~ ! @ # $ % ^ & * () - _ = + | \ } ] { [ " ` : ; / > <., ?. However, you cannot specify a password in the format number-space-anything. The space after the number can cause authentication to fail.

Ensure the neighbor is configured to use the same password.

d) On the **configure neighbor hops** command, click the *options* variable and select one of the following, or click **-** to disable the command if the peer is not multiple hops away (that is, not directly connected to this system). Use these options with care as you can end up with routing loops and oscillating routes: configuring only directly-connected peers is preferred.

- **ebgp-multihop**. Accept and attempt BGP connections to external peers residing on networks that are not directly connected. If you select this option, the following commands are added:
  - **neighbor ebgp-multihop 255**. Click the 255 and enter the time-to-live value in number of hops, from 1-255.
  - **neighbor disable-connected-check**. Click **+** to enable this command to disable connection verification to establish an eBGP peering session with a single-hop peer that uses a loopback interface. Without this command, if the peer is not directly connected to the same network segment, connection verification will prevent the peering session from being established.
- **ttl-security-hop**. Secure a BGP peering session and configure the maximum number of hops that separate two external BGP (eBGP) peers. If you select this option, the following command is added:

  neighbor ttl-security hops *hop-count*. Click the variable and enter the maximum number of hops that separate the peers, from 1-254.

  The **neighbor ttl-security** command provides a lightweight security mechanism to protect BGP peering sessions from CPU utilization-based attacks. These types of attacks are typically brute force Denial of Service (DoS) attacks that attempt to disable the network by flooding the network with IP packets that contain forged source and destination IP addresses in the packet headers.

  This feature leverages designed behavior of IP packets by accepting only IP packets with a TTL count that is equal to or greater than the locally configured value. Accurately forging the TTL count in an IP packet is generally considered to be impossible. Accurately forging a packet to match the TTL count from a trusted peer is not possible without internal access to the source or destination network.

To maximize the effectiveness of this feature, the hop-count value should be strictly configured to match the number of hops between the local and external network. However, you should also take path variation into account when configuring this feature for a multihop peering session. Ensure that you configure this feature on all routers in the network.

e) On the **neighbor version** command, click the *version-number* variable and enter 4 to force the software to use BGP version 4, or click **-** to disable the command. The software uses version 4 by default and negotiates down to version 2 dynamically if necessary. Configuring 4 on this command prevents version negotiation.

f) On the **neighbor transport connection-mode** command, click the *options* variable and select whether the TCP connection is **active** or **passive**, or click **-** to disable the command and leave the mode to default.

g) On the **neighbor transport path-mtu-discovery** command, click the *options* variable and select **blank** to enable path MTU discovery, or **disable** to disable it. Selecting blank is the same as clicking **-** to disable the command, as the system performs path MTU discovery by default. Path MTU discovery makes it possible for the BGP session to take advantage of larger MTU links.

Step 9 (Optional.) Configure the neighbor migration settings.

The migration settings configure the **neighbor local-as** command. The **neighbor local-as** command is used to customize the AS_PATH attribute by adding and removing autonomous system numbers for routes received from eBGP neighbors. The configuration of this command allows a router to appear to external peers as a member of another autonomous system for the purpose of autonomous system number migration. This feature simplifies the process of changing the autonomous system number in a BGP network by allowing the network operator to migrate customers to new configurations during normal service windows without disrupting existing peering arrangements.

You can perform this migration for only true eBGP peering sessions. This command does not work for two peers in different sub-autonomous systems of a confederation.

Caution
BGP prepends the autonomous system number from each BGP network that a route traverses to maintain network reachability information and to prevent routing loops. Configure this command only for autonomous system migration, and de-configure it after the transition has been completed. This procedure should be attempted only by an experienced network operator. You can create routing loops through improper configuration.

a) If you have already configured it, click **...** > **Duplicate** for the **configure neighbor** *neighbor-address* **remote-as** *settings* command, or simply click **+** to enable it if it is not yet in use. Click **Show Disabled** if you cannot see the command.

b) Click *settings* and select **migration**. This adds the following command:

configure *neighbor-address* **local-as** *local-as-number options*

c) Click the *local-as-number* variable and enter the local autonomous system (AS) number to prepend to the AS_PATH attribute, from 1 to 4294967295 (asplain notation) or from 1.0 to 65535.65535 (asdot notation). You cannot specify the autonomous system number from the local BGP routing process or from the network of the remote peer.

d) Click the *options* variable and select one of the following. Note that selecting an item in this list (other than **none**) also selects all options above it in the list. This is expected: the options are not truly independent.

- **none**. Do not configure any of the following options.
- **no-prepend**. Do not prepend the local autonomous system number to any routes received from the eBGP neighbor.

- **replace-as**. 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.
- **dual-as**. 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 10 (Optional. IPv4 only.) Configure the neighbor high availability (HA) settings.

The HA mode settings configure the **neighbor ha-mode graceful-restart** command, which enables or disables the graceful restart capability for an individual BGP neighbor. Use the disable keyword to disable the graceful restart capability when graceful restart has been previously enabled for the BGP peer.

The graceful restart capability is negotiated between nonstop forwarding (NSF)-capable and NSF-aware peers in OPEN messages during session establishment. If you enable the graceful restart capability after a BGP session has been established, you will need to restart the session with a soft or hard reset.

The HA mode setting configures graceful restart for an individual neighbor. Instead, you can use the BGP global settings to enable graceful restart for all neighbors.

a) If you have already configured it, click **...** > **Duplicate** for the **configure neighbor** *neighbor-address* **remote-as** *settings* command, or simply click **+** to enable it if it is not yet in use. Click **Show Disabled** if you cannot see the command.
b) Click *settings* and select **ha-mode**.
c) If you want to disable graceful restart, click *options* on the **neighbor ha-mode graceful-restart** command and select **disable**. Select blank to reverse a previous disable action.

Step 11 (Optional.) Configure neighbor activation options.

When you configure a new neighbor, it is activated by default. You need to enable the activation settings if you want the neighbor to be disabled initially, or to configure other activation settings.

a) Click **+** to enable the **configure neighbor** *neighbor-address* **activate** *activate-options* command. Click **Show Disabled** if you cannot see the command.
b) Click *activate-options* and select **properties**.
c) The **neighbor** *neighbor-address* **activate** command is added in the enabled state. Click **-** to disable the command and configure the neighbor as initially disabled. You will need to edit this object to enable the neighbor when you are ready to communicate with it.

Step 12 (Optional.) Configure filtering in the neighbor activation settings.

a) If you have already configured it, click **...** > **Duplicate** for the **configure neighbor** *neighbor-address* **activate** *settings* command, or simply click **+** to enable it if it is not yet in use. Click **Show Disabled** if you cannot see the command.
b) Click *settings* and select **filtering**.
c) Configure filtering to control the prefixes received from or sent to this neighbor using any combination of the following neighbor commands. Click **-** to disable any that you do not want to use. All of these commands allow filtering in both the inbound and outbound direction: click **...** > **Duplicate** for a command if you want to configure both directions.

   Do not apply both a **neighbor distribute-list** and a **neighbor prefix-list** command to a neighbor in the same direction. These two commands are mutually exclusive, and only one of them can be applied to each inbound or outbound direction.

- **distribute-list** *acl options*. (IPv4 only.) Filter prefixes based on the selected standard access list (ACL). Then, click *options* and select whether to apply the filter in the **in** or **out** direction.
- **route-map** *route-map options*. Filter prefixes based on the selected route map. Then, click *options* and select whether to apply the filter in the **in** or **out** direction. Within the route map, you can configure filtering based on access list, AS path, prefix, and distribution lists.
- **prefix-list** *prefix-list options*. Filter prefixes based on the selected IPv4 or IPv6 prefix list. Then, click *options* and select whether to apply the filter in the **in** or **out** direction.
- **filter-list** *as-path options*. Filter prefixes based on the selected AS Path filter object. Then, click *options* and select whether to apply the filter in the **in** or **out** direction.

d) On the **configure prefix-limit neighbor** *neighbor-address limit-options* command, either click *limit-options* and select one of the following, or click - to disable the command. Selecting any option adds some form of the **neighbor maximum-prefix** command, with additional options you need to configure. Use this command to control how many prefixes can be received from the neighbor.

- **none**. Configure the basic form of the command without additional parameters. Click the variables and configure the following values:
  - max-prefix-limit. The maximum number of prefixes allowed from this neighbor, from 1 - 2147483647. You must also configure this variable if you select any of the other options.
  - **75** (threshold). The percentage of the maximum at which the router starts to generate a warning message, from 1 to 100. The default is 75%.
- **restart**. Stop the peering session with the neighbor when the limit is reached. Click the *restart-interval* variable and configure how long the system should wait before restarting the session, from 1 to 65535 minutes.
- **warning-only**. Do not stop the session when the limit is reached. Instead, simply issue a warning syslog message and continue the session.

Step 13 (Optional.) Configure routes in the neighbor activation settings.

a) If you have already configured it, click **...** > **Duplicate** for the **configure neighbor** *neighbor-address* **activate** *settings* command, or simply click + to enable it if it is not yet in use. Click **Show Disabled** if you cannot see the command.

b) Click *settings* and select **routes**.

c) On the **neighbor advertisement-interval** command, click the *value* variable and enter the minimum route advertisement interval between sending route updates to this neighbor, from 0 to 600 seconds, or click - to disable the command and leave the interval to default to 0, for iBGP and eBGP sessions in a virtual router, or 30 for eBGP sessions not in a virtual router. The value 0 means the system will send updates whenever the routing table changes, regardless of how frequent that might be.

d) On the **neighbor advertise-map** command, either configure the following options to conditionally advertise selected routes to the neighbor, or click - to disable the command, thus sending all route updates to the neighbor unconditionally.

The routes (prefixes) that will be conditionally advertised are defined in two route maps: an advertise map and either an exist map or non-exist map.

The route map associated with the exist map or non-exist map specifies the prefix that the BGP speaker will track.

The route map associated with the advertise map specifies the prefix that will be advertised to the specified neighbor when the condition is met.

If you configure an exist map, the condition is met when the prefix exists in both the advertise map and the exist map.

If you configure a non-exist map, the condition is met when the prefix exists in the advertise map, but does not exist in the non-exist map.

If the condition is not met, the route is withdrawn and conditional advertisement does not occur. All routes that may be dynamically advertised or not advertised need to exist in the BGP routing table for conditional advertisement to occur.

- *advertise-route-map*. Click this variable and select the route map that defines which routes should be advertised if the conditions of the exist map or non-exist map are met.
- *options condition-route-map*. Click options and select one of the following:
  - **exist-map**. Click the variable and select the exist route map.
  - **non-exist-map**. Click the variable and select the non-exist route map.

e) The **neighbor** *neighbor-address* **remove-private-as** command is added in the enabled state. Click - to disable the command. This command removes private autonomous system numbers from the eBGP outbound routing updates. The private AS values are 64512 to 65535.

f) On the **configure neighbor default-originate** command, either click *options* and select one of the following, or click - to disable the command.

- **none**. Allows the system to send a default route to the neighbor unconditionally.
- **route-map**. Have the system conditionally send a default route to the neighbor. When used with a route map, the default route is injected if the route map contains a match IP address clause and there is a route that matches the IP access list exactly. You can use standard or extended access lists in the route map to define the default routes. You must click the *route-map* variable in **neighbor default-originate** command that is added to the object and select the route map.

Step 14 (Optional.) Configure timers in the neighbor activation settings.

If you configure timers for a neighbor, the settings override the timers configured for all BGP neighbors in the global BGP settings.

a) If you have already configured it, click **...** > **Duplicate** for the **configure neighbor** *neighbor-address* **activate** *settings* command, or simply click **+** to enable it if it is not yet in use. Click **Show Disabled** if you cannot see the command.

b) Click *settings* and select **timers**.

c) On the **neighbors timers** command, configure the following variables:

- *keepalive-interval*. The frequency with which the system sends keepalive messages to this neighbor, from 0 to 65535 seconds. The default is 60 seconds if you do not configure this command.
- *hold-time*. The interval after not receiving a keepalive message that the system declares this neighbor dead, 0 to 65535 seconds. The default is 180 seconds if you do not configure this command.
- **0** (minimum hold time). The minimum acceptable hold-time that can be configured on this neighbor, from 0 to 65535 seconds. This value must be less than or equal to the hold time configured for this

system. If the neighbor's hold time is less than this value, the system will not establish a BGP session with the neighbor.

Step 15 (Optional.) Configure advanced neighbor activation settings.

a) If you have already configured it, click **...** > **Duplicate** for the **configure neighbor** *neighbor-address* **activate** *settings* command, or simply click **+** to enable it if it is not yet in use. Click **Show Disabled** if you cannot see the command.

b) Click *settings* and select **advanced**.

c) Decide which of the following **neighbor** commands to leave enabled. Click **-** to disable unwanted options.

- **send-community**. Sends the communities attribute to the neighbor.
- **weight** *value*. Click the variable to assign the initial weight to routes learned from this neighbor, from 0 and 65535. If you do not configure this command, routes learned through another BGP peer have a default weight of 0 and routes sourced by the local router have a default weight of 32768. However, any route weight that is set using a route map overrides the weight configured using this command.
- **next-hop-self**. Configures the router as the next hop for a BGP-speaking neighbor. This command is useful in unmeshed networks (such as Frame Relay or X.25) where BGP neighbors may not have direct access to all other neighbors on the same IP subnet.

Step 16 You can click **...** > **Duplicate** next to the **configure neighbor** command to define another neighbor. Define as many as you require.

Step 17 Click **OK**.

Configure BGP Route Redistribution from Other Routing Protocols

You can control the redistribution of routes into a BGP process from other routing protocols, connected routes, and static routes.

Before you begin

It is best practice to configure the routing process from which you will redistribute routes, and deploy your changes, before you configure redistribution into BGP.

If you want to apply a route map to fine-tune which routes are redistributed, create the Smart CLI route map object. Routes that match the route map are redistributed, and all non-matching routes are not redistributed.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring BGP.

Step 3 Click the **BGP** tab.

Step 4 Add or edit the BGP process object.

The redistribution commands are within the command sets beneath the **configure address family ipv4** or **ipv6** command. You must configure the address family to configure redistribution.

Step 5 Click **Show Disabled** to expose all commands, then click **+** to enable the **configure ipv4/ipv6 redistribution** command.

Step 6 Click the *protocol* variable and select the source process from which you are redistributing routes. You can redistribute **connected** and **static** routes, or routes generated by **eigrp** (IPv4 only), **isis**, **ospf**, or **rip** (IPv4 only).

Step 7 If you select a routing process, click the *identifier* variable and enter the required value:

- **eigrp**. Enter the autonomous system number.
- **ospf**. Enter the process ID number.
- **connected**, **static**, **isis**, **rip**. Enter **none**. Even if you enter a different value, it will be ignored.

Step 8 (Optional; IS-IS only.) On the **redistribute isis level-2** command, click **level-2** and select whether you are redistributing routes learned only within an IS-IS area (**level-1**), between IS-IS areas (**level-2**) or both (**level-1-2**).

Step 9 (Optional; all protocols.) To fine-tune the metrics for redistributed routes, click **+** to enable the following command and configure the options:

redistribute *protocol* **metric** *metric-value*

Click the variable and enter the metric value for the routes being distributed, from 0 to 4294967295.

Step 10 (Optional; all protocols.) To fine-tune which routes are redistributed based on a route map, click **+** to enable the **redistribute route-map** command, click the variable, and select the route map that defines your restrictions.

If you do not apply a route map, all routes for the process (that fit the other commands configured for redistribution), are redistributed.

Step 11 (Optional; OSPF only.) The following commands are enabled by default when you redistribute routes from an OSPF process. You can click **-** to disable unwanted commands.

These commands specify the criteria by which OSPF routes are redistributed into other routing domains.

- **redistribute ospf match external 1**. Routes that are external to the autonomous system, but are imported into OSPF as Type 1 external routes.
- **redistribute ospf match external 2**. Routes that are external to the autonomous system, but are imported into OSPF as Type 2 external routes.
- **redistribute ospf match internal**. Routes that are internal to a specific autonomous system.
- **redistribute ospf match nssa-external 1**. Routes that are external to the autonomous system, but are imported into OSPF as Type 1 external routes and marked as Not-So-Stubby-Area (NSSA) only.
- **redistribute ospf match nssa-external 2**. Routes that are external to the autonomous system, but are imported into OSPF as Type 2external routes and marked as Not-So-Stubby-Area (NSSA) only.

Step 12 You can click **...** > **Duplicate** next to the **configure redistribution** command to configure redistribution for another protocol. Configure redistribution for each protocol that makes sense for your network.

Step 13 Click **OK**.

Monitoring BGP

To monitor and troubleshoot BGP, open the CLI console or log into the device CLI and use the following commands. You can also select some of these commands from the **Commands** menu on the Routing page.

Use **show bgp ?** to get lists of additional options. For example, you can specify autonomous system number, and virtual router, to limit the information you see, as well as other options to target just the information you are looking for. The following list is a summary only.

- **show bgp**

  Displays the entries in the BGP routing table.

- **show bgp cidr-only**

  Displays routes with non-natural network masks (that is, classless interdomain routing, or CIDR).

- **show bgp community**

  Displays routes that belong to specified BGP communities.

- **show bgp community-list**

  Displays routes that are permitted by the BGP community list.

- **show bgp filter-list** *access-list-number*

  Displays routes that conform to a specified filter list.

- **show bgp injected-paths**

  Displays all the injected paths in the BGP routing table.

- **show bgp ipv4 unicast**

  Displays entries in the IP version 4 (IPv4) BGP routing table for unicast sessions.

- **show bgp ipv6 unicast**

  Displays entries in the IPv6 BGP routing table.

- **show bgp neighbors**

  Displays information about BGP and TCP connections to neighbors.

- **show bgp paths**

  Displays all the BGP paths in the database.

- **show bgp prefix-list**

  Displays information about a prefix list or prefix list entries.

- **show bgp regexp** *regexp*

  Displays routes that match the autonomous system path regular expression.

- **show bgp rib-failure**

  Displays BGP routes that failed to install in the Routing Information Base (RIB) table.

- **show bgp summary**

Display the status of all BGP connections.

- **show bgp update-group**

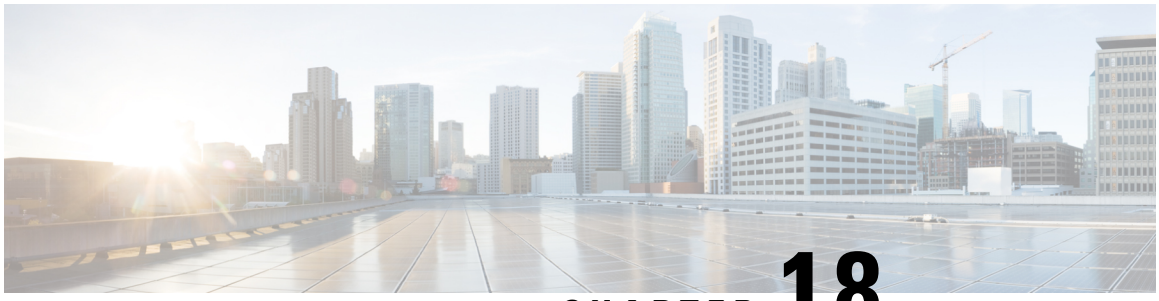
  Display information about the BGP update groups.



PART V

Security Policies

- SSL Decryption, on page 429
- Identity Policies, on page 453
- Security Intelligence, on page 467
- Access Control, on page 473
- Intrusion Policies, on page 509
- Network Address Translation (NAT), on page 539



CHAPTER 18

SSL Decryption

Some protocols, such as HTTPS, use Secure Sockets Layer (SSL) or its follow-on version, Transport Layer Security (TLS), to encrypt traffic for secure transmissions. Because the system cannot inspect encrypted connections, you must decrypt them if you want to apply access rules that consider higher-layer traffic characteristics to make access decisions.

- About SSL Decryption, on page 429
- License Requirements for SSL Decryption, on page 432
- Guidelines for SSL Decryption, on page 433
- How to Implement and Maintain the SSL Decryption Policy, on page 433
- Configuring SSL Decryption Policies, on page 435
- Example: Blocking Older SSL/TLS Versions from the Network, on page 447
- Monitoring and Troubleshooting SSL Decryption, on page 449

About SSL Decryption

Normally, connections go through the access control policy to determine if they are allowed or blocked. However, if you enable the SSL decryption policy, encrypted connections are first sent through the SSL decryption policy to determine if they should be decrypted or blocked. Any unblocked connections, whether or not decrypted, then go through the access control policy for a final allow/block decision.

Note You must enable the SSL decryption policy in order to implement active authentication rules in the identity policy. If you enable SSL decryption to enable identity policies, but do not otherwise want to implement SSL decryption, select Do Not Decrypt for the default action and do not create additional SSL decryption rules. The identity policy automatically generates whatever rules it needs.

The following topics explain encrypted traffic flow management and decryption in more detail.

Why Implement SSL Decryption?

Encrypted traffic, such as HTTPS connections, cannot be inspected.

Many connections are legitimately encrypted, such as connections to banks and other financial institutions. Many web sites use encryption to protect privacy or sensitive data. For example, your connection to the Firewall Device Manager is encrypted.

However, users can also hide undesirable traffic within encrypted connections.

By implementing SSL decryption, you can decrypt connections, inspect them to ensure they do not contain threats or other undesirable traffic, and then re-encrypt them before allowing the connection to proceed. (The decrypted traffic goes through your access control policy and matches rules based on inspected characteristics of the decrypted connection, not on the encrypted characteristics.) This balances your need to apply access control policies with the user's need to protect sensitive information.

You can also configure SSL decryption rules to block encrypted traffic of types you know you do not want on your network.

Keep in mind that decrypting and then re-encrypting traffic adds a processing load on the device, which will reduce overall system performance.

Actions You Can Apply to Encrypted Traffic

When configuring SSL decryption rules, you can apply the actions described in the following topics. These actions are also available for the default action, which applies to any traffic that does not match an explicit rule.

Note Any traffic that passes through the SSL decryption policy must then pass through the access control policy. Except for traffic you drop in the SSL decryption policy, the ultimate allow or drop decision rests with the access control policy.

Decrypt Re-Sign

If you elect to decrypt and re-sign traffic, the system acts as a man-in-the-middle.

For example, the user types in https://www.cisco.com in a browser. The traffic reaches the Firewall Threat Defense device, the device then negotiates with the user using the CA certificate specified in the rule and builds an SSL tunnel between the user and the Firewall Threat Defense device. At the same time the device connects to https://www.cisco.com and creates an SSL tunnel between the server and the Firewall Threat Defense device.

Thus, the user sees the CA certificate configured for the SSL decryption rule instead of the certificate from www.cisco.com. The user must trust the certificate to complete the connection. The Firewall Threat Defense device then performs decryption/re-encryption in both directions for traffic between the user and destination server.



Note If the client does not trust the 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.

If you configure a rule with the Decrypt Re-Sign 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 can select a single re-sign certificate for the SSL decryption policy, this can limit traffic matching for resign rules.

For example, outgoing traffic encrypted with an elliptic curve (EC) algorithm matches a Decrypt Re-Sign rule only if the re-sign certificate is an EC-based CA certificate. Similarly, traffic encrypted with an RSA algorithm matches Decrypt Re-Sign rules only if the global re-sign certificate is RSA; outgoing traffic encrypted with an EC algorithm does not match the rule, even if all other configured rule conditions match.

Decrypt Known Key

If you own the destination server, you can implement decryption with a known key. In this case, when the user opens a connection to https://www.cisco.com, the user sees the actual certificate for www.cisco.com, even though it is the Firewall Threat Defense device that is presenting the certificate.



Your organization must be the owner of the domain and certificate. For the example of cisco.com the only possible way to have the end user see Cisco's certificate would be if you actually own the domain cisco.com (i.e. you are Cisco Systems) and have ownership of the cisco.com certificate signed by a public CA. You can only decrypt with known keys for sites that your organization owns.

The main purpose of decrypting with a known key is to decrypt traffic heading to your HTTPS server to protect your servers from external attacks. For inspecting client side traffic to external HTTPS sites, you must use decrypt re-sign as you do not own the servers.

Note To use known key decryption, you must upload the server's certificate and key as an internal identity certificate, and then add it to the list of known-key certificates in the SSL decryption policy settings. Then, you can write the rule for known-key decryption with the server's address as the destination address. For information on adding the certificate to the SSL decryption policy, see Configure Certificates for Known Key and Re-Sign Decryption, on page 445.

Do Not Decrypt

If you elect to bypass decryption for certain types of traffic, no processing is done on the traffic. The encrypted traffic proceeds to the access control policy, where it is allowed or dropped based on the access control rule it matches.

Block

You can simply block encrypted traffic that matches an SSL decryption rule. Blocking in the SSL decryption policy prevents the connection from reaching the access control policy.

When you block an HTTPS connection, the user does not see the system default block response page. Instead, the user sees the browser's default page for a secure connection failure. 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 will not be obvious from this message that you blocked the connection on purpose.

Automatically Generated SSL Decryption Rules

Whether you enable the SSL decryption policy, the system automatically generates Decrypt Re-sign rules for each identity policy rule that implements active authentication. This is required to enable active authentication for HTTPS connections.

When you enable the SSL decryption policy, you see these rules under the Identity Policy Active Authentication Rules heading. These rules are grouped at the top of the SSL decryption policy. The rules are read only. You can change them only by altering your identity policy.

Handling Undecryptable Traffic

There are several characteristics that make a connection undecryptable. If a connection has any of the following characteristics, the default action is applied to the connection regardless of any rule the connection would otherwise match. If you select Block as your default action (rather than Do Not Decrypt), you might run into issues, including excessive drops of legitimate traffic.

- Compressed session—Data compression was applied to the connection.
- SSLv2 session—The minimum supported SSL version is SSLv3.
- Unknown cipher suite—The system does not recognize the cipher suite for the connection.
- Unsupported cipher suite—The system does not support decryption based on the detected cipher suite.
- Session not cached—The 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.
- Handshake errors—An error occurred during the SSL handshake negotiation.
- Decryption errors—An error occurred during the decryption operation.
- Passive interface traffic—All traffic on passive interfaces (passive security zones) is undecryptable.

License Requirements for SSL Decryption

You do not need a special license to use the SSL decryption policy.

However, you do need the **URL** license to create rules that use URL categories and reputations as match criteria. For information on configuring licenses, see Enabling or Disabling Optional Licenses, on page 89.

Guidelines for SSL Decryption

Keep the following in mind when configuring and monitoring SSL decryption policies:

- The SSL Decryption policy is bypassed for any connections that match access control rules set to trust or block if those rules:
  - Use security zone, network, geolocation, and port only as the traffic matching criteria.
  - Come before any 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.
- When using URL category matching, note that there are cases where the login page for a site is in a different category than the site itself. For example, Gmail is in the "Web based email" category, whereas the login page is in the "Internet Portals" category. To get connections to these sites decrypted, you must include both categories in the rule.
- If a Vulnerability Database (VDB) update removes (deprecates) applications, you must make changes to any SSL decryption rules or application filters that use the application that was deleted. You cannot deploy changes until you fix these rules. In addition, you cannot install system software updates before fixing the issue. On the Application Filters object page, or the Application tab of the rule, these applications say "(Deprecated)" after the application name.
- You cannot disable the SSL decryption policy if you have any active authentication rules. To disable the SSL decryption policy, you must either disable the identity policy, or delete any identity rules that use active authentication.

How to Implement and Maintain the SSL Decryption Policy

You can use SSL decryption policies to turn encrypted traffic into plain text traffic, so that you can then apply URL filtering, intrusion and malware control, and other services that require deep packet inspection. If your policies allow the traffic, the traffic is re-encrypted before it leaves the device.

The SSL decryption policy applies to encrypted traffic only. No unencrypted connections are evaluated against SSL decryption rules.

Unlike some other security policies, you need to monitor and actively maintain the SSL decryption policy, because certificates can expire or even be changed on destination servers. Additionally, changes in client software might alter your ability to decrypt certain connections, because the decrypt re-sign action is indistinguishable from a man-in-the-middle attack.

The following procedure explains the end-to-end process of implementing and maintaining the SSL decryption policy.

Procedure

Step 1 If you will implement Decrypt Re-sign rules, create the required internal CA certificate.

You must use an internal Certificate Authority (CA) certificate. You have the following options. Because users must trust the certificate, either upload a certificate client browsers are already configured to trust, or ensure that the certificate you upload is added to the browser trust stores.

- Create a self-signed internal CA certificate, which is signed by the device itself. See Generating Self-Signed Internal and Internal CA Certificates, on page 153.
- Upload an internal CA certificate and key signed by an external trusted CA or by a CA inside your organization. See Uploading Internal and Internal CA Certificates, on page 151.

Step 2 If you will implement Decrypt Known Key rules, collect the certificate and key from each of the internal servers.

You can use Decrypt Known Key only with servers that you control, because you must obtain the certificate and key from the server. Upload these certificates and keys as internal certificates (not internal CA certificates). See Uploading Internal and Internal CA Certificates, on page 151.

Step 3 Enable the SSL Decryption Policy, on page 436.

When you enable the policy, you also configure some basic settings.

Step 4 Configure the Default SSL Decryption Action, on page 437.

If in doubt, select **Do Not Decrypt** as the default action. Your access control policy can still drop traffic that matches the default SSL decryption rule if appropriate.

Step 5 Configure SSL Decryption Rules, on page 438.

Identify traffic to decrypt and the type of decryption to apply.

Step 6 If you configure known key decryption, edit the SSL decryption policy settings to include those certificates. See Configure Certificates for Known Key and Re-Sign Decryption, on page 445.

Step 7 If necessary, download the CA certificate used for Decrypt Re-sign rules and upload it to the browser on client workstations.

For information on downloading the certificate and distributing it to clients, see Downloading the CA Certificate for Decrypt Re-Sign Rules, on page 446.

Step 8 Periodically, update re-sign and known key certificates.

- Re-sign certificate—Update this certificate before it expires. If you generate the certificate through the Firewall Device Manager, it is valid for 5 years. To check the validity period for a certificate, select **Objects** > **Certificates**, find the certificate in the list, and click the information icon ( ) for it in the Actions column. The information dialog box shows the validity period and some other characteristics. You can also upload a replacement certificate from this page.
- Known-key certificate—For any known-key decryption rules, you need to ensure that you have uploaded the destination server's current certificate and key. Whenever the certificate and key changes on supported servers, you must also upload the new certificate and key (as an internal certificate) and update the SSL decryption settings to use the new certificate.

Step 9 Upload missing trusted CA certificates for external servers.

The system includes a wide range of trusted CA root and intermediate certificates issued by third parties. These are needed when negotiating the connection between the Firewall Threat Defense and the destination servers for decrypt re-sign rules.

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. Upload certificates on the **Objects** > **Certificates** page. See Uploading Trusted CA Certificates, on page 154.

Configuring SSL Decryption Policies

You can use SSL decryption policies to turn encrypted traffic into plain text traffic, so that you can then apply URL filtering, intrusion and malware control, and other services that require deep packet inspection. If your policies allow the traffic, the traffic is re-encrypted before it leaves the device.

The SSL decryption policy applies to encrypted traffic only. No unencrypted connections are evaluated against SSL decryption rules.

Note VPN tunnels are decrypted before the SSL decryption policy is evaluated, so the policy never applies to the tunnel itself. However, any encrypted connections within the tunnel are subject to evaluation by the SSL decryption policy.

The following procedure explains how to configure the SSL decryption policy. For an explanation of the end-to-end process of creating and managing SSL decryption, see How to Implement and Maintain the SSL Decryption Policy, on page 433.

Before you begin

The SSL decryption rules table contains two sections:

- **Identity Policy Active Authentication Rules**—If you enable the identity policy and create rules that use active authentication, the system automatically creates the SSL decryption rules needed to make those policies work. These rules are always evaluated before the SSL decryption rules you create yourself. You can alter these rules only indirectly, by making changes to the identity policy.
- **SSL Native Rules**—These are rules that you have configured. You can add rules to this section only.

Procedure

Step 1 Select **Policies** > **SSL Decryption**.

If you have not yet enabled the policy, click **Enable SSL Decryption** and configure policy settings, as described in Enable the SSL Decryption Policy, on page 436.

Step 2 Configure the default action for the policy.

The safest choice is **Do Not Decrypt**. For more information, see Configure the Default SSL Decryption Action, on page 437.

Step 3 Manage the SSL decryption policy.

After you configure SSL decryption settings, this page lists all rules in order. Rules are matched against traffic from top to bottom with the first match determining the action to apply. You can do the following from this page:

- To disable the policy, click the **SSL Decryption Policy** toggle. You can re-enable it by clicking **Enable SSL Decryption**.
- To edit policy settings, including the list of certificates used in the policy, click the **SSL Decryption Settings** button (). You can also download the certificate used with decrypt re-sign rules so that you can distribute it to clients. See the following topics:
  - Configure Certificates for Known Key and Re-Sign Decryption, on page 445
  - Downloading the CA Certificate for Decrypt Re-Sign Rules, on page 446
- To configure rules:
  - To create a new rule, click the + button. See Configure SSL Decryption Rules, on page 438.
  - To edit an existing rule, click the edit icon () for the rule (in the Actions column). You can also selectively edit a rule property by clicking on the property in the table.
  - To delete a rule you no longer need, click the delete icon () for the rule (in the Actions column).
- To move a rule, edit it and select the new location from the **Order** drop-down list.
- If any rules have problems, for example, because of removed or changed URL categories, click the **See Problem Rules** link next to the search box to filter the table to show only those rules. Please edit and correct (or delete) these rules, so that they will provide the service that you require.

Enable the SSL Decryption Policy

Before you can configure SSL decryption rules, you must enable the policy and configure some basic settings. The following procedure explains how to enable the policy directly. You can also enable it when you enable identity policies. Identity policies require that you enable the SSL decryption policy.

Before you begin

If you upgraded from a release that did not have SSL decryption policies, but you had configured the identity policy with active authentication rules, the SSL decryption policy is already enabled. Ensure that you select the Decrypt Re-Sign certificate you want to use, and optionally enable pre-defined rules.

Procedure

Step 1 Select **Policies** > **SSL Decryption**.

Step 2 Click **Enable SSL Decryption** to configure the policy settings.

- If this is the first time you enabled the policy, the SSL Decryption Configuration dialog box opens. Proceed with the next steps.

- If you have already configured the policy once and then disabled it, the policy is simply enabled again with your previous settings and rules. You can click the **SSL Decryption Settings** button () and configure settings as described in Configure Certificates for Known Key and Re-Sign Decryption, on page 445.

Step 3 In **Decrypt Re-Sign Certificate**, select the internal CA certificate to use for rules that implement decryption with re-signed certificates.

You can use the pre-defined NGFW-Default-InternalCA certificate, or one that you created or uploaded. If the certificate does not yet exist, click **Create Internal CA** to create it.

If you have not already installed the certificate in client browsers, click the download button () to obtain a copy. See the documentation for each browser for information on how to install the certificate. Also see Downloading the CA Certificate for Decrypt Re-Sign Rules, on page 446.

Step 4 (Optional.) Click **+** under **Trusted CA Certificates** and select the certificates or certificate groups that you want the policy to trust.

The default group, Cisco-Trusted-Authorities, includes all of the system-defined trusted CA certificates. If you uploaded additional certificates, you can add them here, or collect them in your own group and select the group here. You can either replace the Cisco-Trusted-Authorities group, or simply add your group. Users will be prompted to accept the certificate for any site whose certificate's signing authority is not represented in this list: access to the site is not blocked simply because the certificate is not trusted.

If you leave the list empty, or select only empty certificate groups, then the SSL decryption policy will trust any certificate.

Step 5 Select the initial SSL decryption rules.

The system includes the following pre-defined rule that you might find useful:

- **Sensitive_Data**—This rule does not decrypt traffic that matches web sites in the Financial Services or Health and Medicine URL categories, which include banks, healthcare services, and so forth. You must enable the URL license to implement this rule.

Step 6 Click **Enable**.

Configure the Default SSL Decryption Action

If an encrypted connection does not match a specific SSL decryption rule, it is handled by the default action for the SSL decryption policy.

Procedure

Step 1 Select **Policies** > **SSL Decryption**.

Step 2 Click anywhere in the **Default Action** field.

Step 3 Select the action to apply to matching traffic.

- **Do Not Decrypt**—Allow the encrypted connection. The access control policy then evaluates the encrypted connection and drops or allows it based on access control rules.

- **Block**—Drop the connection immediately. The connection is not passed on to the access control policy.

Step 4 (Optional.) Configure logging for the default action.

You must enable logging for traffic that matches the default action to be included in dashboard data or Event Viewer. Select from these options:

- **At End of Connection**—Generate an event at the conclusion of the connection.
  - **Send Connection Events To**—If you want to send a copy of the events to an external syslog server, select the server object that defines the syslog server. If the required object does not already exist, click **Create New Syslog Server** and create it. (To disable logging to a syslog server, select **Any** from the server list.)

    Because event storage on the device is limited, sending events to an external syslog server can provide more long term storage and enhance your event analysis.

- **No Logging**—Do not generate any events.

Step 5 Click **Save**.

Configure SSL Decryption Rules

Use SSL decryption rules to determine how to handle encrypted connections. Rules in the SSL decryption policy are evaluated from top to bottom. The rule applied to traffic is the first one where all the traffic criteria are matched.

You can create and edit rules in the SSL Native Rules section only.

Note Traffic for your VPN connections (both site-to-site and remote access) is decrypted before the SSL decryption policy evaluates connections. Thus, SSL decryption rules are never applied to VPN connections, and you do not need to consider VPN connections when creating these rules. However, any use of encrypted connections within a VPN tunnel are evaluated. For example, an HTTPS connection to an internal server through an RA VPN connection is evaluated by SSL decryption rules, even though the RA VPN tunnel itself is not (because it is decrypted already).

Before you begin

If you are creating a decrypt known-key rule, ensure that you upload the certificate and key for the destination server (as an internal certificate) and also edit the SSL decryption policy settings to use the certificate. Known-key rules typically specify the destination server in the destination network criteria of the rule. For more information, see Configure Certificates for Known Key and Re-Sign Decryption, on page 445.

Procedure

Step 1 Select **Policies** > **SSL Decryption**.

If you have not configured any SSL decryption rules (other than the ones automatically generated for active authentication identity rules), you can add pre-defined rules by clicking **Add Pre-Defined Rules**. You are prompted to select the rules that you want.

Step 2 Do any of the following:

- To create a new rule, click the + button.
- To edit an existing rule, click the edit icon ( ) for the rule.

To delete a rule you no longer need, click the delete icon ( ) for the rule.

Step 3 In **Order**, select where you want to insert the rule in the ordered list of rules.

You can insert rules into the **SSL Native Rules** section only. The Identity Policy Active Authentication Rules are automatically generated from your identity policy and are read-only.

Rules are applied on a first-match basis, so you must ensure that rules with highly specific traffic matching criteria appear above policies that have more general criteria that would otherwise apply to the matching traffic.

The default is to add the rule to the end of the list. If you want to change a rule's location later, edit this option.

Step 4 In **Title**, enter a name for the rule.

The name cannot contain spaces. You can use alphanumeric characters and these special characters: + . _ -

Step 5 Select the action to apply to matching traffic.

For a detailed discussion of each option, see the following:

- Decrypt Re-Sign, on page 430
- Decrypt Known Key, on page 431
- Do Not Decrypt, on page 431
- Block, on page 432

Step 6 Define the traffic matching criteria using any combination of the following tabs:

- **Source/Destination**—The security zones (interfaces) through which the traffic passes, the IP addresses or the country or continent (geographical location) for the IP address, or the TCP ports used in the traffic. The default is any zone, address, geographical location, and TCP port. See Source/Destination Criteria for SSL Decryption Rules, on page 440.
- **Application**—The application, or a filter that defines applications by type, category, tag, risk, or business relevance. The default is any encrypted application. See Application Criteria for SSL Decryption Rules, on page 441.
- **URL**—The URL category of a web request. The default is that the URL category and reputation are not considered for matching purposes. See URL Criteria for SSL Decryption Rules, on page 442.
- **Users**—The identity source, user or user group. Your identity policies determine whether user and group information is available for traffic matching. You must configure identity policies to use this criteria. See User Criteria for SSL Decryption Rules, on page 443.
- **Advanced**—The characteristics derived from the certificates used in the connection, such as SSL/TLS version and certificate status. See Advanced Criteria for SSL Decryption Rules, on page 444.

To modify a condition, you click the + button within that condition, select the desired object or element, and click **OK** in the popup dialog box. If the criterion requires an object, you can click **Create New *Object*** if the object you require does not exist. Click the **x** for an object or element to remove it from the policy.

When adding conditions to SSL decryption rules, consider the following tips:

- 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, you can use a single rule to decrypt traffic based on URL category.
- 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 apply application control for up to 50 applications or application filters. Thus, there is an OR relationship among the items in a single condition, but an AND relationship between condition types (for example, between source/destination and application).
- Matching URL category requires the URL filtering license.

Step 7 (Optional.) Configure logging for the rule.

You must enable logging for traffic that matches the rule to be included in dashboard data or Event Viewer. Select from these options:

- **At End of Connection**—Generate an event at the conclusion of the connection.
  - **Send Connection Events To**—If you want to send a copy of the events to an external syslog server, select the server object that defines the syslog server. If the required object does not already exist, click **Create New Syslog Server** and create it. (To disable logging to a syslog server, select Any from the server list.)

    Because event storage on the device is limited, sending events to an external syslog server can provide more long term storage and enhance your event analysis.
- **No Logging**—Do not generate any events.

Step 8 Click **OK**.

Source/Destination Criteria for SSL Decryption Rules

The Source/Destination criteria of an SSL decryption rule define the security zones (interfaces) through which the traffic passes, the IP addresses or the country or continent (geographical location) for the IP address, or the TCP ports used in the traffic. The default is any zone, address, geographical location, and any TCP port. TCP is the only protocol matched to SSL decryption rules.

To modify a condition, you click the + button within that condition, select the desired object or element, and click **OK**. If the criterion requires an object, you can click **Create New *Object*** if the object you require does not exist. Click the **x** for an object or element to remove it from the policy.

You can use the following criteria to identify the source and destination to match in the rule.

Source Zones, Destination Zones

The security zone objects that define the interfaces through which the traffic passes. You can define one, both, or neither criteria: any criteria not specified applies to traffic on any interface.

- To match traffic leaving the device from an interface in the zone, add that zone to the **Destination Zones**.

- To match 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.

Use this criteria when the rule should apply based on where the traffic enters or exits the device. For example, if you want to ensure that all traffic going from outside hosts to inside hosts gets decrypted, you would select your outside zone as the **Source Zones** and your inside zone as the **Destination Zones**.

Source Networks, Destination Networks

The network objects or geographical locations that define the network addresses or locations of the traffic.

- To match traffic from an IP address or geographical location, configure the **Source Networks**.
- To match 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 traffic must originate from one of the specified IP addresses and be destined for one of the destination IP addresses.

When you add this criteria, you select from the following tabs:

- **Network**—Select the network objects or groups that define the source or destination IP addresses for the traffic you want to control.



Note For Decrypt Known-Key rules, select an object with the IP address of the destination server that uses the certificate and key you uploaded.

- **Geolocation**—Select the geographical location to control traffic based on its source or destination country or continent. Selecting a continent selects all countries within the continent. Besides selecting geographical location directly in the rule, you can also select a geolocation object that you created to define the location. Using geographical location, you could easily restrict access to a particular country without needing to know all of the potential IP addresses used there.

Source Ports, Destination Ports/Protocols

The port objects that define the protocols used in the traffic. You can specify TCP protocol and ports only for SSL decryption rules.

- To match traffic from a TCP port, configure the **Source Ports**.
- To match traffic to a TCP port, configure the **Destination Ports/Protocols**.
- To match traffic both originating from specific TCP ports and destined for specific TCP ports, configure both. For example, you could target traffic from port TCP/80 to port TCP/8080.

Application Criteria for SSL Decryption Rules

The Application criteria of an SSL decryption rule defines the application used in an IP connection, or a filter that defines applications by type, category, tag, risk, or business relevance. The default is any application that has the SSL Protocol tag. You cannot match SSL decryption rules to any non-encrypted application.

Although you can specify individual applications in the rule, application filters simplify policy creation and administration. For example, you could create an SSL decryption rule that decrypts or blocks all high risk, low business relevance applications. If a user attempts to use one of those applications, the session is decrypted or blocked.

In addition, Cisco frequently updates and adds additional application detectors via system and vulnerability database (VDB) updates. Thus, a rule for high risk applications can automatically apply to new applications without you having to update the rule manually.

You can specify applications and filters directly in the rule, or create application filter objects that define those characteristics. The specifications are equivalent, although using objects can make it easier to stay within the 50-items-per-criteria system limit if you are creating a complex rule.

To modify the application and filters list, you click the + button within the condition, select the desired applications or application filter objects, which are listed on separate tabs, and click **OK** in the popup dialog box. On either tab, you can click **Advanced Filter** to select filter criteria or to help you search for specific applications. Click the **x** for an application, filter, or object to remove it from the policy. Click the **Save As Filter** link to save the combined criteria that is not already an object as a new application filter object.

For more information about the application criteria and how to configure advanced filters and select applications, see Configuring Application Filter Objects, on page 139.

Consider the following tips when using application criteria in SSL decryption rules.

- 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.
- The system can identify the application only after the server certificate exchange. If traffic exchanged during the SSL handshake matches all other conditions in an 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. After the system completes its identification, the system applies the SSL rule action to the remaining session traffic that matches its application condition.
- If a selected application was removed by a VDB update, "(Deprecated)" appears after the application name. You must remove these applications from the filter, or subsequent deployments and system software upgrades will be blocked.

URL Criteria for SSL Decryption Rules

The URL criteria of an SSL decryption rule defines the category to which the URL in a web request belongs. You can also specify the relative reputation of sites to decrypt, block, or allow without decryption. The default is to not match connections based on URL categories.

For example, you could block all encrypted Gambling sites, or decrypt untrusted Social Networking sites. If a user attempts to browse to any URL with that category and reputation combination, the session is blocked or decrypted. For more information on URL category matching, see Filtering URLs by Category and Reputation, on page 478.

Categories Tab

Click +, select the desired categories, and click **OK**. Click the **x** for a category to remove it from the policy.

The default is to apply the rule to all URLs in each selected category regardless of reputation. To limit the rule based on reputation, click the down arrow for each category, deselect the **Any** checkbox, and then use the **Reputation** slider to choose the reputation level. The left of the reputation slider indicates sites that will be allowed without decryption, the right side are sites that will be decrypted or blocked. How reputation is used depends on the rule action:

- If the rule decrypts or blocks connections, selecting a reputation level also selects all reputations more severe than that level. For example, if you configure a rule to decrypt or block **Questionable sites** (level 2), it also automatically decrypts or blocks **Untrusted** (level 1) sites.
- If the rule allows connections without decryption (do not decrypt), selecting a reputation level also selects all reputations less severe than that level. For example, if you configure a rule to not decrypt **Favorable sites** (level 4), it also automatically does not decrypt **Trusted** (level 5) sites.

Select the **Include Sites with Unknown Reputation** option to have URLs with unknown reputation included in the reputation match. New sites typically are unrated, and there can be other reasons a site's reputation is unknown or cannot be determined.

Check the Category for a URL

You can check on the category and reputation for a particular URL. Enter the URL in the **URL to Check** box and click **Go**. You will be taken to an external website to see the results. If you disagree with a categorization, click the **Submit a URL Category Dispute** link and let us know.

User Criteria for SSL Decryption Rules

The User criteria of an SSL decryption rule defines the user or user group for an IP connection. You must configure identity policies and the associated directory server to include user or user group criteria in a rule.

Your identity policies determine whether user identity is collected for a particular connection. If identity is established, the IP address of the host is associated with the identified user. Thus, traffic whose source IP address is mapped to a user is considered to be from that user. IP packets themselves do not include user identity information, so this IP-address-to-user mapping is the best approximation available.

Because you can add a maximum of 50 users or groups to a rule, selecting groups usually makes more sense than selecting individual users. For example, you could create a rule that decrypts traffic to the Engineering group that comes from the outside network, and create a separate rule that does not decrypt outgoing traffic from that group. Then, to make the rule apply to new engineers, you only need to add the engineer to the Engineering group in the directory server.

You an also select identity sources to apply to all users within that source. Thus, if you support multiple Active Directory domains, you can provide differential decryption based on the domain.

To modify the users list, you click the **+** button within the condition and select the desired users or user groups using one of the following techniques. Click the **x** for a user or group to remove it from the policy.

- **Identity Sources**—Select an identity source, such as an AD realm or the local user database, to apply the rule to all users obtained from the selected sources. If the realm you need does not yet exist, click **Create New Identity Realm** and create it now.
- **Groups**—Select the desired user groups. Groups are available only if you configure groups in the directory server. If you select a group, the rule applies to any member of the group, including subgroups. If you want to treat a sub-group differently, you need to create a separate access rule for the sub-group and place it above the rule for the parent group in the access control policy.
- **Users**—Select individual users. The user name is prefixed with the identity source, such as Realm\username.

There are some built-in users under the Special-Identities-Realm:

- **Failed Authentication**—The user was prompted to authenticate, but failed to enter a valid username/password pair within the maximum number of allowed attempts. Failure to authenticate does not itself prevent the user from accessing the network, but you can write an access rule to limit network access for these users.
- **Guest**—Guest users are like Failed Authentication users, except that your identity rule is configured to call these users Guest. Guest users were prompted to authenticate and failed to do so within the maximum number of attempts.
- **No Authentication Required**—The user was not prompted to authentication, because the user's connections matched identity rules that specified no authentication.
- **Unknown**—There is no user mapping for the IP address, and there is no record of failed authentication yet. Typically, this means that no HTTP traffic has yet been seen from that address.

Advanced Criteria for SSL Decryption Rules

The Advanced traffic matching criteria relate to characteristics derived from the certificates used in the connection. You can configure any or all of the following options.

Certificate Properties

Traffic matches the certificate properties option of the rule if it matches any of the selected properties. You can configure the following:

Certificate Status

Whether the certificate is **Valid** or **Invalid**. Select **Any** (the default) if you do not care about certificate status.

A certificate is considered valid if all of the following conditions are met, otherwise it is invalid:

- The policy trusts the CA that issued the certificate.
- The certificate’s signature can be properly validated against the certificate’s content.
- The issuer CA certificate is stored in the policy’s list of trusted CA certificates.
- None of the policy’s trusted CAs revoked the certificate.
- The current date is between the certificate Valid From and Valid To dates.

Self-Signed

Whether the server certificate contains the same subject and issuer distinguished name. Select one of the following:

- **Self-Signing**—The server certificate is self-signed.
- **CA-Signing**—The server certificate is signed by a Certificate Authority. That is, the issuer and subject are not the same.
- **Any**—Do not consider whether the certificate is self-signed as a match criteria.

Supported Version

The SSL/TLS version to match. The rule applies to traffic that uses the any of the selected versions only. The default is all versions. Select from: **SSL 3.0**, **TLS 1.0**, **TLS 1.1**, **TLS 1.2**.

For example, if you wanted to permit TLSv1.2 connections only, you could create a block rule for the lower versions.

Traffic that uses any version not listed, such as SSL v2.0, is handled by the default action for the SSL decryption policy.

Configure Certificates for Known Key and Re-Sign Decryption

If you implement decryption, either by re-signing or using known keys, you need to identify the certificates that the SSL decryption rules can use. Ensure that all certificates are valid and unexpired.

Especially for known-key decryption, you need to ensure that the system has the current certificate and key for each destination server whose connections you are decrypting. With a decrypt known key rule, you use the actual certificate and key from the destination server for decryption. Thus, you must ensure that the Firewall Threat Defense device has the current certificate and key at all times, or decryption will be unsuccessful.

Upload a new internal certificate and key whenever you change the certificate or key on the destination server in a known key rule. Upload them as an internal certificate (not an internal CA certificate). You can upload the certificate during the following procedure, or go to the **Objects** > **Certificates** page and upload it there.

Procedure

Step 1 Select **Policies** > **SSL Decryption**.

Step 2 Click the **SSL Decryption Settings** button ( ).

Step 3 In **Decrypt Re-Sign Certificate**, select the internal CA certificate to use for rules that implement decryption with re-signed certificates.

You can use the pre-defined NGFW-Default-InternalCA certificate, or one that you created or uploaded. If the certificate does not yet exist, click **Create Internal CA** to create it.

If you have not already installed the certificate in client browsers, click the download button ( ) to obtain a copy. See the documentation for each browser for information on how to install the certificate. Also see Downloading the CA Certificate for Decrypt Re-Sign Rules, on page 446.

Step 4 For each rule that decrypts using a known key, upload the internal certificate and key for the destination server.

a) Click + under **Decrypt Known-Key Certificates**.
b) Select the internal identity certificate, or click **Create New Internal Certificate** to upload it now.
c) Click **OK**.

Step 5 (Optional.) Click + under **Trusted CA Certificates** and select the certificates or certificate groups that you want the policy to trust.

The default group, Cisco-Trusted-Authorities, includes all of the system-defined trusted CA certificates. Following are the main cases where you might want to change this setting:

- You want to use trusted CA certificates that are not in the default group. You would then select both the default group and your new group in the SSL decryption policy settings. You might do this if you have uploaded additional trusted CA certificates.
- You want to use a more limited list of trusted CA certificates than what is in the default group. You would then create a group that has a complete list of trusted certificates, not just your delta, and select it as the sole group in the SSL decryption policy settings.

Users will be prompted to accept the certificate for any site whose certificate's signing authority is not represented in this list: access to the site is not blocked simply because the certificate is not trusted.

If you leave the list empty, or select only empty certificate groups, then the SSL decryption policy will trust any certificate.

Step 6 Click **Save**.

Downloading the CA Certificate for Decrypt Re-Sign Rules

If you decide to decrypt traffic, users must have the internal CA certificate that is used in the encryption process defined as a Trusted Root Certificate Authority in their applications that use TLS/SSL. Typically if you generate a certificate, or sometimes even if you import one, the certificate is not already defined as trusted in these applications. By default in most web browsers, when users send HTTPS requests, they will see a warning message from the client application informing them that there is a problem with the web site's security certificate. Usually, the error message says that the web site's security certificate was not issued by a trusted certificate authority or the web site was certified by an unknown authority, but the warning might also suggest there is a possible man-in-the-middle attack in progress. Some other client applications do not show this warning message to users nor allow users to accept the unrecognized certificate.

You have the following options for providing users with the required certificate:

Inform users to accept the root certificate

You can inform the users in your organization what the new policies are at the company and tell them to accept the root certificate supplied by the organization as a trusted source. Users should accept the certificate and save it in the Trusted Root Certificate Authority storage area so that they are not prompted again the next time they access the site.



Note The user needs to accept and trust the CA certificate that created the replacement certificate. If they instead simply trust the replacement server certificate, they will continue to see warnings for each different HTTPS site that they visit.

Add the root certificate to client devices

You can add the root certificate to all client devices on the network as a trusted root certificate authority. This way, the client applications automatically accept transactions with the root certificate.

You can either make the certificate available to users by E-mailing it or placing it on a shared site, or you could incorporate the certificate into your corporate workstation image and use your application update facilities to distribute it automatically to users.

The following procedure explains how to download the internal CA certificate and install it on Windows clients.

Procedure

Step 1 Download the certificate from the Firewall Device Manager.

a) Select **Policies** > **SSL Decryption**.

b) Click the **SSL Decryption Settings** button ( ).

c) Click the **Download** button ( ).

d) Select a download location, optionally change the file name (but not the extension), and click **Save**.

You can now cancel out of the SSL Decryption Settings dialog box.

Step 2 Install the certificate in the Trusted Root Certificate Authority storage area in web browsers on client systems, or make it available for clients to install themselves.

The process differs depending on the operating system and type of browser. For example, you can use the following process for Internet Explorer and Chrome running on Windows. (For Firefox, install through the **Tools** > **Options** > **Advanced** page.)

a) From the **Start** menu, select **Control Panel** > **Internet Options**.

b) Select the **Content** tab.

c) Click the **Certificates** button to open the Certificates dialog box.

d) Select the **Trusted Root Certificate Authorities** tab.

e) Click **Import**, and follow the wizard to locate and select the downloaded file (<uuid>_internalCA.crt) and add it to the Trusted Root Certificate Authorities store.

f) Click **Finish**.

Messages should indicate that the import was successful. You might see an intermediate dialog box warning you that Windows could not validate the certificate if you generated a self-signed certificate rather than obtaining one from a well-known third-party Certificate Authority.

You can now close out the Certificate and Internet Options dialog boxes.

Example: Blocking Older SSL/TLS Versions from the Network

Some organizations are required to prevent the use of older versions of SSL or TLS either by government regulation or company policy. You can use the SSL Decryption policy to block traffic that uses an SSL/TLS version that you prohibit. Consider placing this rule at the top of the SSL Decryption policy to ensure that you catch the prohibited traffic immediately.

The following example blocks all SSL 3.0 and TLS 1.0 connections.

Before you begin

This procedure assumes you have already enabled the SSL Decryption policy as explained in Enable the SSL Decryption Policy, on page 436.

Procedure

Step 1 Select **Policies** > **SSL Decryption**.

Step 2 Click the + button to create a new rule.

Step 3 In Order, select **1** to place the rule at the top of the policy, or select the number most suitable for your network.

The default is to add the rule at the end of the policy.

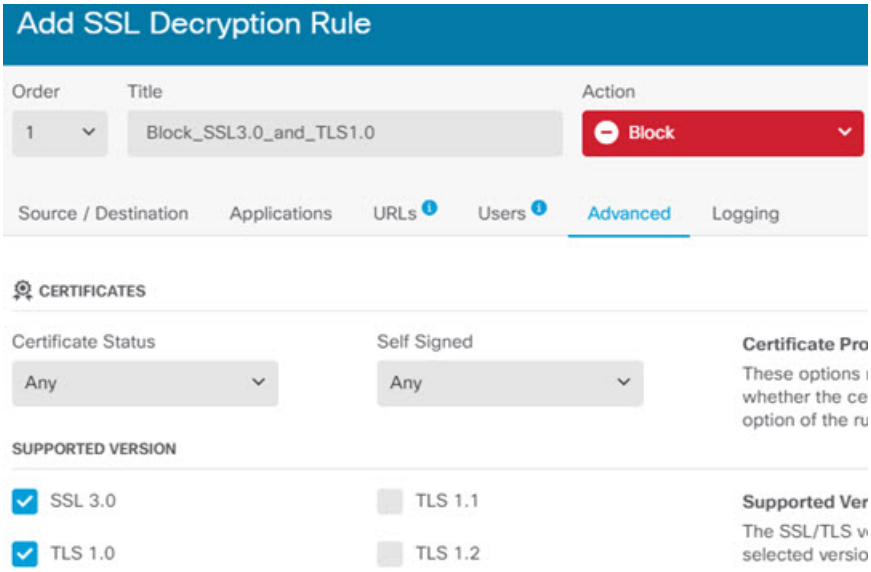
Step 4 In **Title**, enter a name for the rule, for example, Block_SSL3.0_and_TLS1.0.

Step 5 In **Action**, select **Block**. This will immediately drop any traffic that matches the rule.

Step 6 Leave the default values for all options on the following tabs: **Source/Destination**, **Applications**, **URLs**, **Users**.

Step 7 Click the **Advanced** tab and under **Supported Versions**, leave SSL 3.0 and TLS 1.0 selected, but uncheck TLS 1.1, TLS 1.2.

The policy should look like the following:



Step 8 (Optional) Click the **Logging** tab and select **At End of Connection** if you want to dashboards and events to reflect blocked connections. You can also select an external syslog server if you are using one.

Step 9 Click **OK**.

You can now deploy the policy. Once deployed, any SSL 3.0 or TLS 1.0 connection that goes through the system will be dropped.

Note
SSL 2.0 connections are handled by the default action for the policy. If you want to ensure these are also dropped, change the default action to Block.

What to do next

If you implement this rule, we have the following recommendations:

- For any type of decrypt rule, leave the default settings on the Advanced tab, where all SSL/TLS options are selected. By applying to all versions, the handshake process is simplified. However, your initial block rule will still prevent SSL 3.0 and TLS 1.0 connections.
- We normally recommend that you use Do Not Decrypt as the default action for the policy. However, because SSL 2.0 connections are always handled by the default action, you might want to use Block instead. However, if you want to apply Do Not Decrypt as the default action for all decryptable traffic, create a Do Not Decrypt rule at the end of the policy where you accept all default values for traffic matching criteria. This rule would match any supported TLS connection that does not match an earlier rule in the table, and act as the default for those TLS versions.

Monitoring and Troubleshooting SSL Decryption

The following topics explain how to monitor and troubleshoot SSL decryption policies.

Monitoring SSL Decryption

You can view information about decryption in the dashboards and events for traffic that matches rules (or the default action) for which you enabled logging.

SSL Decryption Dashboard

To evaluate overall decryption statistics, view the **Monitoring** > **SSL Decryption** dashboard. The dashboard shows the following information:

- Percentage of encrypted versus plain text traffic.
- How much encrypted traffic is decrypted per SSL rules.

Events

In addition to the dashboard, the event viewer (**Monitoring** > **Events**) includes SSL information for encrypted traffic. Following are some tips in evaluating events:

- For connections that were dropped because they matched an SSL rule (or default action) that blocked matching traffic, the **Action** should be “Block,” and the **Reason** should indicate “SSL Block.”
- The **SSL Actual Action** field indicates the actual action that the system applied to the connection. This can differ from the **SSL Expected Action**, which indicates the action defined on the matching rule. For example, a connection might match a rule that applies decryption, but could not be decrypted for some reason.

Handling Web Sites Where Decrypt Re-sign Works for a Browser but not an App (SSL or Certificate Authority Pinning)

Some apps for smart phones and other devices use a technique called SSL (or Certificate Authority) pinning. The SSL pinning technique embeds the hash of the original server certificate inside the app itself. As a result, when the app receives the resigned certificate from the Firewall Threat Defense device, the hash validation fails and the connection is aborted.

The primary symptom is that users cannot connect to the web site using the site's app, but they can connect using the web browser, even when using the browser on the same device where the app fails. For example, users cannot use the Facebook iOS or Android app, but they can point Safari or Chrome at https://www.facebook.com and make a successful connection.

Because SSL pinning is specifically used to avoid man-in-the-middle attacks, there is no workaround. You must choose between the following options:

- Support app users, in which case you cannot decrypt any traffic to the site. Create a Do Not Decrypt rule for the site's application (on the Application tab for the SSL Decryption rule) and ensure that the rule comes before any Decrypt Re-sign rule that would apply to the connections.
- Force users to use browsers only. If you must decrypt traffic to the site, you will need to inform users that they cannot use the site's app when connecting through your network, that they must use their browsers only.

More Details

If a site works in a browser but not in an app on the same device, you are almost certainly looking at an instance of SSL pinning. However, if you want to delve deeper, you can use connection events to identify SSL pinning in addition to the browser test.

There are two ways an app might deal with hash validation failures:

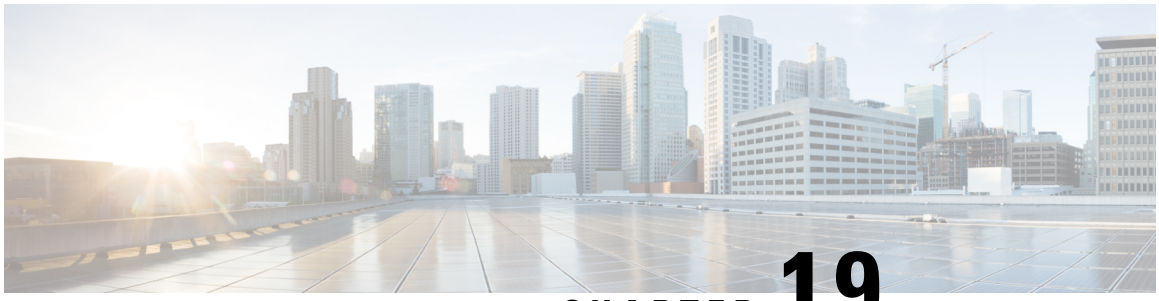
- Group 1 apps, such as Facebook, send an SSL ALERT Message as soon as it receives the SH, CERT, SHD message from the server. The Alert is usually an "Unknown CA (48)" alert indicating SSL Pinning. A TCP Reset is sent following the Alert message. You should see the following symptoms in the event details:
  - SSL Flow Flags include ALERT_SEEN.
  - SSL Flow Flags do not include APP_DATA_C2S or APP_DATA_S2C.
  - SSL Flow Messages typically are: CLIENT_HELLO, SERVER_HELLO, SERVER_CERTIFICATE, SERVER_KEY_EXCHANGE, SERVER_HELLO_DONE.



Note If you enabled SSL hardware acceleration, you would also see the CLIENT_ALERT message.

- Group 2 apps, such as Dropbox, do not send any alerts. Instead they wait until the handshake is done and then send a TCP Reset. You should see the following symptoms in the event:
  - SSL Flow Flags do not include ALERT_SEEN, APP_DATA_C2S, or APP_DATA_S2C.

- SSL Flow Messages typically are: 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.



CHAPTER 19

Identity Policies

You can use identity policies to collect user identity information from connections. You can then view usage based on user identity in the dashboards, and configure access control based on user or user group.

- Identity Policy Overview, on page 453
- How to Implement the Identity Policy, on page 455
- Best Practice for Active Authentication, on page 456
- Configuring Identity Policies, on page 457
- Enabling Transparent User Authentication, on page 463
- Monitoring Identity Policies, on page 466
- Examples for Identity Policies, on page 466

Identity Policy Overview

You can use identity policies to detect the user who is associated with a connection. By identifying the user, you can correlate threat, endpoint, and network intelligence with user identity information. By linking network behavior, traffic, and events directly to individual users, the system can help you identify the source of policy breaches, attacks, or network vulnerabilities.

For example, you can identify who owns the host targeted by an intrusion event, and who initiated an internal attack or port scan. You can also identify high bandwidth users and users who are accessing undesirable web sites or applications.

User detection goes beyond collecting data for analysis. You can also write access rules based on user name or user group name, selectively allowing or blocking access to resources based on user identity.

You can obtain user identity using the following methods:

- Passive authentication—For all types of connections, obtain user identity from other authentication services without prompting for username and password.
- Active authentication—For HTTP connections only, prompt for username and password and authenticate against the specified identity source to obtain the user identity for the source IP address.

The following topics provide more information on user identity.

Establishing User Identity Through Passive Authentication

Passive authentication gathers user identity without prompting the user for username and password. The system obtains the mappings from the identity sources you specify.

You can passively obtain user-to-IP address mappings from the following sources:

- Remote access VPN logins. The following user types are supported for passive identity:
  - User accounts defined in an external authentication server.
  - Local user accounts that are defined in the Firewall Device Manager.
- Cisco Identity Services Engine (ISE); Cisco Identity Services Engine Passive Identity Connector (ISE PIC).

If a given user is identified through more than one source, the RA VPN identity takes precedence.

Establishing User Identity through Active Authentication

Authentication is the act of confirming the identity of a user.

With active authentication, when an HTTP traffic flow comes from an IP address for which the system has no user-identity mapping, you can decide whether to authenticate the user who initiated the traffic flow against the directory configured for the system. If the user successfully authenticates, the IP address is considered to have the identity of the authenticated user.

Failure to authenticate does not prevent network access for the user. Your access rules ultimately decide what access to provide these users.

Dealing with Unknown Users

When you configure the directory server for the identity policy, the system downloads user and group membership information from the directory server. This information is refreshed every 24 hours at midnight, or whenever you edit and save the directory configuration (even if you do not make any changes).

If a user succeeds in authenticating when prompted by an active authentication identity rule, but the user's name is not in the downloaded user identity information, the user is marked as Unknown. You will not see the user's ID in identity-related dashboards, nor will the user match group rules.

However, any access control rules for the Unknown user will apply. For example, if you block connections for Unknown users, these users are blocked even though they succeeded in authenticating (meaning that the directory server recognizes the user and the password is valid).

Thus, when you make changes to the directory server, such as adding or deleting users, or changing group membership, these changes are not reflected in policy enforcement until the system downloads the updates from the directory.

If you do not want to wait until the daily midnight update, you can force an update by editing the directory realm information (from **Objects** > **Identity Sources**, then edit the realm). Click **Save**, then deploy changes. The system will immediately download the updates.



Note You can check whether new or deleted user information is on the system by going to **Policies** > **Access Control**, clicking the **Add Rule (+)** button, and looking at the list of users on the **Users** tab. If you cannot find a new user, or you can find a deleted user, then the system has old information.

How to Implement the Identity Policy

To enable user identity acquisition, so that the user associated with an IP address is known, you need to configure several items. When configured correctly, you will be able to see usernames in the monitoring dashboards and events. You will also be able to use user identity in access control and SSL decryption rules as a traffic match criteria.

The following procedure provides an overview of what you must configure to get identity policies to work.

Procedure

Step 1 Configure the AD identity realm.

Whether you collect user identity actively (by prompting for user authentication) or passively, you need to configure the Active Directory (AD) server that has the user identity information. See Configuring AD Identity Realms, on page 161.

If you configure passive identity, you can create AD realm sequences that allow the system to pull from identities in more than one AD realm. This is useful if you have multiple AD domains in your network.

Step 2 If you want to use passive authentication identity rules, configure the passive identity sources.

You can configure any of the following, based on the services you are implementing in the device and the services available to you in your network.

- Remote access VPN—If you intend to support remote access VPN connections to the device, user logins can provide the identity based on the AD server or on local users (those defined within the Firewall Device Manager). For information on configuring RA VPN, see Configuring Remote Access VPN, on page 678.
- Cisco Identity Services Engine (ISE) or Cisco Identity Services Engine Passive Identity Connector (ISE PIC)—If you use these products, you can configure the device as a pxGrid subscriber, and obtain user identity from ISE. See Configure Identity Services Engine, on page 170.

Step 3 Choose **Policies** > **Identity**, and enable the identity policy. See Configuring Identity Policies, on page 457.

Step 4 Configure Identity Policy Settings, on page 457.

The passive identity sources are automatically selected based on the sources you configured in the system. If you want to configure active authentication, you must configure the certificates for captive portal and SSL re-sign decryption (if you have not already enabled the SSL Decryption policy).

Step 5 Configure the Identity Policy Default Action, on page 459.

If your intention is to use passive authentication only, you can set the default action to passive authentication and there is no need to create specific rules.

Step 6 Configure Identity Rules, on page 459.

Create rules that will collect passive or active user identities from the relevant networks.

Best Practice for Active Authentication

When an identity rule requires active authentication for a user, the user is redirected to the captive portal port on the interface through which they are connected and then they are prompted to authenticate.

Because this redirection is to the interface IP address, the identity policy certificate does not match exactly, and users get an untrusted certificate error. The user must accept the certificate to continue and authenticate to the device. Because this behavior is similar to a man-in-the-middle attack, users are reluctant to accept the untrusted certificate.

To avoid this problem, you can configure active authentication to use the fully-qualified domain name (FQDN) of one interface on the device. With a properly configured certificate, users will not get an untrusted certificate error, and the authentication will be more seamless and appear to be more secure.

Before you begin

Active authentication occurs for HTTP traffic only, and is disruptive to the end user whenever the device does not have a current user mapping for the user's workstation or other client device. You can avoid disruption by instead implementing passive authentication.

Procedure

Step 1 In the DNS server, define a fully-qualified domain name (FQDN) for the interface IP address of the interface you want to use for gathering active authentication.

Also called the captive portal, this must be a routed interface.

Step 2 Using a Certificate Authority (CA), obtain a certificate for this FQDN.

You can create a certificate for the specific FQDN, such as. ftd1.captive-port.example.com. Optionally, you can:

- Obtain a wildcard certificate that can apply to the captive portal interfaces on many different devices, for example, *.captive-port.example.com. The wildcard could also be broader, and apply to a wide class of endpoints, such as *.eng.example.com, or even *.example.com.
- Include multiple Subject Alternate Names (SAN) in the certificate.

Step 3 Select **Objects** > **Certificates** and upload the certificate.

Step 4 Select **Objects** > **Network** and create an FQDN network object for the DNS name.

Step 5 On the **Policies** > **Identity** page, update the identity policy settings with the certificate and FQDN object.

Step 6 Create rules in the identity policy that use active authentication.

Configuring Identity Policies

You can use identity policies to collect user identity information from connections. You can then view usage based on user identity in the dashboards, and configure access control based on user or user group.

The following is an overview of how to configure the elements required to obtain user identity through identity policies.

Procedure

Step 1 Select **Policies** > **Identity**.

If you have not yet defined an identity policy, click **Enable Identity Policy** and configure settings as described in Configure Identity Policy Settings, on page 457.

Step 2 Manage the identity policy.

After you configure identity settings, this page lists all rules in order. Rules are matched against traffic from top to bottom with the first match determining the action to apply. You can do the following from this page:

- To enable or disable the identity policy, click the **Identity Policy** toggle.
- To change the identity policy settings, click the **Identity Policy Configuration** button ( ).
- To change the **Default Action**, click the action and select the desired action. See Configure the Identity Policy Default Action, on page 459.
- To move a rule, edit it and select the new location from the **Order** drop-down list.
- To configure rules:
  - To create a new rule, click the + button.
  - To edit an existing rule, click the edit icon ( ) for the rule (in the Actions column). You can also selectively edit a rule property by clicking on the property in the table.
  - To delete a rule you no longer need, click the delete icon ( ) for the rule (in the Actions column).

For more information on creating and editing identity rules, see Configure Identity Rules, on page 459.

Configure Identity Policy Settings

For identity policies to work, you must configure the sources that provide user identity information. The settings you must configure differ based on the type of rules you will configure: passive, active, or both.

The settings dialog box shows these settings in separate sections. Depending on how you access the dialog box, you will see both sections, or just one section. The dialog box appears automatically if you try to create a rule for an authentication type without having already configured the required settings.

The following procedure covers the full dialog box.

Before you begin

Ensure that time settings are consistent among the directory servers, Firewall Threat Defense device, and clients. A time shift among these devices can prevent successful user authentication. "Consistent" means that you can use different time zones, but the time should be the same relative to those zones; for example, 10 AM PST = 1 PM EST.

Procedure

Step 1 Select **Policies** > **Identity**.

Step 2 Click the **Identity Policy Configuration** button ().

Step 3 Configure the **Passive Authentication** options.

The dialog box shows you the passive authentication sources that you have already configured.

If necessary, you can configure ISE through this dialog box. If you have not configured an ISE object yet, you can click the **Integrate ISE** link and create it now. If the object exists, it is listed along with its state: Enabled or Disabled.

You must have configured at least one enabled passive identity source to create passive authentication rules.

Step 4 Configure the **Active Authentication** options.

When an identity rule requires active authentication for a user, the user is redirected to the captive portal port and then they are prompted to authenticate. Before configuring these settings, please read Best Practice for Active Authentication, on page 456.

- **Server Certificate**—Select the internal certificate to present to users during active authentication. If you have not already created the required certificate, click **Create New Internal Certificate** from the bottom of the drop-down list .

  Users will have to accept the certificate if you do not upload a certificate that their browsers already trust.

- **Redirect to Host Name** (Snort 3.0 only.)—Select the network object that defines the fully-qualified host name of the interface that should be used as the captive portal for active authentication requests. Click **Create New Network** if the object does not exist.

  The FQDN must resolve to the IP address of one of the interfaces on the device. By using an FQDN, you can assign a certificate for active authentication that the client will recognize, thus avoiding the untrusted certificate warning users get when being redirected to an IP address. The certificate can specify the FQDN, a wildcard FQDN, or multiple FQDNs in the Subject Alternate Names (SAN) in the certificate.

  If an identity rule requires active authentication for a user, but you do not specify a redirect FQDN, the user is redirected to the captive portal port on the interface through which they are connected.

- **Port**—The captive portal port. The default is 885 (TCP). If you configure a different port, it must be in the range 1025-65535.

  Note
  If you do not supply a **Redirect to Host Name** FQDN, the HTTP Basic, HTTP Response Page, and NTLM authentication methods redirect the user to the captive portal using the IP address of the interface. However, for HTTP Negotiate, the user is redirected using the fully-qualified DNS name

firewall-hostname.AD-domain-name. If you want to use HTTP Negotiate without a **Redirect to Host Name** FQDN, you must also update your DNS server to map this name to the IP addresses of all inside interfaces where you are requiring active authentication. Otherwise, the redirection cannot complete, and users cannot authenticate. We recommend that you always provide a **Redirect to Host Name** FQDN to ensure consistent behavior regardless of authentication method.

Step 5 (Active authentication only.) In **Decrypt Re-Sign Certificate**, select the internal CA certificate to use for rules that implement decryption with re-signed certificates.

You can use the pre-defined NGFW-Default-InternalCA certificate, or one that you created or uploaded. If the certificate does not yet exist, click **Create Internal CA** to create it.

If you have not already installed the certificate in client browsers, click the download button ( ) to obtain a copy. See the documentation for each browser for information on how to install the certificate. Also see Downloading the CA Certificate for Decrypt Re-Sign Rules, on page 446.

Note
You are prompted for SSL Decryption settings only if you have not already configured the SSL decryption policy. To change these settings after enabling the identity policy, edit the SSL decryption policy settings.

Step 6 Click **Save**.

Configure the Identity Policy Default Action

The identity policy has a default action, which is implemented for any connections that match no individual identity rules.

In fact, having no rules is a valid configuration for your policy. If you intend to use passive authentication on all traffic sources, then simply configure Passive Authentication as your default action.

Procedure

Step 1 Select **Policies** > **Identity**.

Step 2 Click in the **Default Action** and choose one of the following:

- **Passive Auth (Any Identity Source)**—User identity will be determined using all configured passive identity sources for connections that do not match any identity rules. If you do not configure any passive identity sources, using Passive Auth as the default is the same as using No Auth.
- **No Auth (No Authentication Required)**—User identity will not be determined for connections that do not match any identity rules.

Configure Identity Rules

Identity rules determine whether user identity information should be collected for matching traffic. You can configure No Authentication if you do not want to get user identity information for matching traffic.

Keep in mind that regardless of your rule configuration, active authentication is performed on HTTP traffic only. Thus, you do not need to create rules to exclude non-HTTP traffic from active authentication. You can simply apply an active authentication rule to all sources and destinations if you want to get user identity information for all HTTP traffic.

Note Also keep in mind that a failure to authentication has no impact on network access. Identity policies collect user identity information only. You must use access rules if you want to prevent users who failed to authenticate from accessing the network.

Before you begin

Rules are evaluated top-down. For a connection that matches the specified network criteria of a given rule, the user is evaluated against the identity realm specified in the rule. If the user is not part of that realm, they will be marked as unknown, and no further rules in the identity policy will be evaluated. Thus, if you have more than one realm that needs to be evaluated, be sure to use realm sequences instead of a single realm.

Procedure

Step 1 Select **Policies** > **Identity**.

Step 2 Do any of the following:

- To create a new rule, click the + button.
- To edit an existing rule, click the edit icon ( ) for the rule.

To delete a rule you no longer need, click the delete icon ( ) for the rule.

Step 3 In **Order**, select where you want to insert the rule in the ordered list of rules.

Rules are applied on a first-match basis, so you must ensure that rules with highly specific traffic matching criteria appear above policies that have more general criteria that would otherwise apply to the matching traffic.

The default is to add the rule to the end of the list. If you want to change a rule's location later, edit this option.

Step 4 In **Title**, enter a name for the rule.

Step 5 Select the **Action** and if necessary, **AD Identity Source**.

You must select the AD identity realm that includes the user accounts for passive and active authentication rules. If the realm you need does not yet exist, click **Create New Identity Realm** and create it now. For passive authentication, you can select an AD realm sequence rather than a single AD realm object.

- **Passive Auth**—Use passive authentication to determine user identity. All configured identity sources are shown. The rule automatically uses all configured sources.
- **Active Auth**—Use active authentication to determine user identity. Active authentication is applied to HTTP or decrypted HTTPS traffic only. If any other type of traffic matches an identity policy that requires or allows active authentication, then active authentication will not be attempted.
- **No Auth**—Do not obtain user identity. Identity-based access rules will not be applied to this traffic. These users are marked as **No Authentication Required**.

Step 6 (Active Authentication only.) Select the authentication method (**Type**) supported by your directory server.

- **HTTP Basic**—Authenticate users using an unencrypted HTTP Basic Authentication (BA) connection. Users log in to the network using their browser's default authentication popup window. This is the default.
- **NTLM**—Authenticate users using an NT LAN Manager (NTLM) connection. This selection is only available when you select an AD realm. Users log in to the network using their browser's default authentication popup window, although you can configure IE and Firefox browsers to transparently authenticate using their Windows domain login (see Enabling Transparent User Authentication, on page 463).
- **HTTP Negotiate**—Allow the device to negotiate the method between the user agent (the application the user is using to initiate the traffic flow) and the Active Directory server. Negotiation results in the strongest commonly supported method being used, in order, NTLM, then basic. Users log in to the network using their browser's default authentication popup window.
- **HTTP Response Page**—Prompt users to authenticate using a system-provided web page. This is a form of HTTP Basic authentication.

Note
If you do not supply a **Redirect to Host Name** FQDN, the HTTP Basic, HTTP Response Page, and NTLM authentication methods redirect the user to the captive portal using the IP address of the interface. However, for HTTP Negotiate, the user is redirected using the fully-qualified DNS name *firewall-hostname.AD-domain-name*. If you want to use HTTP Negotiate without a **Redirect to Host Name** FQDN, you must also update your DNS server to map this name to the IP addresses of all inside interfaces where you are requiring active authentication. Otherwise, the redirection cannot complete, and users cannot authenticate. We recommend that you always provide a **Redirect to Host Name** FQDN to ensure consistent behavior regardless of authentication method.

Step 7 (Active authentication only.) Select **Fall Back as Guest** > **On/Off** to determine whether users who fail active authentication are labeled as Guest users.

Users get 3 chances to successfully authenticate. If they fail, your selection for this option determines how the user is marked. You can write access rules based on these values.

- **Fall Back as Guest** > **On**—Users are marked as **Guest**.
- **Fall Back as Guest** > **Off**—Users are marked as **Failed Authentication**.

Step 8 Define the traffic matching criteria on the **Source/Destination** tab.

Keep in mind that active authentication will be attempted with HTTP traffic only. Therefore, there is no need to configure No Auth rules for non-HTTP traffic, and there is no point in creating Active Authentication rules for any non-HTTP traffic. However, passive authentication is valid for any type of traffic.

The Source/Destination criteria of an identity rule define the security zones (interfaces) through which the traffic passes, the IP addresses or the country or continent (geographical location) for the IP address, or the protocols and ports used in the traffic. The default is any zone, address, geographical location, protocol, and port.

To modify a condition, you click the **+** button within that condition, select the desired object or element, and click **OK** in the popup dialog box. If the criterion requires an object, you can click **Create New *Object*** if the object you require does not exist. Click the **x** for an object or element to remove it from the policy.

You can configure the following traffic matching criteria.

Source Zones, Destination Zones

The security zone objects that define the interfaces through which the traffic passes. You can define one, both, or neither criteria: any criteria not specified applies to traffic on any interface.

- To match traffic leaving the device from an interface in the zone, add that zone to the **Destination Zones**.
- To match 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.

Use this criteria when the rule should apply based on where the traffic enters or exits the device. For example, if you want to ensure that user identity is collected from all traffic originating from inside networks, select an inside zone as the **Source Zones** while leaving the destination zone empty.

Note
You cannot mix passive and routed security zones in a single rule. In addition, you can specify passive security zones as source zones only, you cannot specify them as destination zones.

Source Networks, Destination Networks

The network objects or geographical locations that define the network addresses or locations of the traffic.

- To match traffic from an IP address or geographical location, configure the **Source Networks**.
- To match 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 traffic must originate from one of the specified IP addresses and be destined for one of the destination IP addresses.

When you add this criteria, you select from the following tabs:

- **Network**—Select the network objects or groups that define the source or destination IP addresses for the traffic you want to control.
- **Geolocation**—Select the geographical location to control traffic based on its source or destination country or continent. Selecting a continent selects all countries within the continent. Besides selecting geographical location directly in the rule, you can also select a geolocation object that you created to define the location. Using geographical location, you could easily restrict access to a particular country without needing to know all of the potential IP addresses used there.

  Note
  To ensure you are using up-to-date geographical location data to filter your traffic, Cisco strongly recommends that you regularly update the geolocation database (GeoDB).

Source Ports, Destination Ports/Protocols

The port objects that define the protocols used in the traffic. For TCP/UDP, this can include ports.

- To match traffic from a protocol or port, configure the **Source Ports**. Source ports can be TCP/UDP only.
- To match traffic to a protocol or port, configure the **Destination Ports/Protocols**.
- To match traffic both originating from specific TCP/UDP ports and destined for specific TCP/UDP ports, configure both. If you add both source and destination ports to a condition, you can only add ports that

share a single transport protocol, TCP or UDP. For example, you could target traffic from port TCP/80 to port TCP/8080.

Step 9 Click **OK**.

Enabling Transparent User Authentication

If you configure the identity policy to allow for active authentication, you can use the following authentication methods to acquire user identity:

HTTP Basic

With HTTP basic authentication, users are always prompted to authenticate with their directory username and password. The password is transmitted in clear text. For that reason, basic authentication is not considered a secure form of authentication.

Basic is the default authentication mechanism.

HTTP Response Page

This is a type of HTTP basic authentication, where the user is presented with a login browser page.

NTLM, HTTP Negotiate (Integrated Windows Authentication for Active Directory)

With integrated Windows authentication, you take advantage of the fact that users log into a domain to use their workstation. The browser tries to use this domain login when accessing a server, including the Firewall Threat Defense captive portal during active authentication. The password is not transmitted. If authentication is successful, the user is transparently authenticated; the user is unaware that any authentication challenge was made or satisfied.

If the browser cannot satisfy an authentication request using the domain login credentials, the user is prompted for username and password, which is the same user experience as basic authentication. Thus, if you configure integrated Windows authentication, it can reduce the need for users to supply credentials when accessing the network or servers in the same domain.

Note that HTTP Negotiate picks the strongest method supported by both the Active directory server and the user agent. If negotiation selects HTTP Basic as the authentication method, you will not get transparent authentication. The order of strength is NTLM, then basic. Negotiation must select NTLM for transparent authentication to be possible.

You must configure client browsers to support integrated Windows authentication to enable transparent authentication. The following sections explain the general requirements and basic configuration of integrated Windows authentication for some commonly used browsers that support it. Users should consult the help for their browser (or other user agent) for more detailed information, because the techniques can change between software releases.

Tip Not all browsers support integrated Windows authentication, such as Chrome and Safari (based on the versions available when this was written). Users will be prompted for username and password. Consult the browser's documentation to determine if support is available in the version you use.

Requirements for Transparent Authentication

Users must configure their browser or user agent to implement transparent authentication. They can do this individually, or you can configure it for them and push the configuration to client workstations using your software distribution tools. If you decide to have users do it themselves, ensure that you provide the specific configuration parameters that work for your network.

Regardless of browser or user agent, you must implement the following general configuration:

- Add the Firewall Threat Defense redirection host name, or the interface through which users connect to the network, to the Trusted Sites list. If you do not use a redirection host name, you can use the IP address or if available, the fully-qualified domain name (for example, inside.example.com). You can also use wildcards or partial addresses to create a generalized trusted site. For example, you can typically cover all internal sites using *.example.com or simply example.com, trusting all servers in your network (use your own domain name). If you add the specific address of the interface, you might need to add several addresses to the trusted sites to account for all user access points to the network.
- Integrated Windows authentication does not work through a proxy server. Therefore, you must either not use a proxy, or you must add the Firewall Threat Defense redirection host name or interface to the addresses excluded from going through the proxy. If you decide that you must use a proxy, users will be prompted for authentication even if you use NTLM.

Tip Configuring transparent authentication is not a requirement, but a convenience to end users. If you do not configure transparent authentication, users are presented with a login challenge for all authentication methods.

Configuring Internet Explorer for Transparent Authentication

To configure Internet Explorer for NTLM transparent authentication:

Procedure

Step 1 Select **Tools** > **Internet Options**.

Step 2 Select the **Security** tab, select the **Local Intranet** zone, then do the following:

a) Click the **Sites** button to open the list of trusted sites.

b) Ensure that at least one of the following options is selected:

- **Automatically detect intranet network**. If you select this option, all other options are disabled.
- **Include all sites that bypass the proxy**.

c) Click **Advanced** to open the Local Intranet Sites dialog box, then paste the URL you want to trust into the **Add Site** box and click **Add**.

Repeat the process if you have more than one URL. Use wildcards to specify a partial URL, such as `http://*.example.com` or simply `*.example.com`.

Close the dialog boxes to return to the Internet Options dialog box.

d) With **Local Intranet** still selected, click **Custom Level** to open the Security Settings dialog box. Find the **User Authentication** > **Logon** setting and select **Automatic logon only in Intranet zone**. Click **OK**.

Step 3 In the Internet Options dialog box, click the **Connections** tab, then click **LAN Settings**.

If **Use a proxy server for your LAN** is selected, you need to ensure that the Firewall Threat Defense interface bypasses the proxy. Do any of the following as appropriate:

- Select **Bypass proxy server for local addresses**.
- Click **Advanced** and enter the address into the **Do not use proxy server for addresses beginning with** box. You can use wildcards, for example, **`*.example.com`**.

Configuring Firefox for Transparent Authentication

To configure Firefox for NTLM transparent authentication:

Procedure

Step 1 Open **about:config**. Use the filter bar to help you locate the preferences that you need to modify.

Step 2 To support NTLM, modify the following preferences (filter on network.automatic):

- **network.automatic-ntlm-auth.trusted-uris**—Double-click the preference, enter the URL, and click **OK**. You can enter multiple URLs by separating them with commas; including the protocol is optional. For example:

  ```
  http://host.example.com, http://hostname, myhost.example.com
  ```

  You can also use partial URLs. Firefox matches the end of the string, not a random substring. Thus, you could include your entire internal network by specifying just your domain name. For example:

  ```
  example.com
  ```

- **network.automatic-ntlm-auth.allow-proxies**—Ensure that the value is **true**, which is the default. Double-click to change the value if it is currently false.

Step 3 Check the HTTP proxy settings. You can find these by selecting **Tools** > **Options**, then click the **Network** tab in the Options dialog box. Click the **Settings** button in the Connection group.

- If **No Proxy** is selected, there is nothing to configure.
- If **Use System Proxy Settings** is selected, you need to modify the **network.proxy.no_proxies_on** property in about:config to add the trusted URIs you included in **network.automatic-ntlm-auth.trusted-uris**.
- If **Manual Proxy Configuration** is selected, update the **No Proxy For** list to include these trusted URIs.
- If one of the other options is selected, ensure that the properties used for those configurations exclude the same trusted URIs.

Monitoring Identity Policies

If identity policies that require authentication are working correctly, you should see user information on the **Monitoring** > **Users** dashboard and other dashboards that include user information.

In addition, events shown in **Monitoring** > **Events** should include user information.

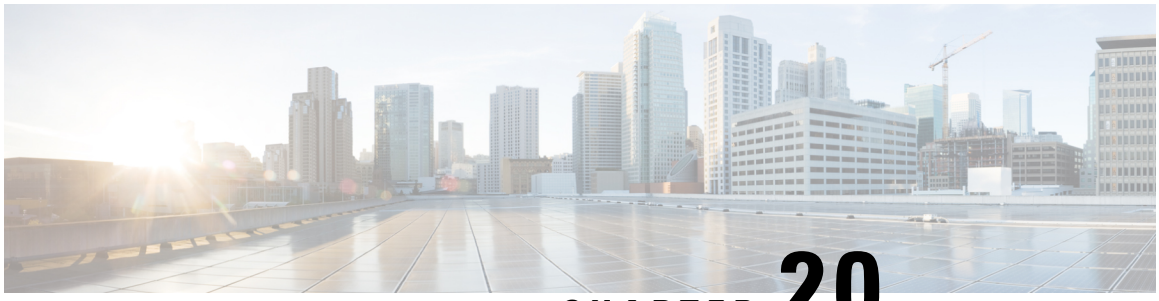
If you do not see any user information, verify that the directory server is functioning correctly. Use the **Test** button in the directory server configuration dialog box to verify connectivity.

If the directory server is functioning and usable, verify that the traffic matching criteria on the identity rules that require active authentication are written in a way that will match your users. For example, ensure that the source zone contains the interfaces through which your user traffic will enter the device. The active authentication identity rules match HTTP traffic only, so users must be sending that type of traffic through the device.

For passive authentication, use the **Test** button in the ISE object if you are using that source. If you are using remote access VPN, verify that the service is functioning correctly and that users can make VPN connections. See the troubleshooting topics for these features for more detailed information on identifying and resolving issues.

Examples for Identity Policies

The use case chapter includes an example of implementing identity policies. Please see How to Gain Insight Into Your Network Traffic, on page 44.



CHAPTER 20

Security Intelligence

The Security Intelligence policy gives you an early opportunity to drop unwanted traffic based on source/destination IP address or destination URL. The following topics explain how to implement Security Intelligence.

- About Security Intelligence, on page 467
- License Requirements for Security Intelligence, on page 469
- Configuring Security Intelligence, on page 469
- Monitoring Security Intelligence, on page 470
- Examples for Security Intelligence, on page 471

About Security Intelligence

The Security Intelligence policy gives you an early opportunity to drop unwanted traffic based on source/destination IP address or destination URL. The system drops this unwanted traffic before evaluating it with the access control policy, thus reducing the amount of system resources used.

You can block traffic based on the following:

- Cisco Talos Intelligence Group (Talos) feeds—Talos provides access to regularly updated security 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. The system downloads feed updates regularly, and thus new threat intelligence is available without requiring you to redeploy the configuration.

Note Talos feeds are updated by default every hour. You can change the update frequency, and even update the feeds on demand, from the **Device** > **Updates** page.

- Network and URL objects—If you know of specific IP addresses or URLs you want to block, you can create objects for them and add them to the blocked list or the exception list. Note that you cannot use network objects with FQDN or range specifications.

You create separate lists for IP addresses (networks) and URLs.

Note If an HTTP/HTTPS request is to a URL that uses an IP address instead of a hostname, the system looks up the IP address reputation in the network address lists. You do not need to duplicate IP addresses in the network and URL lists.

Making Exceptions to the Block Lists

For each block list, you can create an associated exception list, also known as the do not block list. The only purpose of the exception list is to exempt IP addresses or URLs that appear in the block list. That is, if you find an address or URL you need to use, and you know to be safe, is in a feed configured on the block list, you can exempt that network/URL without completely removing the category from the block list.

Exempted traffic is subsequently evaluated by the access control policy. The ultimate decision on whether the connections are allowed or dropped is based on the access control rule the connections match. The access rule also determines whether intrusion or malware inspection is applied to the connection.

Security Intelligence Feed Categories

The following table describes the categories available in the Cisco Talos Intelligence Group (Talos) feeds. These categories are available for both network and URL blocking.

These categories can change over time, so a newly-downloaded feed might have category changes. When configuring Security Intelligence, you can click the info icon next to a category name to see a description.

Table 11: 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
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)

Security Intelligence Category	Description
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

License Requirements for Security Intelligence

You must enable the **Threat** license to use Security Intelligence. See Enabling or Disabling Optional Licenses, on page 89.

Configuring Security Intelligence

The Security Intelligence policy gives you an early opportunity to drop unwanted traffic based on source/destination IP address or destination URL. Any allowed connections are still evaluated by access control policies and might eventually be dropped.

Procedure

Step 1 Select **Policies** > **Security Intelligence**.

Step 2 If the policy is not enabled, click the **Enable Security Intelligence** button.

You can disable the policy at any time by clicking the **Security Intelligence** toggle to **Off**. Your configuration is preserved, so that when you enable the policy again you do not need to reconfigure it.

Step 3 Configure Security Intelligence.

There are separate block lists for Networks (IP addresses) and URLs.

a) Click the **Network** or **URL** tab to display the list you want to configure.

b) In the block/drop list, click **+** to select the objects or feeds whose connections you want to drop immediately.

The object selector organizes the objects and feeds on separate tabs by type. If the object you want does not yet exist, click the **Create New Object** link at the bottom of the list and create it now. For a description of the Cisco Talos Intelligence Group (Talos) feeds, click the **i** button next to the feed. See also Security Intelligence Feed Categories, on page 468.

Note
Security Intelligence ignores IP address blocks using a /0 netmask. This includes the any-ipv4 and any-ipv6 network objects. Do not select these objects for network blocking.

c) In the do not block list, click **+** and select any exceptions to the block list.

The only reason to configure this list is to make exceptions for IP addresses or URLs that are in the block list. Exempted connections are subsequently evaluated by your access control policy, and might be dropped anyway.

d) Repeat the process to configure the other block list.

Step 4 (Optional.) Click the **Edit Logging Settings** button (⚙) to configure logging.

If you enable logging, any matches to block list entries are logged. Matches to exception entries are not logged, although you get log messages if exempted connections match access control rules with logging enabled.

Configure the following settings:

- **Connection Events Logging**—Click the toggle to enable or disable logging.
- **Syslog**—If you want to send a copy of the events to an external syslog server, select this option and select the server object that defines the syslog server. If the required object does not already exist, click **Add Syslog Server** and create it.

  Because event storage on the device is limited, sending events to an external syslog server can provide more long term storage and enhance your event analysis.

Monitoring Security Intelligence

If you enable logging for the Security Intelligence policy, the system generates Security Intelligence events for each connection that matches an item on a block list. There are matching connection events for these connections.

Statistics for dropped connections appear in the various dashboards available on the Monitoring page.

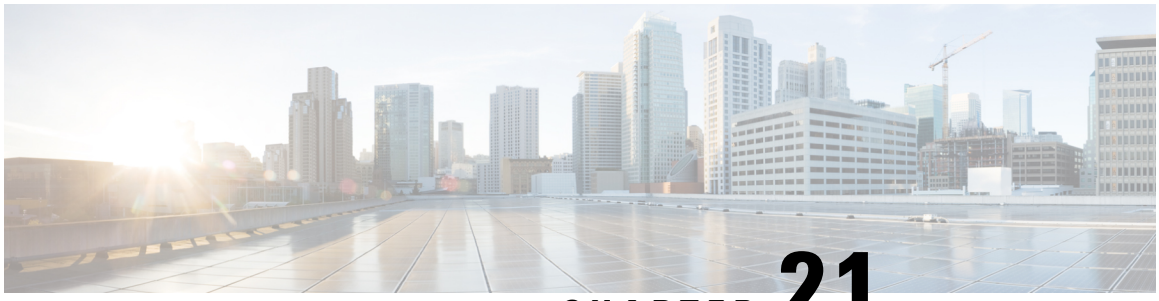
The **Monitoring** > **Access and SI Rules** dashboard shows the top access rules and Security Intelligence rule-equivalents that are matching traffic.

In addition, you can select **Monitoring** > **Events**, then the **Security Intelligence** view, to see the Security Intelligence events, as well as the related connection events on the **Connection** tab.

- The SI Category ID field in an event indicates the object matched in the block list, such as a network or URL object or feed.
- The Reason field in a connection event explains why the action shown in the event was applied. For example, a Block action paired with reasons such as IP Block or URL Block indicates that a connection was dropped by Security Intelligence.

Examples for Security Intelligence

The use case chapter includes an example of implementing Security Intelligence policies. Please see How to Block Threats, on page 51.



CHAPTER 21

Access Control

The following topics explain access control rules. These rules control which traffic is allowed to pass through the device, and apply advanced services to the traffic, such as intrusion inspection.

- Best Practices for Access Control, on page 473
- Access Control Overview, on page 476
- License Requirements for Access Control, on page 486
- Guidelines and Limitations for Access Control Policies, on page 487
- Configuring the Access Control Policy, on page 489
- Monitoring Access Control Policies, on page 500
- Examples for Access Control, on page 502

Best Practices for Access Control

The access control policy is your primary tool for protecting your internal networks and preventing your users from accessing undesirable external network resources, such as inappropriate web sites. Thus, we recommend that you pay special attention to this policy and fine-tune it to get the level of protection and connectivity that you need.

The following procedure provides an overview of the basic things you should do with the access control policy. This is an overview, and does not provide exhaustive steps for performing each task.

To get to the access control policy, choose **Policies** > **Access Control**.

Procedure

Step 1 Configure the default action for the policy.

The default action handles connections that do not match the specific rules in the policy. By default, this action is **Block**, so that anything you miss in the rules is blocked. Thus, you simply need to write access control rules that allow desirable traffic. This is the traditional way to configure the access control policy.

You can do the opposite, where you allow traffic by default, and write rules that drop known undesirable traffic, so that you do not need to have rules for everything you want to allow. This makes it easy for new services to be used, but opens the risk that new undesirable traffic will get through before you notice it.

Step 2 Click the **Access Policy Settings** ( ) button, and enable the **TLS Server Identity Discovery** option.

This option improves the initial application detection and URL category and reputation identification for TLS 1.3 connections. If you do not enable this option, TLS 1.3 traffic might not match your intended rules. This option can also improve the efficacy of decryption rules.

Step 3 Create as few access control rules as possible.

With traditional firewalls, you might end up with tens of thousands of rules for various combinations of IP address and port. With a next-generation firewall, you can use advanced inspection and avoid some of these detailed rules. The fewer rules you have, the faster the system can evaluate traffic, and the easier it will be for you to find and fix problems within your rule set.

Step 4 Enabling logging on your access control rules.

Statistics are collected for matching traffic only if you enable logging. Your monitoring dashboards will be inaccurate if you do not enable logging.

Step 5 Put very specific rules at the top of the policy, and ensure that specific rules are above any more general rule that would match the connections a specific rule would also match.

The policy is evaluated top-down, first match wins. Thus, if you put in a rule that blocks all traffic to a specific subnet, then follow it with a rule that allows access to a single IP address within the subnet, traffic to that address will not be allowed, because the first rule will block it.

In addition, place rules that target traffic based only on traditional criteria such as ingress/egress interface, and source/destination IP address, port, or Geolocation, ahead of rules that require deep inspection, such as those that apply to user criteria, URL filtering, or application filtering. Because these rules do not require inspection, putting them early can get you quicker access control decisions for matching connections.

For more suggestions, see Best Practices for Access Control Rule Order, on page 485.

Step 6 Pair Block and Allow rules to target subsets of traffic.

For example, it is likely that you want to allow a lot of HTTP/HTTPS traffic, yet block access to some undesirable sites such as pornography or gambling. You can accomplish this by creating the following rules and keeping them sequential within the policy (for example, rules 11 and 12).

- A URL filtering Block rule that targets undesirable URL categories, applied to the inside security zone (source) and outside security zone (destination), and any IP address, port, or Geolocation. For example, Block Botnets, Child Abuse Content, Cryptojacking, DNS Tunneling, Ebanking Fraud, Exploits, Extreme, Filter Avoidance, Gambling, Hacking, Hate Speech, High Risk Sites and Locations, Illegal Activities, Illegal Downloads, Illegal Drugs, Malicious Sites, Malware Sites, Mobile Threats, P2P Malware Node, Phishing, Pornography, Spam, Spyware and Adware.
- An application filtering Allow rule for the HTTP and HTTPS applications, applied to the inside security zone (source) and outside security zone (destination), and any IP address, port, or Geolocation. After the URL filtering rule blocks access to unwanted web resources, this rule allows all other HTTP/HTTPS access.

Step 7 Use advanced next-generation firewall features to target traffic regardless of IP address or port.

Attackers or other bad actors can frequently change IP addresses and ports to evade traditional access control traffic-matching criteria. Instead, use the following next-generation features:

- User criteria—Configure the Identity policy to obtain information about the user who is initiating the traffic. Ideally, your Active Directory server organizes users into groups, and you can create access control rules that allow or block traffic based on user group membership. For example, allow Engineers to access your development subnets, but implicitly block anyone who is not in the Engineers group. Use

groups rather than individual user names, so you do not need to continually update the rules as people are added to the network.

- Application criteria—Use application filtering criteria to allow or block types of applications. This, if a user changes ports for an HTTP connection, the system can recognize that it is HTTP, even though it isn't going to port 80. For more suggestions, see Best Practices for Application Filtering, on page 477.
- URL category and reputation criteria—Use URL filtering based on category to dynamically allow or block sites based on the type of site. Within the site type, or category, you can fine-tune your rule based on whether the site has a reputation as a good actor or a bad actor. By using category and reputation, you will not need to constantly adjust your rules as sites change URLs, which you would have to do if you tried to manually block sites by URL. For more suggestions, see Best Practices for Effective URL Filtering, on page 482.

  You can also apply URL category/reputation filtering rules to the FQDN in a DNS lookup request. The system can prevent the DNS response for blocked category/reputation, effectively blocking the user's connection attempt. For details, see Filtering DNS Requests Based on URL Category and Reputation, on page 484.

Step 8 Apply intrusion inspection to all of your Allow rules.

One of the powerful aspects of next-generation firewalls is that you can apply intrusion inspection and access control using the same device. Apply an intrusion policy to each Allow rule, so that if an attack does enter your network through a normally benign path, you can catch it and drop the attacking connection.

If your default action is Allow, you can also apply intrusion protection for traffic that matches the default action.

Step 9 Also configure the Security Intelligence policy to block unwanted IP addresses and URLs.

The Security Intelligence policy is applied before the access control policy, so it can block unwanted connections before your access control rules are even evaluated. This can provide an early block and help you reduce the complexity of your access control rules.

Step 10 Consider implementing the SSL Decryption policy.

The system cannot perform deep inspection on encrypted traffic. If you configure the SSL Decryption policy, the access control policy is applied to a decrypted version of the traffic. Thus, deep inspection can identify attacks (using the intrusion policy), and rule matching is better because application and URL filtering can be applied more effectively. Any traffic that the access control policy allows is then re-encrypted before being sent out of the device, so the end user does not lose the protection of encryption.

Step 11 Enable object group search to simplify the deployment of your rules.

Enabling object group search reduces memory requirements for access control policies that include network objects. However, it is important to note that object group search might also decrease rule lookup performance and thus increase CPU utilization. You should balance the CPU impact against the reduced memory requirements for your specific access control policy. In most cases, enabling object group search provides a net operational improvement.

You can set this option using FlexConfig by issuing the **object-group-search access-control** command; use the **no** form of the command in the negate template.

Access Control Overview

The following topics explain access control policies.

Access Control Rules and the Default Action

Use the access control policy to allow or block access to network resources. The policy consists of a set of ordered rules, which are evaluated from top to bottom. The rule applied to traffic is the first one where all the traffic criteria are matched.

You can control access based on:

- Traditional network characteristics such as source and destination IP addresses, protocol, ports, and interfaces (in the form of security zones).
- The fully-qualified domain name (FQDN) of the source or destination (in the form of a network object). Traffic matching is based on the IP address returned for the name from a DNS lookup.
- The security group tag (SGT) assigned to the source or destination by the Cisco Identity Services Engine (ISE).
- The application that is being used. You can control access based on the specific application, or you can create rules that cover categories of applications, applications tagged with a particular characteristic, the type of application (client, server, web), or the application's risk or business relevance rating.
- The destination URL of a web request, including the generalized category of the URL. You can refine category matches based on the public reputation of the target site.
- The URL category and reputation of an FQDN in a DNS lookup request. You can block the DNS response for unwanted categories or poor reputations, effectively preventing the subsequent connection attempt.
- The user who is making the request, or the user groups to which the user belongs.

For unencrypted traffic that you allow, you can apply IPS inspection to check for threats and block traffic that appears to be an attack. You can also use file policies to check for prohibited files or malware.

Any traffic that does not match an access rule is handled by the access control **Default Action**. If you allow traffic by default, you can apply intrusion inspection to the traffic. However, you cannot perform file or malware inspection on traffic handled by the default action.

Application Filtering

You can use access control rules to filter traffic based on the application used in the connection. The system can recognize a wide variety of applications, so that you do not need to figure out how to block one web application without blocking all web applications.

For some popular applications, you can filter on different aspects of the application. For example, you could create a rule that blocks Facebook Games without blocking all of Facebook.

You can also create rules based on general application characteristics, blocking or allowing entire groups of applications by selecting risk or business relevance, type, category, or tag. **However, as you select categories in an application filter, look over the list of matching applications to ensure you are not including**

unintended applications. For a detailed explanation of the possible groupings, see Application Criteria, on page 494.

Application Control for Encrypted and Decrypted Traffic

If an application uses encryption, the system might not be able to identify the application.

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.

Use the application filters dialog box to determine if your application requires decryption by selecting the following Tags, then examining the list of applications.

- **SSL Protocol**—You do not need to decrypt traffic tagged as SSL Protocol. The system can recognize this traffic and apply your access control action. Access control rules for the listed applications should match to expected connections.
- **Decrypted Traffic**—The system can recognize this traffic only if you first decrypt the traffic. Configure SSL decryption rules for this traffic.

Filtering on Common Industrial Protocol (CIP) and Modbus Applications (ISA 3000)

You can enable the Common Industrial Protocol (CIP) and Modbus pre-processors on Cisco ISA 3000 devices, and filter on CIP and Modbus applications in access control rules. All CIP application names start with "CIP," such as CIP Write. There is only one application for Modbus.

To enable the pre-processors, you must go into expert mode in a CLI session (SSH or Console) and issue the following command to enable one or both of these Supervisory Control and Data Acquisition (SCADA) applications.

sudo /usr/local/sf/bin/enable_scada.sh {cip | modbus | both}

For example, to enable both pre-processors:

```
> expert
admin@firepower:~$ sudo /usr/local/sf/bin/enable_scada.sh both
```

Note You must issue this command after every deployment. These pre-processors are disabled during deployment.

Best Practices for Application Filtering

Please keep the following recommendations in mind when designing your application filtering access control rules.

- To handle traffic referred by a web server, such as advertisement traffic, match the referred application rather than the referring application.
- Avoid combining application and URL criteria in the same rule, especially for encrypted traffic.
- If you write a rule for traffic that is tagged **Decrypted Traffic**, ensure that you have an SSL Decryption rule that will decrypt the matching traffic. These applications can be identified in decrypted connections only.

- TLS 1.3 encrpyts most handshake messages, so certificate information is not readily available. For traffic encrypted with TLS 1.3 to match access rules that use application or URL filtering effectively, the system must obtain a clear-text certificate for the server. We recommend that you enable **TLS 1.3 Certificate Visibility** in the access control settings. If you enable this option, the system checks whether a certificate for the site is stored in cache based on the IP address and Server Name Indication (SNI) in the client hello packet. If one is not available, the system uses a TLS 1.2 probe to obtain the certicate, which can then be used for application/URL category and reputation identification without decrypting the connection.
- The system can detect multiple types of Skype application traffic. 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.
- To control access to Zoho mail, select both the Zoho and Zoho Mail applications.

URL Filtering

You can use access control rules to filter traffic based on the URL used in an HTTP or HTTPS connection. Note that URL filtering for HTTP is more straight-forward than it is for HTTPS, because HTTPS is encrypted.

You can use the following techniques to implement URL filtering.

- Category and reputation-based URL filtering—With a URL filtering license, you can control access to web sites based on the URL's general classification (category) and risk level (reputation). This is by far the easiest and most effective way to block unwanted sites.
- Manual URL filtering—With any license, you can manually specify individual URLs, and groups of URLs, to achieve granular, custom control over web traffic. The main purpose of manual filtering is to create exceptions to category-based block rules, but you can use manual rules for other purposes.

The following topics provide more information on URL filtering.

Filtering URLs by Category and Reputation

With a URL filtering license, you can control access to web sites based on the category and reputation of the 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 Untrusted (level 1) to Trusted (level 5). Reputation is determined by the hostname, e.g. all pages on www.cisco.com are considered to have the same reputation.

URL categories and reputations help you quickly configure URL filtering. For example, you can use access control to block untrusted URLs in the Illegal Drugs category.

For a description of the categories, see https://www.talosintelligence.com/categories.

Using category and reputation data also simplifies policy creation and administration. Sites that represent security threats, or that serve undesirable content, might appear and disappear faster than you can update and deploy new policies. As Cisco updates the URL database with new sites, changed classifications, and changed reputations, your rules automatically adjust to the new information. You do not need to edit your rules to account for new sites.

If you enable regular URL database updates, you can ensure that the system uses up-to-date information for URL filtering. You can also enable communications with Cisco Collective Security Intelligence (CSI) to obtain the latest threat intelligence for URLs with unknown category and reputation. For more information, see Configuring URL Filtering Preferences, on page 777.

Note To see URL category and reputation information in events and application details, you must create at least one rule with a URL condition.

Looking Up the Category and Reputation for a URL

You can check on the category and reputation for a particular URL. You can go to the URL tab of an access control rule, or SSL decryption rule, or go to **Device** > **System Settings** > **URL Filtering Preferences**. There, you can enter the URL in the **URL to Check** field and click **Go**.

You will be taken to a web site that shows the results of the lookup. You can use this information to help you check the behavior of your category- and reputation-based URL filtering rules.

If you disagree with the categorization, you can click the **Submit a URL Category Dispute** in the Firewall Device Manager to tell us what you think.

Manual URL Filtering

You can supplement or selectively override category and reputation-based URL filtering by manually filtering individual URLs or groups of URLs. You can perform this type of URL filtering without a special license.

For example, you might use access control to block a category of web sites that are not appropriate for your organization. However, if the category contains a web site 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.

To configure manual URL filtering, you create a URL object with the destination URL. How this URL is interpreted is based on the following rules:

- If you do not include a path (that is, there are no / characters in the URL), the match is based on the server's hostname only. If you include one or more / character, the entire URL string is used for a substring match. Then, a URL is considered a match if any of the following are true:
  - The string is at the beginning of the URL.
  - The string follows a dot.
  - The string contains a dot in the beginning.
  - The string follows the :// characters.

  For example, ign.com matches ign.com or www.ign.com, but not versign.com.

Note We recommend that you do not use manual URL filtering to block or allow individual web pages or parts of sites (that is, URL strings with / characters), as servers can be reorganized and pages moved to new paths.

- 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.

Filtering HTTPS Traffic

Because HTTPS traffic is encrypted, performing URL filtering directly on HTTPS traffic is not as straight-forward as it is on HTTP traffic. For that reason, you should consider using SSL Decryption policies to decrypt all HTTPS traffic that you intend to filter. That way, the URL filtering access control policies work on decrypted traffic, and you get the same results you would get for regular HTTP traffic.

However, if you do intend to allow some HTTPS traffic to pass undecrypted into the access control policy, you need to understand that rules match HTTPS traffic differently than they do for HTTP traffic. To filter encrypted traffic, the system determines the requested URL based on information passed during the SSL handshake: the subject common name in the public key certificate used to encrypt the traffic. There might be little or no relationship between the web site hostname in the URL and the subject common name.

You can improve HTTPS matching for category/reputation rules if you enable DNS request filtering. The system can determine the category and reputation during the DNS resolution phase, and block the DNS reply for unwanted combinations, before the user can start the HTTPS connection attempt. For allowed DNS responses, the system will have the category/reputation information available for subsequent HTTPS connections. See DNS Request Filtering, on page 483.

HTTPS filtering, unlike HTTP filtering, disregards subdomains within the subject common name. Do not include subdomain information when manually filtering HTTPS URLs. For example, use example.com rather than www.example.com. Also, review the content of the certificates used by the site to ensure you have the right domain, the one used in the subject common name, and that this name will not conflict with your other rules (for example, the name for a site you want to block might overlap with one you want to allow). For example, the subject common name in the certificate for youtube.com is *.google.com (this of course might change at any time).



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.

Controlling Traffic by Encryption Protocol

The system disregards the encryption protocol (HTTP vs HTTPS) when performing URL filtering. This occurs for both manual and reputation-based URL conditions. In other words, URL filtering treats traffic to the following web sites identically:

- http://example.com
- https://example.com

To configure a rule that matches only HTTP or HTTPS traffic, but not both, either specify the TCP port in the Destination condition or 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 TCP port or application, and URL, condition.

The first rule allows HTTPS traffic to the website:

Action: Allow
TCP port or Application: HTTPS (TCP port 443)
URL: example.com

The second rule blocks HTTP access to the same website:

Action: Block
TCP port or Application: HTTP (TCP port 80)
URL: example.com

Comparing URL and Application Filtering

URL and application filtering have similarities. But you should use them for very distinct purposes:

- URL filtering is best used to block or allow access to an entire web server. For example, if you do not want to allow any type of gambling on your network, you can create a URL filtering rule to block the Gambling category. With this rule, users cannot get to any pages on any web server within the category.
- Application filtering is useful for blocking specific applications regardless of the hosting site, or for blocking specific features of an otherwise allowable web site. For example, you could block just the Facebook Games application without blocking all of Facebook.

Because combining application and URL criteria can lead to unexpected results, especially for encrypted traffic, it is a good policy to create separate rules for URL and application criteria. If you do need to combine application and URL criteria in a single rule, you should place these rules after straight-forward 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. Because URL filtering block rules are more broad than application filtering, you should place them above application-only rules.

If you do combine application and URL criteria, you might need to monitor your network more carefully to ensure that you are not allowing access to unwanted sites and applications.

Best Practices for Effective URL Filtering

Please keep the following recommendations in mind when designing your URL filtering access control rules.

- Use category and reputation blocking whenever possible. This ensures that new sites get blocked automatically as they are added to the categories, and that blocking based on reputation is adjusted if a site becomes more (or less) reputable.
- When using URL category matching, note that there are cases where the login page for a site is in a different category than the site itself. For example, Gmail is in the Web-based Email category, whereas the login page is in the Search Engines and Portals category. If you have different rules with different actions for the categories, you might get unintended results.
- Use URL objects to target entire web sites and to make exceptions to category blocking rules. That is, to allow specific sites that would otherwise get blocked in a category rule.
- If you want to manually block a web server (using a URL object), it is much more effective to do so in the Security Intelligence policy. The Security Intelligence policy drops connections before the access control rules are evaluated, so you get a faster, more efficient, block.
- For the most effective filtering of HTTPS connections, implement SSL decryption rules to decrypt traffic for which you are writing an access control rule. Any decrypted HTTPS connections are filtered as HTTP connections in the access control policy, so you avoid all of the limitations for HTTPS filtering.
- TLS 1.3 encrpyts most handshake messages, so certificate information is not readily available. For traffic encrypted with TLS 1.3 to match access rules that use application or URL filtering effectively, the system must obtain a clear-text certificate for the server. We recommend that you enable **TLS 1.3 Certificate Visibility** in the access control settings. If you enable this option, the system checks whether a certificate for the site is stored in cache based on the IP address and Server Name Indication (SNI) in the client hello packet. If one is not available, the system uses a TLS 1.2 probe to obtain the certicate, which can then be used for application/URL category and reputation identification without decrypting the connection.
- Place URL blocking rules before any application filtering rules, because URL filtering blocks entire web servers, whereas application filtering targets specific application usage regardless of the web server.
- If you want to block high risk sites whose category is unknown, select the Uncategorized category and adjust the reputation slider to Questionable or Untrusted.
- You can improve overall URL filtering efficacy by also enabling DNS request filtering. When you use DNS request filtering, the system determines the FQDN's URL category and reputation at DNS lookup time, so the information is available if a subsequent HTTP/HTTPS request comes for the same destination. Additionally, if you block the category/reputation, the attempted connection gets stopped at the DNS request stage, rather than the web-session establishment stage. See DNS Request Filtering, on page 483.

What the User Sees When You Block Web Sites

When you block web sites with URL filtering rules, what the user sees differs based on whether the site is encrypted.

- HTTP connections—The user sees a system default block response page instead of the normal browser page for timed out or reset connections. This page should make it clear that you blocked the connection on purpose.
- HTTPS (encrypted) connections—The user does not see the system default block response page. Instead, the user sees the browser's default page for a secure connection failure. 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 will not be obvious from this message that you blocked the connection on purpose.

In addition, web sites might be blocked by other access control rules that are not explicitly URL filtering rules, or even by the default action. For example, if you block entire networks or geolocations, any web sites on that network or in that geographic location are also blocked. Users blocked by these rules may, or may not, get a response page as described in the limitations below.

If you implement URL filtering, consider explaining to end users what they might see when a site is intentionally blocked, and what types of site you are blocking. Otherwise, they might spend a good deal of time troubleshooting blocked connections.

Limitations of HTTP Response Pages

HTTP response pages do not always appear when the system blocks web traffic.

- 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).
- The system does not display a response page when web traffic is blocked before the system identifies the requested URL.
- The system does not display a response page for encrypted connections blocked by access control rules.

DNS Request Filtering

You can apply the URL category and reputation database to DNS lookup requests, even for connection attempts that are not HTTP/HTTPS.

For example, if a user tries to make an FTP connection to www.example.com, you can configure the system to look up the category and reputation of www.example.com when it sees the DNS lookup request for that fully-qualified domain name (FQDN). If your DNS/URL filtering rule for the returned category/reputation is a block rule, the system blocks the DNS reply. Thus, the user does not get an IP address for the FQDN, and their connection attempt fails.

By enabling DNS lookup request filtering, you can extend your URL filtering rules to protocols other than HTTP/HTTPS, and prevent FTP, TFTP, SCP, ICMP, and any other protocols from establishing a connection to a site that you are blocking for web access. This works so long as the user uses an FQDN name and thus requires a DNS lookup. If the user uses an IP address, there is no DNS request and DNS request blocking is not possible.

For HTTP/HTTPS traffic, doing the category/reputation lookup at DNS request time might improve system performance, as it can prevent the connection before the attempt to establish the web session. This might be especially beneficial for HTTPS, which is encrypted. By denying at the DNS request stage, the system never sees an HTTPS connection, and thus your decryption rules do not need to be evaluated, nor does the system need to perform the more difficult task of matching an encrypted session to the right access control rule.

Guidelines for DNS Request Filtering

Keep the following in mind as you configure DNS request filtering:

- DNS request filtering works on the DNS session only. If you allow the DNS reply (that is, the URL filtering rule action is Allow), then the subsequent connection that the user establishes with the returned IP address will be matched against your access control rules separately. The connection might match a

different rule, and thus be blocked or allowed for other reasons. For example, if you allow an FTP attempt to get an IP address through a DNS lookup, you might have another access control rule that prohibits FTP connections, and the connection will ultimately be blocked.

- DNS lookup requests that match access control rules that come before your URL/DNS request filtering rules will be allowed or blocked according to the matching rule. Category/reputation lookup will not be done for these connections.
- This feature requires that you implement URL filtering based on category/reputation. You must have the URL filtering license for this type of URL filtering. If you have no URL filtering rules based on category/reputation, then DNS request filtering is not relevant and you should not enable it.
- Connection events generated by DNS filtering include the following fields of special interest: DNS Query, URL Category, and URL Reputation. The DNS Query field shows the fully-qualified domain name (FQDN) for the lookup request. For DNS filtering events, the URL field will be blank.
- DNS request filtering uses the URL category and reputation database only. Any URL objects or other manual URL filtering defined in a matching access control rule are ignored. If you want to implement manual DNS name blocking, use the Security Intelligence DNS policy.

Filtering DNS Requests Based on URL Category and Reputation

The following procedure explains how to implement DNS lookup request filtering.

Before you begin

You must enable the URL license if it is not already enabled.

Procedure

Step 1 Select **Policies** > **Access Control**.

Step 2 If necessary, click the **Access Policy Settings** ( ) button, select the **Reputation Enforcement on DNS Traffic** option, and click **OK**.

This option enables DNS request filtering for the access control policy. The option is enabled by default.

Step 3 Evaluate existing URL filtering rules, or create new ones, to implement filtering based on URL category and reputation that will also apply to DNS requests.

URL filtering normally applies to HTTP/HTTPS traffic only, so there is no reason to restrict these rules based on application or port. However, if you do have these restrictions, ensure that the rule can also apply to DNS requests:

- On the **Source/Destination** tab, if the **Destination Ports** field has **Any**, no change is needed. If you specified ports, add **DNS over UDP** and **DNS over TCP** to the list.
- On the **Applications** tab, if the application list simply has **Any**, no change is needed. If you specified any applications or application filters, add the **DNS** application to the list or filter. The other DNS-related options are not relevant for this purpose.

For information on creating access control rules, see Configuring Access Control Rules, on page 491.

Step 4 Evaluate preceding rules to ensure that DNS requests do not match those rules.

Category and reputation determination happens only if the DNS request matches a URL filtering rule that has category and reputation specifications. Any DNS requests that match rules earlier in the access control policy than your URL filtering rule bypass DNS request filtering. Such DNS requests are handled according to the matching rule, either blocked or allowed.

Intrusion, File, and Malware Inspection

Intrusion and file policies work together as the last line of defense before traffic is allowed to its destination:

- Intrusion policies govern the system's intrusion prevention capabilities.
- File policies govern the system's file control and malware defense capabilities.

All other traffic handling occurs before network traffic is examined for intrusions, prohibited files, and malware. 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.

You can configure intrusion and file policies on rules that **allow** traffic only. Inspection is not performed on rules set to **trust** or **block** traffic. In addition, if the default action for the access control policy is **allow**, you can configure an intrusion policy but not a file policy.

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. Until a file is detected and blocked in a session, packets from the session may be subject to intrusion inspection.



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. Inspection works with unencrypted traffic only.

Best Practices for Access Control Rule Order

Rules are applied on a first-match basis, so you must ensure that rules with highly specific traffic matching criteria appear above policies that have more general criteria that would otherwise apply to the matching traffic. Consider the following recommendations:

- Specific rules should come before general rules, especially when the specific rules are exceptions to general rules.
- Any 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. We recommend they come before any rule that requires inspection, such as those with application or URL criteria, because Layer-3/4 criteria can be evaluated quickly and without inspection. Of course, any exceptions to these rules must be placed above them.
- Whenever possible, put specific drop rules near the top of the policy. This ensures the earliest possible decision on undesirable traffic.

- Any rules that include both application and URL criteria should come after straight-forward 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. Combining application and URL criteria can lead to unexpected results, especially for encrypted traffic, so we recommend that you create separate rules for URL and application filtering whenever possible.

NAT and Access Rules

Access rules always use the real IP addresses when determining an access rule match, even if you configure NAT. For example, if you configure NAT for an inside server, 10.1.1.5, so that it has a publicly routable IP address on the outside, 209.165.201.5, then the access rule to allow the outside traffic to access the inside server needs to reference the server's real IP address (10.1.1.5), and not the mapped address (209.165.201.5).

How Other Security Policies Impact Access Control

Other security policies can affect how access control rules function and match connections. As you configure your access rules, keep the following in mind:

- **SSL Decryption** policy—The SSL decryption rules are evaluated before access control. Thus, if an encrypted connection matches an SSL decryption rule that applies some type of decryption, it is the plain text (decrypted) connection that is evaluated by the access control policy. The access rules do not see the encrypted version of the connection. Additionally, any connections that match SSL decryption rules that drop traffic are never seen by the access control policy. Finally, any encrypted connection that matches a Do Not Decrypt rule is evaluated in its encrypted state.
- **Identity** policy—Connections are matched to users (and thus, user groups) only if there is a user mapping for the source IP address. Access rules that key on user or group membership can match only those connections for which user identity was successfully collected by your identity policy.
- **Security Intelligence** policy—Any connection that is dropped is never seen by the access control policy. Connections that match the do not block list are subsequently matched to access control rules and, ultimately, it is the access control rule that determines how the connection is handled (allowed or dropped).
- **VPN** (site-to-site or remote access)—VPN traffic is always evaluated against the access control policy, and connections are allowed or dropped based on the matching rule. However, the VPN tunnel itself is decrypted before the access control policy is evaluated. The access control policy evaluates the connections that are embedded within the VPN tunnel, not the tunnel itself.

License Requirements for Access Control

You do not need a special license to use the access control policy.

However, you do need the following licenses for specific features within the access control policy. For information on configuring licenses, see Enabling or Disabling Optional Licenses, on page 89.

- **URL** license—To create rules that use URL categories and reputations as match criteria.
- **Threat** license—To configure an intrusion policy on an access rule or the default action. You also need this license to use a file policy (the Malware license is also required).

- **Malware** license—To configure a file policy on an access rule. The Threat is also required for file policies.

Guidelines and Limitations for Access Control Policies

Following are some additional limitations for access control. Please consider them when evaluating whether you are getting the expected results from your rules.

- If a URL database update includes added (new, incoming), deprecated (outgoing), or deleted categories, there is a grace period for you to make changes to any access control rules that are affected. Impacted rules are marked with informational messages, with explanations about the issues that impact the rule and links to the Cisco Talos Intelligence Group (Talos) web site for more information about the category changes. You need to update the rule so that it uses the appropriate categories available in the latest URL database.

  To accommodate the grace period, add the newly added incoming categories to the appropriate rules without removing the outgoing deprecated categories: your rules should contain both the new and old categories. The new categories will be effective when the old categories are marked for deletion. When the old categories are finally deleted, you need to edit the rules to remove the deleted categories and redeploy the configuration. You will be blocked from deploying the configuration until you fix any rules that use deleted categories. Click the **See Problem Rules** link above the table to filter on rules that need attention.

- Firewall Device Manager can download information on up to 50,000 users from the directory server. If your directory server includes more than this number of user accounts, you will not see all possible names when selecting users in an access rule or when viewing user-based dashboard information. You can write rules on only those names that were downloaded.

  The limit also applies to the names associated with groups. If a group has more than the limit of members, only the first downloaded names up to the limit can be matched against the group membership.

- If a Vulnerability Database (VDB) update removes (deprecates) applications, you must make changes to any access control rules or application filters that use the application that was deleted. You cannot deploy changes until you fix these rules. In addition, you cannot install system software updates before fixing the issue. On the Application Filters object page, or the Application tab of the rule, these applications say “(Deprecated)” after the application name.

- To use fully-qualified domain name (FQDN) network objects as source or destination criteria, you must also configure DNS for the data interfaces on **Device** > **System Settings** > **DNS Server**. The system does not use the management DNS server setting to do lookups for FQDN objects used in access control rules. For information on troubleshooting FQDN resolution, see Troubleshooting General DNS Problems, on page 764.

  Note that controlling access by FQDN is a best-effort mechanism. Consider the following points:

  - Because DNS replies can be spoofed, only use fully trusted internal DNS servers.
  - Some FQDNs, especially for very popular servers, can have hundreds if not thousands of IP addresses, and these can frequently change. Because the system uses cached DNS lookup results, users might get addresses that are not yet in the cache, and their connections will not match the FQDN rule. Rules that use FQDN network objects function effectively only for names that resolve to fewer than 100 addresses.

We recommend that you do not create network object rules for an FQDN that resolves to more than 100 addresses, as the likelihood of the address in a connection being one that has been resolved and available in the DNS cache on the device is low. For these cases, use a URL-based rule instead of an FQDN network object rule.

- For popular FQDNs, different DNS servers can return a different set of IP addresses. Thus, if your users use a different DNS server than the one you configure, FQDN-based access control rules might not apply to all IP addresses for the site that are used by your clients, and you will not get the intended results for your rules.
- Some FQDN DNS entries have very short time to live (TTL) values. This can result in frequent recompliation of the lookup table, which can impact overall system performance.

- If you edit a rule that is actively in use, the changes do not apply to established connections that are no longer being inspected by Snort. The new rule is used to match against future connections. In addition, if Snort is actively inspecting a connection, it can apply the changed matching or action criteria to an existing connection. If you need to ensure that your changes apply to all current connections, you can log into the device CLI and use the **clear conn** command to end established 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.
- It takes 3 to 5 packets for the system to identify the application or URL in a connection. Thus, the correct access control rule might not be matched immediately for a given connection. However, once the application/URL is known, the connection is handled based on the matching rule. For encrypted connections, this happens after the server certificate exchange in the SSL handshake.
- The system applies the default policy action to packets that do not have a payload in a connection where an application is identified.
- Leave matching criteria empty whenever possible, especially those for security zones, network objects, and port objects. For example, the system can more efficiently match traffic for all interfaces if you simply leave the security zone criteria blank, rather than if you create zones that contain all interfaces. When you specify multiple criteria, the system must match against every combination of the contents of the criteria you specify.
- If you specify IP addresses for source or destination criteria, do not mix IPv4 and IPv6 addresses in the same rule. Create separate rules for IPv4 and IPv6 addresses.
- While operating, the Firewall Threat Defense device expands access control rules into multiple access control list entries based on the contents of any network objects used in the access rule. You can reduce the memory required to search access control rules by enabling object group search. With object group search enabled, the system does not expand network objects, but instead searches access rules for matches based on those group definitions. Object group search does not impact how your access rules are defined or how they appear in the Firewall Device Manager. It impacts only how the device interprets and processes them while matching connections to access control rules.

  Enabling object group search reduces memory requirements for access control policies that include network objects. However, it is important to note that object group search might also decrease rule lookup performance and thus increase CPU utilization. You should balance the CPU impact against the reduced memory requirements for your specific access control policy. In most cases, enabling object group search provides a net operational improvement.

  You can set this option using FlexConfig by issuing the **object-group-search access-control** command; use the **no** form of the command in the negate template.

- GRE tunnels that violate the related RFCs will be dropped. For example, if a GRE tunnel contains non-zero values in the reserved bits, contrary to the RFCs, it is dropped. If you need to allow non-compliant GRE tunnels, you need to use a remote manager and configure a prefilter rule that trusts the sessions. You cannot configure prefilter rules using the Firewall Device Manager.

Configuring the Access Control Policy

Use the access control policy to control access to network resources. The policy consists of a set of ordered rules, which are evaluated from top to bottom. The rule applied to traffic is the first one where all the traffic criteria are matched. If no rules match the traffic, the default action shown at the bottom of the page is applied.

To configure the access control policy, select **Policies** > **Access Control**.

The access control table lists all rules in order. For each rule:

- Click the > button next to the rule number in the left-most column to open the rule diagram. The diagram can help you visualize how the rule controls traffic. Click the button again to close the diagram.
- Most cells allow inline editing. For example, you can click the action to select a different one, or click a source network object to add or change the source criteria.
- To move a rule, hover over the rule until you get the move icon (), then click, drag, and drop the rule to the new location. You can also move a rule by editing it and selecting the new location in the **Order** list. It is critical that you put the rules in the order that you want them processed. Specific rules should be near the top, especially for rules that define exceptions to more general rules
- The right-most column contains the action buttons for a rule; mouse over the cell to see the buttons. You can edit () or delete () a rule.
- Click the **Access Control Settings** () button to configure settings that apply to the access control policy, rather than to specific rules within the policy.
- Click the **Toggle Hit Counts** icon () above the table to add or remove the Hit Counts column in the table. The Hit Count column appears to the right of the Name column with the total hit count for the rule and the date and time of the last hit. The hit count information is fetched at the time you click the toggle button. Click the **refresh** icon () to get the latest information.
- If any rules have problems, for example, because of removed or changed URL categories, click the **See Problem Rules** link next to the search box to filter the table to show only those rules. Please edit and correct (or delete) these rules, so that they will provide the service that you require.

The following topics explain how to configure the policy.

Configuring the Default Action

If a connection does not match a specific access rule, it is handled by the default action for the access control policy.

Procedure

Step 1 Select **Policies** > **Access Control**.

Step 2 Click anywhere in the **Default Action** field.

Step 3 Select the action to apply to matching traffic.

- **Trust**—Allow traffic without further inspection of any kind.
- **Allow**—Allow the traffic subject to the intrusion policy.
- **Block**—Drop the traffic unconditionally. The traffic is not inspected.

Step 4 If the action is **Allow**, select an intrusion policy.

For an explanation of the policy options, see Intrusion Policy Settings, on page 498.

Step 5 (Optional.) Configure logging for the default action.

You must enable logging for traffic that matches the default action to be included in dashboard data or Event Viewer. See Logging Settings, on page 499.

Step 6 Click **OK**.

Configuring Access Control Policy Settings

You can configure settings that apply to the access control policy, rather than to specific rules within the policy.

Procedure

Step 1 Select **Policy** > **Access Control**.

Step 2 Click the **Access Policy Settings** ( ) button.

Step 3 Configure the settings:

- **TLS Server Identity Discovery**—TLS 1.3 encrpyts most handshake messages, so certificate information is not readily available. For traffic encrypted with TLS 1.3 to match access rules that use application or URL filtering, the system must have a clear-text certificate for the server. If you enable this option, the system checks whether a certificate for the site is stored in cache based on the IP address and Server Name Indication (SNI) in the client hello packet. If one is not available, the system uses a TLS 1.2 probe to obtain the certicate, which can then be used for application/URL category and reputation identification. We recommend that you enable this option to ensure encrypted connections are matched to the right access control rule. The setting obtains the certificate only; the connection remains encrypted. Enabling this option is sufficient to obtain TLS 1.3 certificates; you do not need to create a corresponding SSL decryption rule. However, the cached certificates are also used for more effective decryption rule processing in addition to access control processing.
- **Reputation Enforcement on DNS Traffic**—Enable this option to apply your URL filtering category and reputation rules to DNS lookup requests. If the fully-qualified domain name (FQDN) in the lookup request has a category and reputation that you are blocking, the system blocks the DNS reply. Because

the user does not receive a DNS resolution, the user cannot complete the connection. Use this option to apply URL category and reputation filtering to non-web traffic. For more information, see DNS Request Filtering, on page 483.

Step 4 Click **OK**.

Configuring Access Control Rules

Use access control rules to control access to network resources. Rules in the access control policy are evaluated from top to bottom. The rule applied to traffic is the first one where all the traffic criteria are matched.

Procedure

Step 1 Select **Policies** > **Access Control**.

Step 2 Do any of the following:

- To create a new rule, click the + button.
- To edit an existing rule, click the edit icon ( ) for the rule.

To delete a rule you no longer need, click the delete icon ( ) for the rule.

Step 3 In **Order**, select where you want to insert the rule in the ordered list of rules.

Rules are applied on a first-match basis, so you must ensure that rules with highly specific traffic matching criteria appear above policies that have more general criteria that would otherwise apply to the matching traffic.

The default is to add the rule to the end of the list. If you want to change a rule's location later, edit this option.

Step 4 In **Title**, enter a name for the rule.

The name cannot contain spaces. You can use alphanumeric characters and these special characters: + . _ -

Step 5 Select the action to apply to matching traffic.

- **Trust**—Allow traffic without further inspection of any kind.
- **Allow**—Allow the traffic subject to the intrusion and other inspection settings in the policy.
- **Block**—Drop the traffic unconditionally. The traffic is not inspected.

Step 6 Define the traffic matching criteria using any combination of the following tabs:

- **Source/Destination**—The security zones (interfaces) through which the traffic passes, the IP addresses or the country or continent (geographical location) for the IP address, the security group tags (SGT) assigned to the address, or the protocols and ports used in the traffic. The default is any zone, address, geographical location, SGT, protocol, and port. See Source/Destination Criteria, on page 492.
- **Application**—The application, or a filter that defines applications by type, category, tag, risk, or business relevance. The default is any application. See Application Criteria, on page 494.
- **URL**—The URL or URL category of a web or DNS lookup request. The default is any URL. See URL Criteria, on page 496.

- **Users**—The identity source, user or user group. Your identity policies determine whether user and group information is available for traffic matching. You must configure identity policies to use this criteria. See User Criteria, on page 497.

To modify a condition, you click the **+** button within that condition, select the desired object or element, and click **OK** in the popup dialog box. If the criterion requires an object, you can click **Create New *Object*** if the object you require does not exist. Click the **x** for an object or element to remove it from the policy.

When adding conditions to access control rules, consider the following tips:

- 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, you can use a single rule to perform URL filtering for specific hosts or networks.
- 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 apply application control for up to 50 applications or application filters. Thus, there is an OR relationship among the items in a single condition, but an AND relationship between condition types (for example, between source/destination and application).
- Some features require that you enable the appropriate license.

Step 7 (Optional.) For policies that use the Allow action, you can configure further inspection on unencrypted traffic. Click one of the following links:

- **Intrusion Policy**—Select **Intrusion Policy** > **On** and select the intrusion inspection policy to inspect traffic for intrusions and exploits. See Intrusion Policy Settings, on page 498.
- **File Policy**—Select the file policy to inspect traffic for files that contain malware and for files that should be blocked. See File Policy Settings, on page 498.

Step 8 (Optional.) Configure logging for the rule.

By default, connection events are not generated for traffic that matches a rule, although file events are generated by default if you select a file policy. You can change this behavior. You must enable logging for traffic that matches the policy to be included in dashboard data or Event Viewer. See Logging Settings, on page 499.

Intrusion events are always generated for intrusion rules set to drop or alert regardless of the logging configuration on the matching access rule.

Step 9 Click **OK**.

Source/Destination Criteria

The Source/Destination criteria of an access rule define the security zones (interfaces) through which the traffic passes, the IP addresses or the country or continent (geographical location) for the IP address, the security group tags (SGT) assigned to the address, or the protocols and ports used in the traffic. The default is any zone, address, geographical location, SGT, protocol, and port.

To modify a condition, you click the **+** button within that condition, select the desired object or element, and click **OK**. If the criterion requires an object, you can click **Create New *Object*** if the object you require does not exist. Click the **x** for an object or element to remove it from the policy.

You can use the following criteria to identify the source and destination to match in the rule.

Source Zones, Destination Zones

The security zone objects that define the interfaces through which the traffic passes. You can define one, both, or neither criteria: any criteria not specified applies to traffic on any interface.

- To match traffic leaving the device from an interface in the zone, add that zone to the **Destination Zones**.
- To match 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.

Use this criteria when the rule should apply based on where the traffic enters or exits the device. For example, if you want to ensure that all traffic going to inside hosts gets intrusion inspection, you would select your inside zone as the **Destination Zones** while leaving the source zone empty. To implement intrusion filtering in the rule, the rule action must be **Allow**, and you must select an intrusion policy in the rule.



Note You cannot mix passive and routed security zones in a single rule. In addition, you can specify passive security zones as source zones only, you cannot specify them as destination zones.

Source Networks, Destination Networks

The network objects or geographical locations that define the network addresses or locations of the traffic.

- To match traffic from an IP address or geographical location, configure the **Source Networks**.
- To match 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 traffic must originate from one of the specified IP addresses and be destined for one of the destination IP addresses.

When you add this criteria, you select from the following tabs:

- **Network**—Select the network objects or groups that define the source or destination IP addresses for the traffic you want to control. You can use objects that define the address using the fully-qualified domain name (FQDN); the address is determined through a DNS lookup.
- **Geolocation**—Select the geographical location to control traffic based on its source or destination country or continent. Selecting a continent selects all countries within the continent. Besides selecting geographical location directly in the rule, you can also select a geolocation object that you created to define the location. Using geographical location, you could easily restrict access to a particular country without needing to know all of the potential IP addresses used there.



Note To ensure that you are using up-to-date geographical location data to filter your traffic, Cisco strongly recommends that you regularly update the geolocation database (GeoDB).

Source Ports, Destination Ports/Protocols

The port objects that define the protocols used in the traffic. For TCP/UDP, this can include ports. For ICMP, it can include codes and types.

- To match traffic from a protocol or port, configure the **Source Ports**. Source ports can be TCP/UDP only.
- To match traffic to a protocol or port, configure the **Destination Ports/Protocols**. If you add only destination ports to a condition, you can add ports that use different transport protocols. ICMP and other non-TCP/UDP specifications are allowed in destination ports only; they are not allowed in source ports.
- To match traffic both originating from specific TCP/UDP ports and destined for specific TCP/UDP ports, configure both. If you add both source and destination ports to a condition, you can only add ports that share a single transport protocol, TCP or UDP. For example, you could target traffic from port TCP/80 to port TCP/8080.

Source SGT Groups, Destination SGT Groups

The Security Group Tag (SGT) group objects that identify the SGTs assigned to the traffic, as downloaded from the Identity Services Engine (ISE). You can use these objects only if you define an ISE identity source; otherwise, this section will not appear. For detailed information on how to use SGTs for access control, see How to Control Network Access Using TrustSec Security Group Tags, on page 503.

- To match traffic whose source has one of the SGT defined in the group, configure the **Source SGT Groups**.
- To match traffic to a destination that has one of the SGT defined in the group, configure the **Destination SGT Groups**.
- If you add both source and destination SGT conditions to a rule, matching traffic must originate from a source with one of the specified tags and be destined for one of the destination tags.

Application Criteria

The Application criteria of an access rule defines the application used in an IP connection, or a filter that defines applications by type, category, tag, risk, or business relevance. The default is any application.

Although you can specify individual applications in the rule, application filters simplify policy creation and administration. For example, you could 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 session is blocked.

In addition, Cisco frequently updates and adds additional application detectors via system and vulnerability database (VDB) updates. Thus, a rule blocking high risk applications can automatically apply to new applications without you having to update the rule manually.

You can specify applications and filters directly in the rule, or create application filter objects that define those characteristics. The specifications are equivalent, although using objects can make it easier to stay within the 50-items-per-criteria system limit if you are creating a complex rule.

To modify the application and filters list, you click the + button within the condition, select the desired applications or application filter objects, which are listed on separate tabs, and click **OK** in the popup dialog box. On either tab, you can click **Advanced Filter** to select filter criteria or to help you search for specific applications. Click the **x** for an application, filter, or object to remove it from the policy. Click the **Save As Filter** link to save the combined criteria that is not already an object as a new application filter object.

Note If a selected application was removed by a VDB update, "(Deprecated)" appears after the application name. You must remove these applications from the filter, or subsequent deployments and system software upgrades will be blocked.

You can use the following **Advanced Filter** criteria to identify the application or filter to match in the rule. These are the same elements used in application filter objects.

Note Multiple selections within a single filter criteria have an OR relationship. For example, Risk is High OR Very High. The relationship between filters is AND, so Risk is High OR Very High, AND Business Relevance is Low OR Very Low. As you select filters, the list of applications in the display updates to show only those that meet the criteria. You can use these filters to help you find applications that you want to add individually, or to verify that you are selecting the desired filters to add to the rule.

Risks

The likelihood that the application is used for purposes that might be against your organization's security policy, from very low to very high.

Business Relevance

The likelihood that the application is used within the context of your organization's business operations, as opposed to recreationally, from very low to very high.

Types

The type of application:

- **Application Protocol**—Application protocols such as HTTP and SSH, which represent communications between hosts.
- **Client Protocol**—Clients such as web browsers and email clients, which represent software running on the host.
- **Web Application**—Web applications such as MPEG video and Facebook, which represent the content or requested URL for HTTP traffic.

Categories

A general classification for the application that describes its most essential function.

Tags

Additional information about the application, similar to category.

For encrypted traffic, the system can identify and filter traffic using only the applications tagged **SSL Protocol**. Applications without this tag can only be detected in unencrypted or decrypted traffic. Also, the system assigns the **decrypted traffic** tag to applications that the system can detect in decrypted traffic only, not encrypted or unencrypted.

Applications List (bottom of the display)

This list updates as you select filters from the options above the list, so you can see the applications that currently match the filter. Use this list to verify that your filter is targeting the desired applications when

you intend to add filter criteria to the rule. If your intention is to add specific applications, select them from this list.

URL Criteria

The URL criteria of an access rule defines the URL used in a web request, or the category to which the requested URL belongs. For category matches, you can also specify the relative reputation of sites to allow or block. The default is to allow all URLs.

If you enable DNS lookup request filtering, the category and reputation settings also apply to the fully-qualified domain name (FQDN) in the lookup request. Only the category and reputation settings apply for DNS request filtering. Manual URL filtering is ignored.

URL categories and reputations allow you to quickly create URL conditions for access control rules. For example, you could block all Gambling sites, or untrusted Social Networking sites. If a user attempts to browse to any URL with that category and reputation combination, the session is blocked.

Using category and reputation data also simplifies policy creation and administration. It grants you assurance that the system will control web traffic as expected. Finally, because Cisco's threat intelligence is continually updated with new URLs, as well as new categories and risks for existing URLs, you can ensure that the system uses up-to-date information to filter requested URLs. Malicious sites that represent security threats such as malware, spam, botnets, and phishing may appear and disappear faster than you can update and deploy new policies.

To modify the URL list, you click the **+** button within the condition and select the desired categories or URLs using one of the following techniques. Click the **x** for a category or object to remove it from the policy.

URL Tab

Click **+**, select URL objects or groups, and click **OK**. You can click **Create New URL** if the object you require does not exist.



Note Before configuring URL objects to target specific sites, carefully read the information on manual URL filtering.

Categories Tab

Click **+**, select the desired categories, and click **OK**.

For a description of the categories, see https://www.talosintelligence.com/categories.

The default is to apply the rule to all URLs in each selected category regardless of reputation. To limit the rule based on reputation, click the down arrow for each category, deselect the **Any** checkbox, and then use the **Reputation** slider to choose the reputation level. The left of the reputation slider indicates sites that will be allowed, the right side are sites that will be blocked. How reputation is used depends on the rule action:

- If the rule blocks or monitors web access, selecting a reputation level also selects all reputations more severe than that level. For example, if you configure a rule to block or monitor **Questionable sites** (level 2), it also automatically blocks or monitors **Untrusted** (level 1) sites.
- If the rule allows web access, selecting a reputation level also selects all reputations less severe than that level. For example, if you configure a rule to allow **Favorable sites** (level 4), it also automatically allows **Trusted** (level 5) sites.

Select the **Include Sites with Unknown Reputation** option to have URLs with unknown reputation included in the reputation match. New sites typically are unrated, and there can be other reasons a site's reputation is unknown or cannot be determined.

Check the Category for a URL

You can check on the category and reputation for a particular URL. Enter the URL in the **URL to Check** box and click **Go**. You will be taken to an external website to see the results. If you disagree with a categorization, click the **Submit a URL Category Dispute** link and let us know.

User Criteria

The User criteria of an access rule defines the user or user group for an IP connection. You must configure identity policies and the associated directory server to include user or user group criteria in an access rule.

Your identity policies determine whether user identity is collected for a particular connection. If identity is established, the IP address of the host is associated with the identified user. Thus, traffic whose source IP address is mapped to a user is considered to be from that user. IP packets themselves do not include user identity information, so this IP-address-to-user mapping is the best approximation available.

Because you can add a maximum of 50 users or groups to a rule, selecting groups usually makes more sense than selecting individual users. For example, you could create a rule allowing the Engineering group access to a development network, and create a subsequent rule that denies all other access to the network. Then, to make the rule apply to new engineers, you only need to add the engineer to the Engineering group in the directory server.

You an also select identity sources to apply to all users within that source. Thus, if you support multiple Active Directory domains, you can provide differential access to resources based on the domain.

To modify the users list, you click the **+** button within the condition and select the desired identities using one of the following techniques. Click the **x** for an identity to remove it from the policy.

- **Identity Sources**—Select an identity source, such as an AD realm or the local user database, to apply the rule to all users obtained from the selected sources. If the realm you need does not yet exist, click **Create New Identity Realm** and create it now.
- **Groups**—Select the desired user groups. Groups are available only if you configure groups in the directory server. If you select a group, the rule applies to any member of the group, including subgroups. If you want to treat a sub-group differently, you need to create a separate access rule for the sub-group and place it above the rule for the parent group in the access control policy.
- **Users**—Select individual users. The user name is prefixed with the identity source, such as Realm\username.

  There are some built-in users under the Special-Identities-Realm:

  - **Failed Authentication**—The user was prompted to authenticate, but failed to enter a valid username/password pair within the maximum number of allowed attempts. Failure to authenticate does not itself prevent the user from accessing the network, but you can write an access rule to limit network access for these users.
  - **Guest**—Guest users are like Failed Authentication users, except that your identity rule is configured to call these users Guest. Guest users were prompted to authenticate and failed to do so within the maximum number of attempts.
  - **No Authentication Required**—The user was not prompted to authentication, because the user's connections matched identity rules that specified no authentication.

- **Unknown**—There is no user mapping for the IP address, and there is no record of failed authentication yet. Typically, this means that no HTTP traffic has yet been seen from that address.

Intrusion Policy Settings

Cisco delivers several intrusion policies with the firewall system. The several intrusion policies delivered by Cisco Cisco Talos Intelligence Group (Talos) are designed by the Cisco.Talos set the intrusion and preprocessor rule states and advanced settings. For access control rules that allow traffic, you can select an intrusion policy to inspect traffic for intrusions and exploits. An intrusion policy examines decoded packets for attacks based on patterns, and can block or alter malicious traffic.

When running Snort 2, these are the only policies available and you cannot modify them. However, you can change the action to take for a given rule, as described in Changing Intrusion Rule Actions (Snort 2), on page 535.

When running Snort 3, you can select one of these policies, or you can create your own intrusion policies.

To enable intrusion inspection, select **Intrusion Policy** > **On** and select the desired policy. Click the info icon for a policy in the drop-down list to see a description for each policy.

For more information on the pre-defined policies, see System-Defined Network Analysis and Intrusion Policies, on page 510.

File Policy Settings

Use file policies to detect malicious software, or *malware*, using malware defense. You can also use file policies to perform file control, which allows control over all files of a specific type regardless of whether the files contain malware.

Malware defense uses the AMP Cloud to retrieve dispositions for possible malware detected in network traffic, and to obtain local malware analysis and file pre-classification updates. The management interface must have a path to the Internet to reach the AMP Cloud and perform malware lookups. When the device detects an eligible file, it uses the file's SHA-256 hash value to query the AMP Cloud for the file's disposition. The possible dispositions are:

- Malware—The AMP Cloud categorized the file as malware. An archive file (e.g. a zip file) is marked as malware if any file within it is malware.
- Clean—The AMP Cloud categorized the file as clean, containing no malware. An archive file is marked as clean if all files within it are clean.
- Unknown—The AMP Cloud has not assigned a disposition to the file yet. An archive file is marked as unknown if any file within it is unknown.
- Unavailable—The system could not query the AMP Cloud to determine the file's disposition. You may see a small percentage of events with this disposition; this is expected behavior. If you see a number of "unavailable" events in succession, ensure that the Internet connection for the management address is functioning correctly.

Available File Policies

You can select one of the following file policies:

- **None**—Do not evaluate transmitted files for malware and do no file-specific blocking. Select this option for rules where file transmissions are trusted or where they are unlikely (or impossible), or for rules where you are confident your application or URL filtering adequately protects your network.
- **Block Malware All**—Query the AMP Cloud to determine if files traversing your network contain malware, then block files that represent threats.
- **Cloud Lookup All**—Query the AMP Cloudto obtain and log the disposition of files traversing your network while still allowing their transmission.
- **(Custom File Policy)**—You can create your own file policies using the Firewall Threat Defense API filepolicies resource, and the other FileAndMalwarePolicies resources (such as filetypes, filetypecategories, ampcloudconfig, ampservers, and ampcloudconnections). After you create the policies, and deploy changes, you can select your policies when editing an access control rule in the Firewall Device Manager. The policy description appears below the policy when you select it.

Logging Settings

The logging settings for an access rule determine whether connection events are issued for traffic that matches the rule. You must enable logging to see events related to the rule in the Event Viewer. You must also enable logging for matching traffic to be reflected in the various dashboards you can use to monitor the system.

You should log connections according to the security and compliance needs of your organization. If your goal is to limit the number of events you generate and improve performance, only enable logging for the connections critical to your analysis. However, if you want a broad view of your network traffic for profiling purposes, you can enable logging for additional connections.



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 a Block rule, consider whether the rule is for an Internet-facing interface or other interface vulnerable to DoS attack.

You can configure the following logging actions.

Select Log Action

You can select one of the following actions:

- **Log at Beginning and End of Connection**—Issue events at the start and end of a connection. Because end-of-connection events contain everything that start-of-connection events contain, plus all of the information that could be gleaned during the connection, Cisco recommends that you do not select this option for traffic that you are allowing. Logging both events can impact system performance. However, this is the only option allowed for blocked traffic.
- **Log at End of Connection**—Select this option if you want to enable connection logging at the end of the connection, which is recommended for allowed or trusted traffic.
- **No Logging at Connection**—Select this option to disable logging for the rule. This is the default.

Note When an intrusion policy invoked by an access control rule detects an intrusion and generates an intrusion event, the system automatically logs the end of the connection where the intrusion occurred, regardless of the logging configuration of the rule. For connections where an intrusion was blocked, the action for the connection in the connection log is **Block**, with a reason of **Intrusion Block**, even though to perform intrusion inspection you must use an Allow rule.

File Events

Select **Log Files** if you want to enable logging of prohibited files or malware events. You must select a file policy in the rule to configure this option. The option is enabled by default if you select a file policy for the rule. Cisco recommends you leave this option enabled.

When the system detects a prohibited file, it automatically logs one of the following types of event:

- *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.

For connections where a file was blocked, the action for the connection in the connection log is **Block** even though to perform file and malware inspection you must use an Allow rule. The connection's Reason is either **File Monitor** (a file type or malware was detected), or **Malware Block** or **File Block** (a file was blocked).

Send Connection Events To

If you want to send a copy of the events to an external syslog server, select the server object that defines the syslog server. If the required object does not already exist, click **Create New Syslog Server** and create it. (To disable logging to a syslog server, select **Any** from the server list.)

Because event storage on the device is limited, sending events to an external syslog server can provide more long term storage and enhance your event analysis.

This setting applies to connection events only. To send intrusion events to syslog, configure the server in the intrusion policy settings. To send file/malware events to syslog, configure the server on **Device** > **System Settings** > **Logging Settings**.

Monitoring Access Control Policies

The following topics explain how you can monitor the access control policy.

Monitoring Access Control Statistics in the Dashboards

Most of the data on the **Monitoring** dashboards are directly related to your access control policy. See Monitoring Traffic and System Dashboards, on page 102.

- **Monitoring** > **Access And SI Rules** shows the most-hit access rules and Security Intelligence rule-equivalents and related statistics.
- You can find general statistics on the **Network Overview**, **Destinations**, and **Zones**, dashboards.

- You can find URL filtering results on the **URL Categories** and **Destinations** dashboards. You must have at least one URL filtering policy to see any information on the **URL Categories** dashboard.
- You can find application filtering results on the **Applications** and **Web Applications** dashboards.
- You can find user-based statistics on the **Users** dashboard. You must implement identity policies to collect user information.
- You can find intrusion policy statistics on the **Attackers** and **Targets** dashboards. You must apply an intrusion policy to at least one access control rule to see any information on these dashboards.
- You can find file policy and malware filtering statistics on the **File Logs** and **Malware** dashboards. You must apply a file policy to at least one access control rule to see any information on these dashboards.
- **Monitoring** > **Events** also shows events for connections and data related to the access control rules.

Examining Rule Hit Counts

You can view the hit count for each access control rule. The hit count indicates how often connections matched the rule. You can use this information to identify your most active rules and the rules that are less active.

The count is from the last time the system rebooted, whether from your action or from a system upgrade, or from when you reset the hit count for a rule or all rules.

You can also see rule hit count information in the device CLI using the **show rule hits** command.

Procedure

Step 1 Select **Policies** > **Access Control**.

Step 2 Click the **Toggle Hit Counts** icon ().

The Hit Count column appears to the right of the Name column with the total hit count for the rule and the date and time of the last hit. The hit count information is fetched at the time you click the toggle button.

You can do the following with the hit count information:

- To the left of the button, you will see information on when the hit count was last updated. Click the **refresh** icon () to get the latest numbers.
- To open a detailed view of the hit count for a given rule, click the hit count number in the table to open the Hit Count dialog box. The hit count information includes the number of hits and the date and time of the last connection that matched the rule. Click the **Reset** link to reset the counter to zero.

  If you want to reset the hit count for all rules at once, open an SSH session to the device and issue the **clear rule hits** command.
- Click the **Toggle Hit Counts** icon () again to remove the hit count column from the table.

Monitoring Syslog Messages for Access Control

In addition to seeing events in the Event Viewer, you can configure access control rules, intrusion policies, file/malware policies, and Security Intelligence policies to send events to a syslog server. The events use the following message IDs:

- 430001—Intrusion event.
- 430002—Connection event logged at the beginning of a connection.
- 430003—Connection event logged at the end of a connection.
- 430004—File events.
- 430005—Malware events.

Monitoring Access Control Policies in the CLI

You can also open the CLI console or log into the device CLI and use the following commands to get more detailed information about access control policies and statistics.

- **show access-control-config** displays summary information about the access control rules along with per-rule hit counts.
- **show access-list** displays the access control lists (ACLs) that were generated from the access control rules. The ACLs provide an initial filter and attempt to provide quick decisions whenever possible, so that connections that should be dropped do not need to be inspected (and thus consume resources unnecessarily). This information includes hit counts.
- **show rule hits** displays consolidated hit counts that are more accurate than the counts shown with **show access-control-config** and **show access-list**. If you want to reset the hit count, use the **clear rule hits** command.
- **show snort statistics** displays information about the Snort inspection engine, which is the main inspector. Snort implements application filtering, URL filtering, intrusion protection, and file and malware filtering.
- **show conn** displays information about the connections currently established through the interfaces.
- **show traffic** displays statistics about traffic flowing through each interface.
- **show ipv6 traffic** displays statistics about IPv6 traffic flowing through the device.

Examples for Access Control

The use case chapter includes several examples of implementing access control rules. Please see the following examples:

- How to Gain Insight Into Your Network Traffic, on page 44. This example shows some basic ideas for collecting overall connection and user information.
- How to Block Threats, on page 51. This example shows how to apply intrusion policies.
- How to Block Malware, on page 56. This example shows how to apply file policies.

- How to Implement an Acceptable Use Policy (URL Filtering), on page 59. This example shows how to perform URL filtering.
- How to Control Application Usage, on page 64. This example shows how to perform application filtering.
- How to Add a Subnet, on page 67. This example shows how to integrate a new subnet into your overall network, including the access rules needed to allow traffic flow.
- How to Passively Monitor the Traffic on a Network, on page 72

Following are additional examples.

How to Control Network Access Using TrustSec Security Group Tags

If you use Cisco Identity Services Engine (ISE) to define and use security group tags (SGT) for classifying traffic in a Cisco TrustSec network, you can write access control rules that use SGT as matching criteria. Thus, you can block or allow access based on security group membership rather than IP addresses directly.

About Security Group Tags (SGT)

In Cisco Identity Services Engine (ISE), you can create security group tags (SGT) and assign host or network IP addresses to each tag. You can also assign SGTs to user accounts, and the SGT is assigned to the user's traffic. If the switches and routers in the network are configured to do so, these tags are then assigned to packets as they enter the network controlled by ISE, the Cisco TrustSec cloud.

When you configure an ISE identity source in the Firewall Device Manager, the Firewall Threat Defense system automatically downloads the list of SGTs from ISE. You can then use SGT as a traffic matching condition in access control rules.

For example, you could create a Production Users tag, and associate the 192.168.7.0/24 network to the tag. This would be appropriate if you use that network for user end points, such as laptops, Wi-Fi clients, and similar devices. You could create a separate tag for Production Servers, and assign the IP addresses of the relevant servers or subnet to the tag. Then, in the Firewall Threat Defense, you could allow or block access from the user network to the production servers based on the tag. If you later alter the host or network addresses associated with the tag in ISE, you do not need to change the access control rule defined for the Firewall Threat Defense device.

When the Firewall Threat Defense evaluates SGT as a traffic matching criteria for an access control rule, it uses the following priority:

1. The source SGT tag defined in the packet, if any. For the SGT tag to be in the packet, the switches and routers in the network must be configured to add them. See the ISE documentation for information on how to implement this method.
2. The SGT assigned to the user session, as downloaded from the ISE session directory. You need to enable the option to listen to session directory information for this kind of SGT matching, but this option is on by default when you first create the ISE identity source. The SGT can be matched to source or destination. Although not required, you would also normally set up a passive authentication identity rule, using the ISE identity source along with an AD realm, to collect user identity information.
3. The SGT-to-IP address mapping downloaded using SXP. If the IP address is within the range for an SGT, then the traffic matches the access control rule that uses the SGT. The SGT can be matched to source or destination.

ISE uses the Security-group eXchange Protocol (SXP) to propagate the IP-to-SGT mapping database to network devices. When you configure the Firewall Threat Defense device to use an ISE server, you must turn on the option to listen to the SXP topic from ISE. Thus, the Firewall Threat Defense device learns about the security group tags and mappings directly from ISE, and is notified whenever ISE publishes updated security group tags and mappings. This ensures that the list of security group tags and mappings stay up-to-date on the device, so that the Firewall Threat Defense can effectively enforce the policy defined in ISE.

Configure security groups and SXP publishing in ISE

This task enables static SGT-to-IP address mappings for Secure Firewall Threat Defense devices, allowing TrustSec policies and access control based on source and destination security group tags (SGT).

There is a lot of configuration that you must do in Cisco Identity Services Engine (ISE) to create the TrustSec policy and security group tags (SGT). Please look at the ISE documentation for more complete information on implementing TrustSec.

The following procedure picks out the highlights of the core settings you must configure in ISE for the Firewall Threat Defense device to be able to download and apply static SGT-to-IP address mappings, which can then be used for source and destination SGT matching in access control rules. See the ISE documentation for detailed information.

The figures in this procedure are based on ISE 2.4. The exact paths to these features might change in subsequent releases, but the concepts and requirements will be the same. Although ISE 2.4 or later is recommended, and preferably 2.6 or later, the configuration should work starting with ISE 2.2 patch 1.

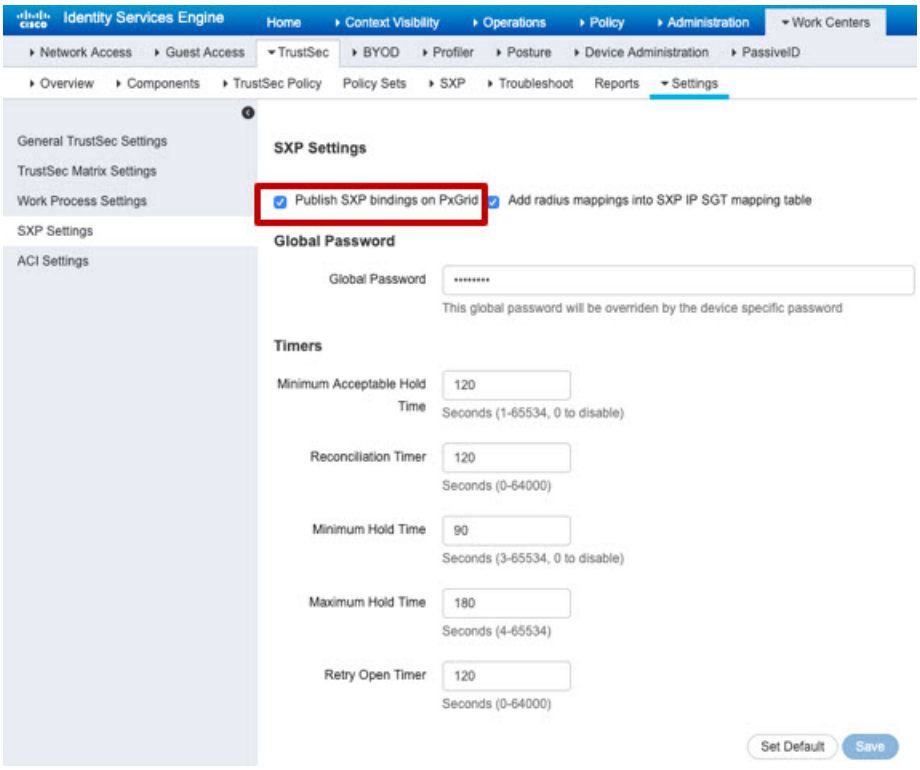
Before you begin

You must have the ISE Plus license to publish SGT-to-IP address static mappings and to get user session-to-SGT mappings so that the Firewall Threat Defense device can receive them.

Procedure

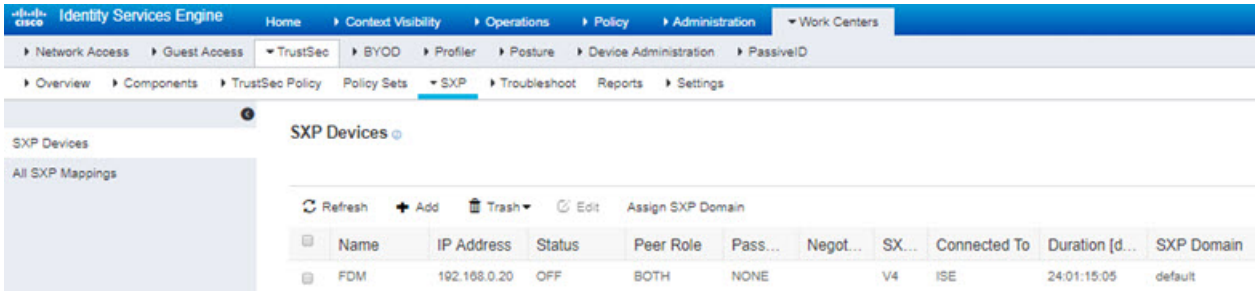
Step 1 Choose **Work Centers** > **TrustSec** > **Settings** > **SXP Settings**, and select the **Publish SXP Bindings on PxGrid** option.

This option makes ISE send the SGT mappings out using SXP. You must select this option for the Firewall Threat Defense device to "hear" anything from listing to the SXP topic. This option must be selected for the Firewall Threat Defense device to get static SGT-to-IP address mapping information. It is not necessary if you simply want to use SGT tags defined in the packets, or SGTs that are assigned to a user session.

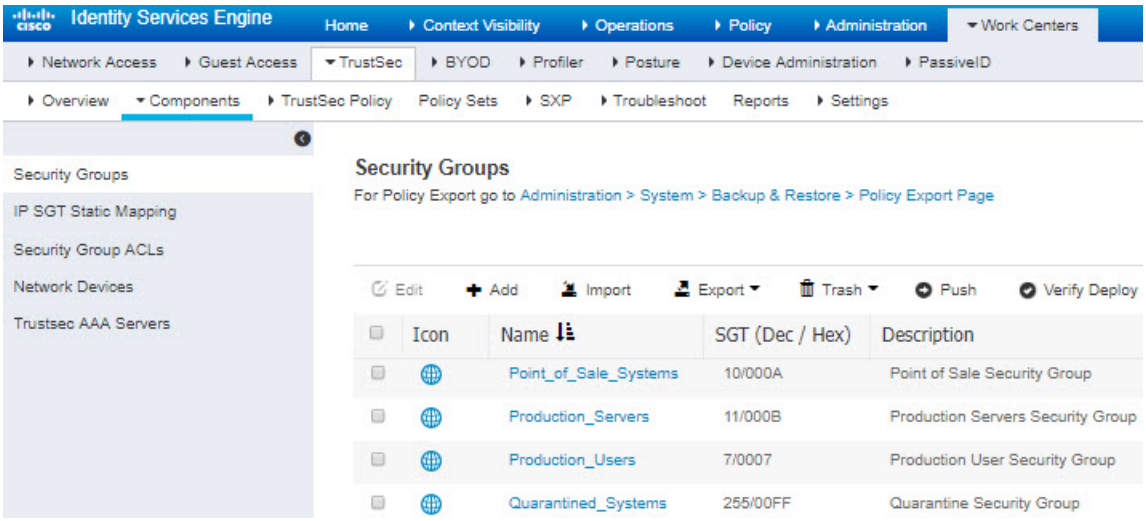


Step 2 Choose **Work Centers** > **TrustSec** > **SXP** > **SXP Devices**, and add a device.

This does not have to be a real device, you can even use the management IP address of the Firewall Threat Defense device. The table simply needs at least one device to induce ISE to publish the static SGT-to-IP address mappings. This step is not necessary if you simply want to use SGT tags defined in the packets, or SGTs that are assigned to a user session.

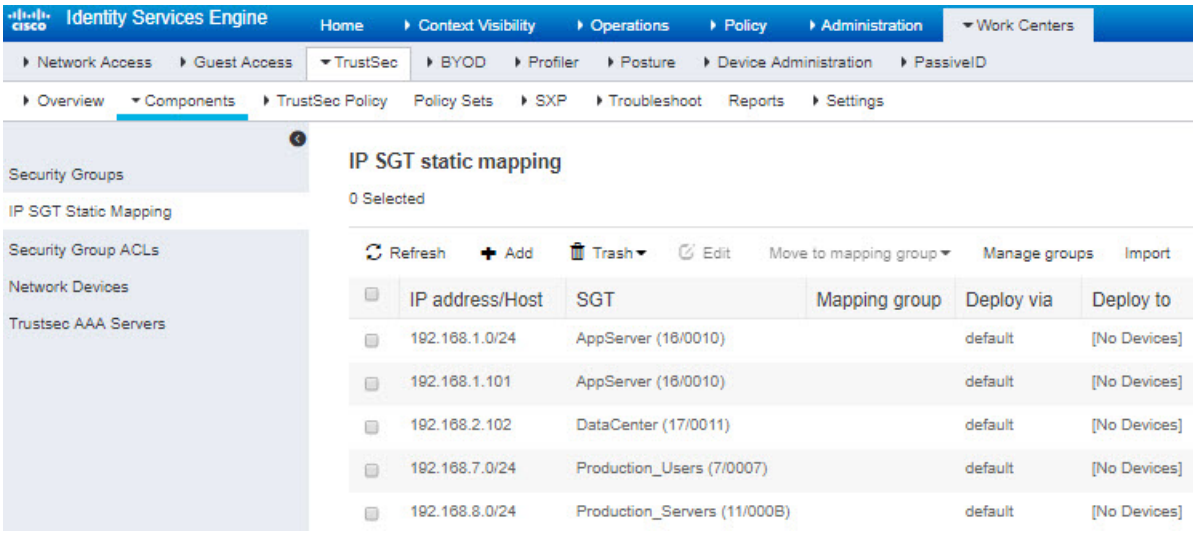


Step 3 Choose **Work Centers** > **TrustSec** > **Components** > **Security Groups** and verify there are security group tags defined. Create new ones as necessary.



Step 4 Choose **Work Centers** > **TrustSec** > **Components** > **IP SGT Static Mapping** and map host and network IP addresses to the security group tags.

This step is not necessary if you simply want to use SGT tags defined in the packets, or SGTs that are assigned to a user session.



You have configured ISE with the necessary security groups, SXP publishing settings, and static SGT-to-IP mappings required for Secure Firewall Threat Defense devices to receive and apply TrustSec policies.

Configure Access Control Based on Security Group Tag (SGT)

To configure access control rules that use security group tags (SGT) as matching criteria, you must first configure the device to obtain the SGT mappings from an ISE server.

The following procedure explains the end-to-end process based on the assumption that you want to get all mappings that are defined in ISE, including SGT-to-IP address mappings published through SXP. Alternatively:

- If you want to use SGT information in the packets only, and not use mappings downloaded from ISE, simply create SGT group dynamic objects and use them as source SGT criteria in access control rules. Note that in this case, you can use SGT tags as a source condition only; these tags will never match destination criteria.
- If you want to use SGT in packets plus user session SGT mappings only, you do not need to turn on the option to subscribe to the SXP topic in the ISE identity source, nor do you need to configure ISE to publish SXP mappings. You can use this information for both source and destination matching conditions.

Before you begin

The assumption here is that you have already configured Cisco TrustSec in your network, and you are simply adding the Firewall Threat Defense device as a policy enforcement point. If you have not deployed Cisco TrustSec, please start with ISE and configure your network, then return to this procedure. Explaining Cisco TrustSec is outside the scope of this document.

Procedure

Step 1 Ensure that SGTs are defined, that ISE is configured correctly to publish the SXP topic, and that any needed static mappings are in place.

See Configure security groups and SXP publishing in ISE, on page 504.

Step 2 Update the Identity Services Engine object to listen to the SXP topic.

You can use ISE to obtain user session SGT mappings, static SGT-to-IP address mappings through SXP, or both. By default, when you configure the ISE identity source, it obtains user session mappings only; you must turn on the option to listen to the SXP topic from ISE.

a) Choose **Objects** > **Identity Sources**.
b) Edit the ISE object. If you have not configured one yet, click **+** > **Identity Services Engine** and see Configure Identity Services Engine, on page 170.
c) Under **Subscribe To**, select **SXP Topic**.

Ensure that **Session Directory Topic** is also selected if you are using passive authentication or want user-to-SGT mappings.



d) Click **OK**.

Step 3 Deploy your changes and wait for the system to download tags and mappings from ISE.

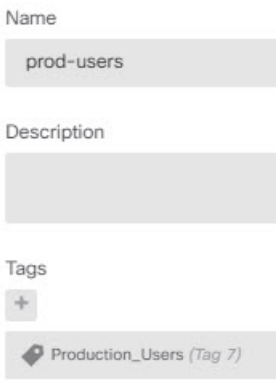
After you configure the ISE identity source and deploy changes, the system retrieves security group tag (SGT) information from the ISE server. The download will not happen until after you deploy changes.

Step 4 Create the SGT group objects that are required for your access control rules.

You cannot use the information retrieved from ISE directly in an access control rule. Instead, you need to create SGT groups, which refer to the downloaded SGT information. Your SGT groups can refer to more than one SGT, so you can apply policy based on relevant collections of tags if that is appropriate.

The number and contents of the objects depend on the access control rules that you intend to write. Repeat the following process to create all the objects that you need.

a) Choose **Objects** > **SGT Groups**.
b) Click **+** to add a new object, or edit an existing object.
c) For new objects, enter a name and optionally a description.
d) Under **Tags**, click **+** and select all the tags that should be included in the group.



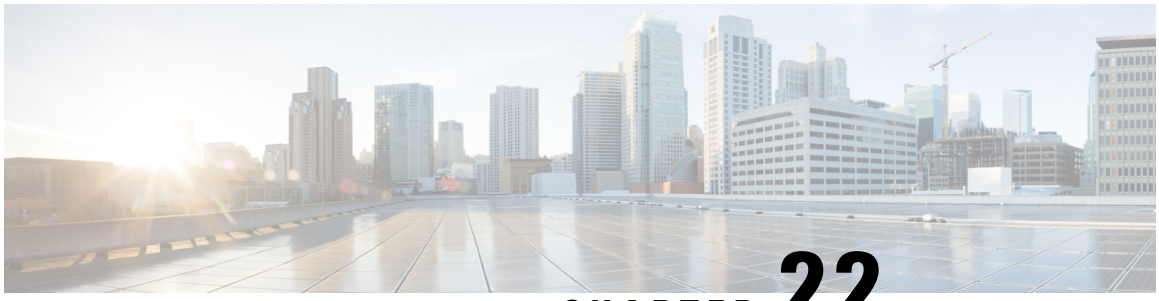
e) Click **OK**.

Step 5 Create access control rules that use the SGT group objects.

As an example, the following rule allows traffic from production users to production servers. The rule depends completely on SGT; it is not limited by source/destination interface or any other criteria. Thus, the rule will dynamically apply to traffic as it comes from different interfaces, and as you change security group membership in ISE. If the packet does not explicitly contain a source SGT, source/destination matching will be based on the packet IP addresses compared to the SGT-to-IP address mappings obtained from user session information or from SXP-published mappings.

a) Choose **Policies** > **Access Control**.
b) Click **+** to create a new rule, or edit an existing rule.
c) Enter a rule name and select **Allow** as the action.
d) On the **Source/Destination** tab, click **+** under **Source** > **SGT Groups** and select the object you created for production users.
e) On the **Source/Destination** tab, click **+** under **Destination** > **SGT Groups**, and select the object you created for production servers.
f) Configure other options as needed. For example, you can enable logging and apply an intrusion policy.
g) Click **OK**.

Step 6 Deploy the configuration.



CHAPTER 22

Intrusion Policies

The following topics explain intrusion policies and the closely associated network analysis policies (NAP). Intrusion policies include rules that check traffic for threats and block traffic that appears to be an attack. Network analysis policies control traffic preprocessing, which prepares traffic to be further inspected by normalizing traffic and identifying protocol anomalies.

Because preprocessing and intrusion inspection are so closely related, the network analysis and intrusion policies examining a single packet must complement each other.

- About Intrusion and Network Analysis Policies, on page 509
- License Requirements for Intrusion Policies, on page 515
- Applying Intrusion Policies in Access Control Rules, on page 515
- Switching Between Snort 2 and Snort 3, on page 516
- Configuring Syslog for Intrusion Events, on page 517
- Configuring the Network Analysis Policy (Snort 3), on page 517
- Managing Intrusion Policies (Snort 3), on page 522
- Managing Intrusion Policies (Snort 2), on page 534
- Monitoring Intrusion Policies, on page 536
- Examples for Intrusion Policies, on page 537

About Intrusion and Network Analysis Policies

Network analysis and intrusion policies work together to detect and prevent intrusion threats.

- A network analysis policy (NAP) governs how traffic is decoded and preprocessed so that it can be further evaluated, especially for anomalous traffic that might signal an intrusion attempt.
- An intrusion policy uses intrusion and preprocessor rules, which are collectively known as intrusion rules, to examine the decoded packets for attacks based on patterns. The rules can either prevent (drop) the threatening traffic and generate an event, or simply detect (alert) it and generate an event only.

As the system analyzes traffic, the network analysis decoding and preprocessing phase occurs before and separately from the intrusion prevention 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.

System-Defined Network Analysis and Intrusion Policies

The system includes several pairs of same-named network analysis and intrusion policies that complement and work with each other. For example there are both NAP and intrusion policies named "Balanced Security and Connectivity," which are meant to be used together. The system-provided policies are configured by the Cisco Talos Intelligence Group (Talos). For these policies, Talos sets the intrusion and preprocessor rule states and provides the initial configurations for preprocessors and other advanced settings.

As new vulnerabilities become known, Talos releases intrusion 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 might also delete rules from system-provided policies and provide new rule categories, as well as modify the default variable set.

You can manually update the rules database, or configure a regular update schedule. You must deploy an update for it to take effect. For more information on updating system databases, see Updating System Databases, on page 788.

The following are the system-provided policies:

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 networks and deployment types. The system uses the Balanced Security and Connectivity network analysis policy as the default.

Connectivity Over Security network analysis and intrusion policies

These policies are built for networks where connectivity, the ability 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 networks 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 networks 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.

Inspection Mode: Prevention vs. Detection

By default, all intrusion policies operate in Prevention mode to implement an Intrusion Prevention System (IPS). In the Prevention inspection mode, if a connection matches an intrusion rule whose action is to drop traffic, the connection is actively blocked.

If you instead want to test the effect of the intrusion policy on your network, you can change the mode to Detection, which implements an Intrusion Detection System (IDS). In this inspection mode, drop rules are treated like alert rules, where you are notified of matching connections, but the action result becomes Would Have Blocked, and connections are never in fact blocked.

You change the inspection mode per intrusion policy, so you can have a mix of prevention and detection.

The Snort 3 network analysis policy (NAP) also has an inspection mode. Unlike the intrusion policy, the NAP policy is global, so you must run all NAP processing in either Prevention or Detection mode. You should use the same mode you use for your intrusion policies. If you have a mix of Prevention and Detection policies, select Prevention to match your most restrictive intrusion policies.

Intrusion and Preprocessor Rules

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.

The system includes the following types of rules created by Cisco Talos Intelligence Group (Talos):

- Intrusion rules, which are subdivided into shared object rules and standard text rules
- Preprocessor rules, which are rules associated with preprocessors and packet decoder detection options in the network analysis policy. Most preprocessor rules are disabled by default.

The following topics explain intrusion rules in more depth.

Intrusion Rule Attributes

When you view an intrusion policy, you see a list of all the intrusion rules available for identifying threats.

The list of rules for each policy is the same. The difference is in the action configured for each rule. Because there are over 30,000 rules, scrolling through the list will take time. Rules are revealed as you scroll through the list.

Following are the attributes that define each rule:

> (Signature Description)

Click the > button in the left column to open the signature description. The description is the actual code used by the Snort inspection engine to match traffic against the rule. Explaining the code is out of scope for this document, but it is explained in detail in *Management Center Configuration Guide*; select the book for your software version from http://www.cisco.com/c/en/us/support/security/defense-center/products-installation-and-configuration-guides-list.html. Look for information on intrusion rule editing.

The signatures contain variables for certain items. For more information, see Default Intrusion Variable Set, on page 512.

GID

Generator Identifier (ID). This number indicates which system component evaluates the rule and generates events. A 1 indicates a standard text intrusion rule, a 3 indicates a shared object intrusion rule. (The difference in these rule types is not meaningful for a Firewall Device Manager user.) These are the main rules of interest when configuring an intrusion policy. For information on the other GIDs, see Generator Identifiers, on page 513.

SID

Snort Identifier (ID), also called signature ID. Snort IDs lower than 1000000 were created by the Cisco Talos Intelligence Group (Talos).

Action

The state of this rule in the selected intrusion policy. For each rule, "(Default)" is added to the action that is the default action for the rule within this policy. To return a rule to its default setting, you select this action. Possible actions are:

- **Alert**—Create an event when this rule matches traffic, but do not drop the connection.
- **Drop**—Create an event when this rule matches traffic, and also drop the connection.
- **Disabled**—Do not match traffic against this rule. No events are generated.

Status

For Snort 2 rules, Status is a separate column. If you change the default action for a rule, this column shows "Overridden." Otherwise, the column is empty.

For Snort 3 rules, the "Overridden" status is shown at the bottom of the Action attribute, if you changed it.

Message

This is the name of the rule, which also appears in events triggered by the rule. The message typically identifies the threat that the signature matches. You can search the Internet for more information on each threat.

Default Intrusion Variable Set

The intrusion rule signatures contain variables for certain items. Following are the default values for the variables, with $HOME_NET and $EXTERNAL_NET being the most commonly used variables. Note that the protocol is specified separately from port numbers, so port variables are numbers only.

- $DNS_SERVERS = $HOME_NET (meaning any IP address).
- $EXTERNAL_NET = any IP address.
- $FILE_DATA_PORTS = $HTTP_PORTS, 143, 110.
- $FTP_PORTS = 21, 2100, 3535.
- $GTP_PORTS = 3386, 2123, 2152.
- $HOME_NET = any IP address.
- $HTTP_PORTS = 144 ports numbered: 36, 80-90, 311, 383, 443, 555, 591, 593, 631, 666, 801, 808, 818, 901, 972, 1158, 1212, 1220, 1414, 1422, 1533, 1741, 1830, 1942, 2231, 2301, 2381, 2578, 2809, 2980, 3029, 3037, 3057, 3128, 3443, 3507, 3702, 4000, 4343, 4848, 5000, 5117, 5222, 5250, 5450, 5600, 5814, 6080, 6173, 6767, 6988, 7000, 7001, 7005, 7071, 7080, 7144, 7145, 7510, 7770, 7777-7779, 8000, 8001, 8008, 8014, 8015, 8020, 8028, 8040, 8060, 8080-8082, 8085, 8088, 8118, 8123, 8161, 8180-8182, 8222, 8243, 8280, 8300, 8333, 8344, 8400, 8443, 8500, 8509, 8787, 8800, 8888, 8899, 8983, 9000, 9002, 9060, 9080, 9090, 9091, 9111, 9290, 9443, 9447, 9710, 9788, 9999, 10000, 11371, 12601, 13014, 15489, 19980, 23472, 29991, 33300, 34412, 34443, 34444, 40007, 41080, 44449, 50000, 50002, 51423, 53331, 55252, 55555, 56712.
- $HTTP_SERVERS = $HOME_NET (meaning any IP address).
- $ORACLE_PORTS = any
- $SHELLCODE_PORTS = 180.

- $SIP_PORTS = 5060, 5061, 5600
- $SIP_SERVERS = $HOME_NET (meaning any IP address).
- $SMTP_SERVERS = $HOME_NET (meaning any IP address).
- $SNMP_SERVERS = $HOME_NET (meaning any IP address).
- $SQL_SERVERS = $HOME_NET (meaning any IP address).
- $SSH_PORTS = 22.
- $SSH_SERVERS = $HOME_NET (meaning any IP address).
- $TELNET_SERVERS = $HOME_NET (meaning any IP address).

Generator Identifiers

The generator identifier (GID) identifies the subsystem that evaluates an intrusion rule and generates events. Standard text intrusion rules have a generator ID of 1, and shared object intrusion rules have a generator ID of 3. There are also several sets of rules for various preprocessors. The following table explains the GIDs.

Table 12: Generator IDs

ID	Component
1	Standard Text Rule.
2	Tagged Packets. (Rules for the Tag generator, which generates packets from a tagged session.)
3	Shared Object Rule.
102	HTTP Decoder.
105	Back Orifice Detector.
106	RPC Decoder.
116	Packet Decoder.
119, 120	HTTP Inspect Preprocessor. (GID 120 rules relate to server-specific HTTP traffic.)
122	Portscan Detector.
123	IP Defragmentor.
124	SMTP Decoder. (Exploits against SMTP verbs.)
125	FTP Decoder.
126	Telnet Decoder.

ID	Component
128	SSH Preprocessor.
129	Stream Preprocessor.
131	DNS Preprocessor.
133	DCE/RPC Preprocessor.
134	Rule Latency, Packet Latency. (Events for these rules are generated when rule latency suspends (SID 1) or re-enables (SID 2) a group of intrusion rules, or when the system stops inspecting a packet because the packet latency threshold is exceeded (SID 3).)
135	Rate-Based Attack Detector. (Excessive connections to hosts on the network.)
137	SSL Preprocessor.
138, 139	Sensitive Data Preprocessor.
140	SIP Preprocessor.
141	IMAP Preprocessor.
142	POP Preprocessor.
143	GTP Preprocessor.
144	Modbus Preprocessor.
145	DNP3 Preprocessor.

Network Analysis Policies

Network analysis policies control traffic preprocessing. Preprocessors prepare traffic to be further inspected by normalizing traffic and identifying protocol anomalies. Network analysis-related preprocessing occurs after Security Intelligence drops and SSL decryption, but before access control and intrusion or file inspection.

By default, the system uses the Balanced Security and Connectivity network analysis policy to preprocess all traffic handled by the access control policy. However, if you configure an intrusion policy on any access control rule, the system uses the network analysis policy that matches the most aggressive intrusion policy applied. For example, if you use both Security over Connectivity and Balanced policies in your access control rules, the system uses the Security over Connectivity NAP for all traffic. For Snort 3 custom intrusion policies, this assignment is done according to the base template policy assigned to the intrusion policy.

When using Snort 3, you can select a policy explicitly and optionally customize its settings. We recommend that you select the policy whose name matches the intrusion policy that you use for most traffic that goes through the device, whether you use the intrusion policy directly, or you use it as the base policy in your custom intrusion policies. You can then change the inspection mode, or adjust specific inspector or binder settings to account for the traffic on your network.

In addition, consider whether you have enabled preprocessor rules in the intrusion policy. If you enable rules that require a preprocessor, ensure that you also enable the corresponding inspector in the NAP. For each inspector, you can also adjust the attributes for the inspector, including the ports examined (the binders), to customize inspector behavior for your network.



Note If you are using Snort 2, the system uses the same-named NAP policy as the most restrictive intrusion policy you apply in any access control rule, and you cannot edit inspector or binder settings.

License Requirements for Intrusion Policies

You must enable the **Threat** license to apply intrusion policies in access control rules. For information on configuring licenses, see Enabling or Disabling Optional Licenses, on page 89.

No extra license is needed for network analysis policies.

Applying Intrusion Policies in Access Control Rules

To apply intrusion policies to network traffic, you select the policy within an access control rule that allows traffic. You do not directly assign intrusion policies.

You can assign different intrusion policies to provide variable intrusion protection based on the relative risks of the networks you are protecting. For example, you might use the more stringent Security over Connectivity policy for traffic between your inside network and external networks. On the other hand, you might apply the more lenient Connectivity over Security policy for traffic between inside networks.

You can also simplify your configuration by using the same policy for all networks. For example, the Balanced Security and Connectivity policy is design to provide good protection without excessively impacting connectivity.

Procedure

Step 1 Select **Policies** > **Access Control**.

Step 2 Either create a new rule, or edit an existing rule, that **allows** traffic.

If the default action is allow, you can also specify an intrusion policy in the default action.

You cannot apply intrusion policies to rules that trust or block traffic.

Step 3 Click the **Intrusion Policy** tab.

Step 4 Select **Intrusion Policy** > **On** and select the intrusion inspection policy to use on matching traffic.

Switching Between Snort 2 and Snort 3

Snort is the main inspection engine for the product. Although you can switch Snort versions freely, some intrusion rules in Snort 2.0 might not exist in Snort 3.0, and vice versa. If you changed the rule action for one of these rules, that change is not preserved if you switch to Snort 3 and then back to Snort 2, or back again to Snort 3. Your changes to rule actions for rules that exist in both versions are preserved. Note that the mapping between rules in Snort 3 vs. Snort 2 can be one-to-one or one-to-many, so preservation of changes is done on a best effort basis.

If you change the Snort version, the system will perform an automatic deployment to implement the change. You can see the progress in the Task list. The tasks are Snort Version Change and Automatic Deployment—Snort version toggle. Because of the deployment and the fact that Snort must be stopped and restarted, all existing connections, including VPN, are dropped and must be reestablished, which will cause a momentary traffic loss.



Note If you try to switch Snort versions and the switch fails, you will have pending changes that you cannot discard, and a subsequent toggle attempt will not be allowed. If this happens, you must complete the switch using the ToggleInspectionEngine API, which you can use from the API Explorer. You must set the bypassPendingChangeValidation attribute to TRUE.

Before you begin

To determine which Snort version is currently enabled, use this procedure, or choose **Policies** > **Intrusion**. Look for the **Snort Version** line above the table. The current version is the first number in the complete version number. For example, 2.9.17-95 is a Snort 2 version.

If the device is in an air-gapped network, consider manually uploading the latest rules package for the new version before switching.

If you downgrade to 2.0, any custom intrusion policies that you created are converted to the base policy used in the custom policy. As far as possible, rule action overrides are retained. If more than one custom policy uses the same base policy, the overrides of the custom policy that is used in the most access control policies are retained, and the overrides for the other custom policies are lost. Access control rules that used these "duplicate" policies will now use the base policy created from your most-used custom policy. All custom policies are deleted. If you want to preserve the custom policies so that you can import them later, after switching back to Snort 3, use the Firewall Threat Defense API to export the configuration.

In addition, a downgrade to 2.0 removes any NAP customizations, and the system switches to use the most appropriate NAP based on the intrusion policies used in access control rules.

Hostname redirect in active authentication also requires Snort 3, and will be removed if you switch to Snort 2.

You must deploy any pending changes before you can switch Snort versions.

Procedure

Step 1 Choose **Device**, then click **View Configuration** in the Updates summary.

Look at the **Intrusion Rule** group. The current Snort version is indicated.

Step 2 In the **Intrusion Rule** group, you can change the Snort version by clicking **Upgrade to Snort 3.0** or **Downgrade to Snort 2.0**.

Step 3 When prompted to confirm your action, select the option to get the latest intrusion rules package, then click **Yes**.

We recommend that you get the latest rules package. The system downloads packages for the active Snort version only, so it is unlikely that you have the latest package installed for the Snort version you are switching to.

You must wait until the task to switch versions completes before you can edit intrusion policies.

Configuring Syslog for Intrusion Events

You can configure an external syslog server for intrusion policies to send intrusion events to your syslog server. You must configure the syslog server on the intrusion policy to get intrusion events sent to the server. Configuring a syslog server on an access rule sends connection events only to the syslog server, not intrusion events.

If you select multiple syslog servers, events are sent to each of the servers.

Intrusion events have the message ID 430001.

Procedure

Step 1 Select **Policies** > **Intrusion**.

Step 2 Click the **Intrusion Policy Settings** button ( ) to configure syslog.

Step 3 Click the **+** button under **Send Intrusion Events To** and select the server objects that define the syslog servers. If the required objects do not already exist, click **Create New Syslog Server** and create them.

Step 4 Click **OK**.

Configuring the Network Analysis Policy (Snort 3)

The Network Analysis Policy (NAP) is applied to all allowed connections on the device. The NAP determines which inspectors are enabled, and the values of the attributes used by the inspectors. The binders determine the ports and protocols that should be associated with the various inspectors.

Coordinate the NAP with the intrusion policies you apply in access control rules:

- If you use a single intrusion policy in your access control rules, select the same-named NAP. Then, make adjustments to inspectors and attributes based on the settings in your intrusion policy. For example, if you enable intrusion rules for a particular inspector, such as CIP, ensure that you enable that inspector in the NAP.

- If you use multiple intrusion policies, select the NAP that matches the most strict intrusion policy that you use.
- If you use custom intrusion policies, make your NAP selection based on the base intrusion policy for your custom intrusion policies.
- If you do not need to customize any inspectors or binders, consider configuring the system to automatically select the most appropriate NAP based on your intrusion policy usage. This is the default option.

Before you begin

Unless you prevent it, the system regularly downloads LSP updates to the inspection rules. These updates can add or remove inspectors and attributes, and change default settings for attributes. If you have made overrides to removed inspectors, these overrides are preserved and you will see warnings that the inspector is no longer supported. In this case, delete the inspector and make any other flagged adjustments to ensure your NAP is completely valid.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Verify that the Snort version shown above the table is 3.x.

Step 2 Click the **Intrusion Policy Settings** button ( ).

Step 3 In **Default Network Analysis Policy**, select one of the following:

- **Auto**—Automatically select the NAP that matches the most-used intrusion policy (or base policy for custom rules) applied in access control rules. If you do not apply any intrusion policies, the Balanced Security and Connectivity NAP is used. The NAP is run in Prevention mode, and you cannot customize intrusion or binder settings. The remainder of this procedure does not apply when running in automatic mode.
- **Custom**—Explicitly select the NAP that should be used. Click the **Edit** link next to the policy name to select a different policy. You can then select the inspection mode, and customize inspector and binder settings, as described in the following steps.

Step 4 In the Edit Network Analysis Policy dialog box, select the policy and configure its settings.

a) In **Network Analysis Policy**, select the policy that should apply globally to all allowed connections.

b) Select the **Inspection Mode**.

The inspection mode determines how non-compliant traffic is handled. Use the same inspection mode as you use in your intrusion policies for optimal results.

- **Prevention**—Block any decoder, normalization, or protocol anomalies based on the settings in the policy. You must use this option if you enable the SSL Decryption policy, or if you enable the **TLS Server Identity Discovery** option in the access control policy settings.
- **Detection**—Simply issue alerts for decoder, normalization, or protocol anomalies. Do not block any traffic.

c) (Optional.) Configure and manage overrides to the inspectors and binders:

- To edit overrides, see Configuring Inspector and Binder Overrides, on page 519.
- To download the schema or overrides, see Downloading Overrides and the Schema, on page 521.
- To upload overrides, see Uploading Overrides, on page 521.
- To reset all the overrides, click the **Reset Inspector/Binder Overrides** link above the NAP file. You are asked to confirm the reset. The deletion is limited to inspectors or binders, as indicated in the command name. For example, deleting all binder overrides leaves your inspector overrides unchanged.
- To reverse all changes to the selected inspector, click **Reset *Inspector* to Defaults**.
- To filter the view so you see only those inspectors that have overrides, click **Show Only Overrides**. Click **Show All Inspectors** to remove the filter.

d) Click **OK**.

Configuring Inspector and Binder Overrides

When you select a base NAP, you are selecting inspector settings contained in that baseline policy. In most cases, these are the appropriate settings.

However, you can override the settings in the selected NAP. For example, you can enable or disable individual inspectors, or change the value for an attribute or a binder.

The following procedure explains how to configure overrides directly. Alternatively, you can download the schema, make changes offline, then upload your overrides. You can also upload the overrides you downloaded from another device.

Before you begin

Explaining each inspector, binder, and attribute is beyond the scope of this document. For detailed information, including examples, see *Snort 3 Inspector Reference* at https://www.cisco.com/c/en/us/td/docs/security/secure-firewall/snort3-inspectors/snort-3-inspector-reference.html.

Procedure

Step 1 Choose **Policies** > **Intrusion**, click the **Intrusion Policy Settings** button ( ), select **Custom** for the NAP settings, then click the **Edit** link next to the policy name.

Step 2 Click the tab that contains the setting you want to change:

- **Inspectors**—The inspectors examine specific types of traffic, such as FTP, for protocol anomalies.
- **Binders**—The Binder inspector determines when a service inspector needs to be used to inspect traffic. The configurations in the binder inspector include the ports, hosts, CIDRs, and services that define when another inspector in the network analysis policy needs to inspect traffic.

Step 3 Edit the settings as needed.

- Use the following to control the view in the JSON editor:

- Use the **Filter** edit box to do a full-text search of the JSON file.
- Click the **Expand All Fields** button ( ) to open all of the folders in the JSON file.
- Click the **Collapse All Fields** ( ) button to close all folders in the JSON file.
- Click the **Undo Last Action** ( ) button to reverse your most recent change.
- Click the **Redo** ( ) button to redo your last reversed change.
- Select **Tree** to see a formatted view of the JSON file, which includes action menus, error flags, and other features to guide your edits.
- Select **Code** to see the raw JSON file.

- In Tree view, click the **Menu** ( ) button to manipulate the contents of the file. You can:
  - **Insert** attributes. Use Auto to have the editor determine the appropriate data type. Otherwise, add an Array, Object, or String. If you add an invalid attribute, the system will mark the inspector or binder as having a problem that you must resolve.
  - **Append** attributes. This action does the same as Insert, but puts the attribute at the end of the section.
  - **Duplicate** the selected attribute.
  - **Remove** (delete) the selected attribute. When editing an attribute, a popup message might also provide a **Delete** command.

- To enable an inspector that is currently disabled, or change the setting of any Boolean attribute, click the checkbox in front of the attribute value. For example, to enable an inspector, change the **enabled : false** attribute to:



- To change the value of a string or numeric attribute, click in the attribute and edit the value as needed. If your entry violates the rules for the field, error messages explain the discrepancy. For example, a numeric value indicates the valid range of values if you enter a value outside the range.
- To reset overrides:
  - Click **Reset Inspector/Binder Overrides** to remove all of your changes to all inspectors or binders and return to default values. The deletion is limited to inspectors or binders, as indicated in the command name. For example, deleting all binder overrides leaves your inspector overrides unchanged.
  - Click **Reset *Inspector* to Defaults** to reverse all changes to the selected inspector only.
- To filter the view so you see only those inspectors that have overrides, click **Show Only Overrides**. Click **Show All Inspectors** to remove the filter.
- If an inspector is no longer supported, the inspector is flagged with a message. Click the **Delete Inspector** link in the message to remove the inspector.

Step 4 Click **OK** when finished.

Downloading Overrides and the Schema

You can download the NAP schema, or download the overrides you have configured for the policy.

Downloading overrides is recommended whenever you change the base NAP, in case you want to go back to your previous settings. In addition, you could use the JSON editor on one device to implement the overrides you want to use on all devices, download the overrides, then upload that override file to other devices.

Downloading the schema is useful if you want to edit the file offline, then upload your overrides to this device or to multiple devices. You should copy/paste just the sections you need to change, rather than upload the entire file, to ensure only those changes you make are considered overrides.

Procedure

Step 1 Choose **Policies** > **Intrusion**, click the **Intrusion Policy Settings** button ( ), select **Custom** for the NAP settings, then click the **Edit** link next to the policy name.

Step 2 Do one of the following:

- To download the schema for the currently-select NAP, click the gear icon ( ) and select **Download** > **Policy Schema**.
- To download the saved set of overrides, as they existed prior to the current editing session, click the gear icon ( ) and select **Download** > **Last Saved Overrides**. The file includes overridden attributes plus their containing objects.
- To download the overrides you have created in the current editing session, click the gear icon ( ) and select **Download** > **Current Unsaved Overrides**. The file includes overridden attributes plus their containing objects.

Uploading Overrides

Rather than editing attributes using the embedded JSON editor, you can download the NAP policy schema, edit the file offline, then upload the file. Any overrides configured in the uploaded file are then applied to the selected NAP.

You can also upload a file you downloaded after configuring overrides on another device.

By uploading your overrides, you can upload the same file to multiple devices and easily apply the same overrides.

Before you begin

For overriding an inspector configuration in the network analysis policy, you should upload only the changes that you require. You should not upload the entire configuration because it makes the overrides sticky in nature and therefore, any subsequent changes to the default values or configuration as part of the LSP updates would not be applied. Ensure that your uploaded overrides are tightly focused on just those attributes you want to change.

Procedure

Step 1 Choose **Policies** > **Intrusion**, click the **Intrusion Policy Settings** button ( ), select **Custom** for the NAP settings, then click the **Edit** link next to the policy name.

Step 2 Click the gear icon ( ) and select **Upload** > **Overrides**.

Step 3 (Optional.) Click one of the **Download** links to save a copy of your existing overrides.

You can either download the last saved overrides (those made before the current editing session) or the current unsaved overrides (those made during the current editing session).

Step 4 Click **Yes** on the Confirm Upload Overrides dialog box to confirm you want to continue.

Step 5 Click **Browse**, or drag and drop, to select the JSON file that contains your overrides, and click **OK**.

Managing Intrusion Policies (Snort 3)

When you use Snort 3 as the inspection engine, you can create your own intrusion policies and customize them for your purposes. The system comes with pre-defined policies that are based on the same-named Cisco Talos Intelligence Group (Talos)-defined policies. Although you can edit these policies, a better practice is to create your own policy based on the underlying Talos policy and change that if you need to adjust rule actions.

Each of these pre-defined policies includes the same list of intrusion rules (also known as signatures), but they differ in the actions taken for each rule. For example, a rule might be enabled in one policy, but disabled in another policy.

If you find that a particular rule is giving you too many false positives, where the rule is blocking traffic that you do not want blocked, you can disable the rule without needing to switch to a less-secure intrusion policy. You could alternatively change it to alert on matches without dropping traffic.

Conversely, if you know you need to protect against a specific attack, but the related rule is disabled in your chosen intrusion policy, you can enable the rule without changing to a more secure policy.

Use the intrusion related dashboards, and the Event Viewer (both on the **Monitoring** page), to evaluate how intrusion rules are impacting traffic. Keep in mind that you will see intrusion events and intrusion data only for traffic that matches intrusion rules set to alert or drop; disabled rules are not evaluated.

Note If you switch to Snort 2, you cannot create custom policies, and the use of intrusion policies differs slightly. Instead of this topic, see Managing Intrusion Policies (Snort 2), on page 534.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Verify that the Snort version shown above the table is 3.x.

Step 2 Do any of the following:

- Use the **Search/Filter** box to find a policy. You can search by name only.
- Click the gear icon ( ) to enable logging to a syslog server. See Configuring Syslog for Intrusion Events, on page 517.
- Click the gear icon ( ) to configure the network analysis policy (NAP). See Configuring the Network Analysis Policy (Snort 3), on page 517.
- Click + to create a new policy. See Configuring a Custom Intrusion Policy (Snort 3), on page 523.
- Click the edit icon ( ) to see the properties and rules in the policy, and to edit them. See Viewing or Editing Intrusion Policy Properties (Snort 3), on page 524.
- Click the delete icon ( ) to delete a policy.

Configuring a Custom Intrusion Policy (Snort 3)

You can create new intrusion policies to customize rule behavior if the pre-defined policies do not fit your needs. In general, it is good practice to create custom policies based on the pre-defined policies rather than to alter those policies. This ensures you can easily implement one of the Cisco Talos defined policies if you find that your customizations do not deliver the results that you need.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Do one of the following:

- To create a new policy, click +.
- To edit an existing policy, click the edit icon ( ) for the policy. When you are shown the policy details, click the **Edit** link in the policy properties section at the top of the page.

Step 3 Enter a **Name**, and optionally, a description for the policy.

Step 4 Configure the **Inspection Mode** for the policy.

- **Prevention**—Intrusion rule actions are always applied. Connections that match a drop rule are blocked.
- **Detection**—Intrusion rules generate alerts only. A connection that matches a drop rule will generate alert messages, but the connection will not be blocked.

Step 5 Select the **Base Template** for the policy.

The base templates are provided by Cisco Talos. Click the info icon for each to see more information about the policies. Note that policy names can change, and new policies appear, when a new rules package gets installed.

- **Maximum Detection (Cisco Talos)**—This policy places all emphasis on security. Network connectivity and throughput is not guaranteed and false positives are likely. This policy should only be used for high security areas and security monitors must be prepared to investigate alerts to determine their validity.
- **Security Over Connectivity (Cisco Talos)**—This policy places an emphasis on security, at the possible expense of network connectivity and throughput. Traffic is inspected more deeply, more rules are evaluated, and both false positives and increased latency are expected but within reason.
- **Balanced Security and Connectivity (Cisco Talos)**—(Default.) This policy attempts to strike the delicate balance between network connectivity and throughput and the needs of security. While not as strict as Security Over Connectivity, this policy attempts to keep users secure while being less obtrusive about normal traffic.
- **Connectivity Over Security (Cisco Talos)**—This policy places an emphasis on network connectivity and throughput, at the possible expense of security. Traffic is inspected less deeply, and fewer rules are evaluated.
- **No Rules Active (Cisco Talos)**—This policy is a basic policy that configures typical preprocessor settings but does not have any rules or built-in alerts enabled. Use this policy as a base if you want to ensure that only those policies you want to apply are enabled.

Step 6 Click **OK**.

You are returned to the intrusion policy list. You can now view the new policy and adjust rule actions as needed.

Viewing or Editing Intrusion Policy Properties (Snort 3)

The Intrusion Policy page shows a list of the policies, including both pre-defined and user-defined policies, and their descriptions. To edit a policy, you must first view the policy's properties.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

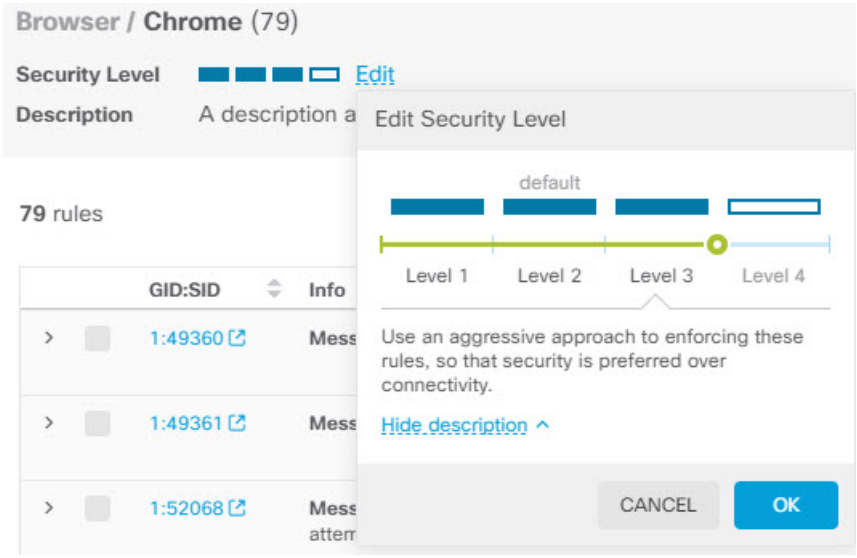
Step 2 Click the edit icon ( ) for a policy.

The policy contains the following sections:

- **Policy Name** drop-down list.
  - You can easily switch to a different policy by selecting it from the drop-down list, or return to the list of policies by clicking the back button ( ).
  - You can delete this policy by clicking the delete icon next to the policy name ( ).
- **General properties**. This section shows the intrusion mode, base policy, and description. Click **Edit** to change these properties or the policy name.
- **Rule Group** table of contents. This list displays all the rule groups that have active rules in the policy. The groups have a hierarchy, with parent groups containing child groups that organize subsets of rules

within the larger parent group. Each group is a logical collection of rules, and a given rule can appear in more than one group.

- To add a group that currently has no active rules in the policy, click + > **Add Existing Rule Group** and select the group. See Adding or Removing Rule Groups in an Intrusion Policy (Snort 3), on page 526.
- To change the security level of a group, select the child group in the list. The rule list changes to show the security level at the top, with the rules in the group listed below. Click the **Edit** link next to the security level and select a new level. Click **View Description** when editing to get information about each security level. Note that changing the level can change which rules are active, and also the action for a given rule, with more secure levels tending to have more active rules and more rules with the Drop action. Click **OK** to confirm the change. (Security Level does not apply to custom rule groups.)



- To remove all the rules in a group, select the child group in the list. Then, click the **Exclude** link to the far right of the group name and confirm that you want to exclude the group. Excluding the group simply disables all the rules in the group. It does not delete the group.

  However, if the group includes rules that are shared with other groups that are enabled, the shared rules retain whatever actions are applied by the still-active group. In all cases, we retain your most aggressive setting for an individual rule, regardless of group membership.
- To add a new custom rule group of custom rules, click + > **Upload Custom Rules**. For details, see Uploading Custom Intrusion Rules, on page 531.
- To change the name or description of a custom rule group, click **Edit**.
- To delete a custom rule group, click **Delete**. For more information, see Managing Custom Intrusion Rules and Rule Groups, on page 530.
- To add a new custom rule in a custom rule group, click + above the rule table. See Configuring Individual Custom Intrusion Rules, on page 533.
- To edit, duplicate, delete, or manage group membership for a custom rule, mouse over the right of the rule and click the appropriate button or command. For more information, see Configuring Individual Custom Intrusion Rules, on page 533.

- **List of rules**. You can use the search field to help you find rules using full-text search. You can also select filtering items to search on any combination of GID or SID, show only user defined rules (those you added), or simply show rules based on their actions (disabled, alert, drop). The rules are lazy-loaded, so it will take quite a bit of time to scroll through the entire unfiltered list. When filtering the list, click the refresh button to reload the filtered view.
  - To change the action for a rule, click the **Action** cell for the rule and select the new action, to **Alert** only, to **Block** matching traffic, or to **Disable** the rule. The default action for each rule is indicated.
  - To change the action for more than one rule at a time, click the check box in the left column of the rules you want to change, then select the new action from the **Action** drop-down list above the rules table. Click the check box in the GID:SID header to select all the rules in the list. You can change up to 5000 rules at a time.
  - To update the rules within a custom rule group, click **Upload Rule File**. For more information, see Uploading Custom Intrusion Rules, on page 531.
  - To get more information about a rule, click the link in the **GID:SID** cell. The link takes you to Snort.org.
  - To change the rules listed, you can click a child group from the rule group table of contents (not a parent group). You can return to the all-rules list by clicking **ALL RULES** at the top of the rule group list.
  - To change the sorting order, click the table header for a column. The default sorting for the rules is overridden rules first, then drop rules, then alert rules.
  - To see what changes where made in an intrusion rule (LSP) update, select **LSP Update** in the filter field, then select the updates whose changes you want to see, and specify whether you want to see all changes, or just additions or changes to the rules.

Adding or Removing Rule Groups in an Intrusion Policy (Snort 3)

Intrusion rules are organized in local groups. There is a hierarchy to the groups, with parent groups containing related child groups. The rules themselves appear in child groups only: parent groups are simply an organizational construct. A given rule can appear in more than one group.

Any custom rule groups you create are in the User Defined Groups folder. Custom rule groups do not have a hierarchy.

The easiest way to add or remove rules in an intrusion policy is to add or remove groups. Because the rules in a group are logically related, it is highly likely that you will want to use most if not all rules within a given group.

The following procedure explains how to add groups and change the group's security level.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Click the edit icon ( ) for the policy you want to change.

Step 3 (Adding groups.) If the group is not shown in the list of rule groups, click + > **Add Existing Rule Group** and do the following:

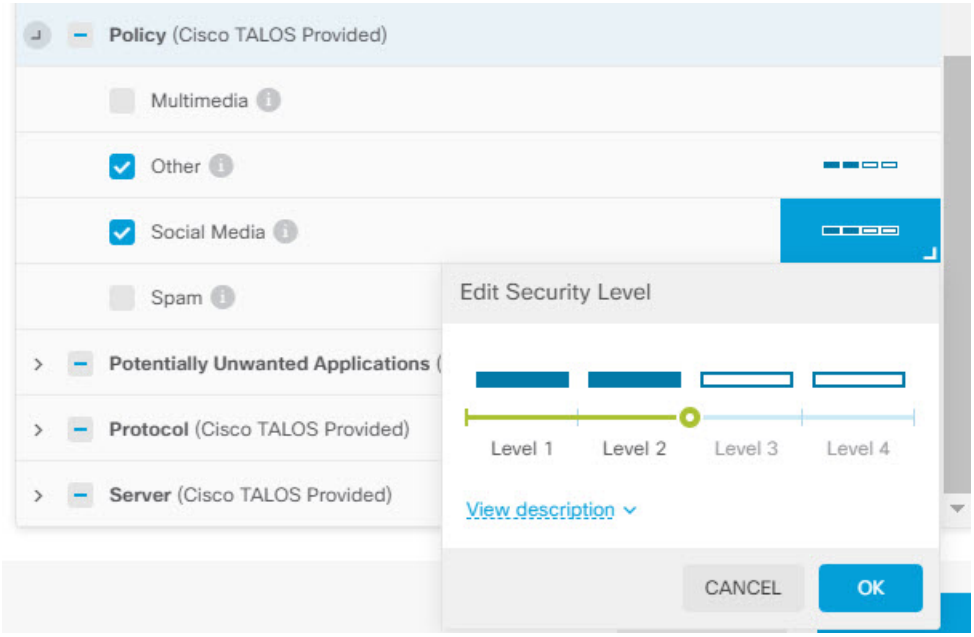
a) Find the child group.

- A check mark next to a parent group name indicates that all child groups in the parent group are already selected.
- A minus mark next to a parent group name indicates that one or more child group has no enabled rules for this policy. These are the groups you can add.
- A check mark next to a child group name indicates the group is already selected.

b) Select the group you want to add (that is, check its check box).

c) (Optional, does not apply to custom rule groups.) Each group has a default security level depending on the base policy used for the custom policy. If you want to change it, click the security level icon, select a new level, and click **OK**.

Level 1 is the least secure posture, emphasizing connectivity over security, whereas level 4 is the most aggressive, providing maximum security. You can click **View Description** to see an explanation of each level as you select it.



d) Continue selecting (or deselecting) groups until you have made all of your changes.

e) Click **OK**.

Step 4 (Removing groups.) If you want to disable all the rules within a group, you can use any of the following methods:

- Select the group, then click the **Exclude** link to the far right of the group name, above the list of rules.
- Use the method for adding a group, but instead, deselect the unwanted group (that is, uncheck its check box), and click **OK**.

- You can delete a custom rule group to remove it entirely from the system and from all intrusion policies that use it. Select the group, then click **Delete**.

Changing Intrusion Rule Actions (Snort 3)

Each intrusion policy has the same rules. The difference is the action taken for each rule can be different from policy to policy.

By changing the rule action, you can disable rules that are giving you too many false positives, or you can change whether the rule alerts on or drops matching traffic. You can also enable disabled rules to alert or drop matching traffic.

The easiest method to change rule actions is to change the security level of a rule group. When you change a group's security level, the action of the rules within the group change. This can mean that some rules become enabled (or disabled), or the action can change between alert and drop, based on the security posture you select. However, you can change an individual rule action if that is what you need.



Note The default action for a given rule is based on the overall selection of group and severity. Changing the severity of the group, or excluding the group, can change the default action for the rule.

Before you begin

Custom rule groups do not have security levels. You cannot use the security level technique to change rule actions for custom rules.

Procedure

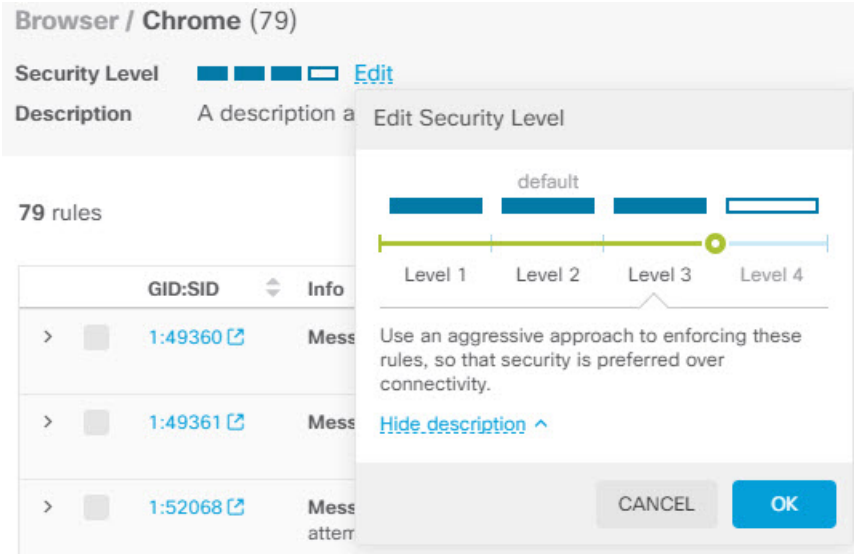
Step 1 Choose **Policies** > **Intrusion**.

Step 2 Click the view icon () for the policy whose rule actions you want to change.

Step 3 (Recommended method.) Change the security level for a rule group.

a) Click the child rule group in the rule group list.

b) Above the list of rules, click **Edit** next to the security level for the group.



Note
If you want to disable all the rules in the group, do not click **Edit**. Instead, click **Exclude** and confirm that you want to exclude the group. The group is not deleted, its rules are simply disabled. Skip the remaining steps.

c) Select the new level for the group. Click **View Description** to see an explanation of each level as you pick it.

Level 1 is the least secure posture, emphasizing connectivity over security, whereas level 4 is the most aggressive, providing maximum security.

d) Click **OK**.

Step 4 (Manual method.) Change the action for one or more rule.

a) Find the rule whose action you want to change.

Use the **Search/Filter** box to search on strings within the rule information. You can also select filtering items to search on any combination of GID or SID, or simply show rules based on their actions (disabled, alert, drop). The rules are lazy-loaded, so it will take quite a bit of time to scroll through the entire unfiltered list. When filtering the list, click the refresh button to reload the filtered view.

Ideally, you can get the Snort identifier (SID) and generator identifier (GID) from an event or from Cisco Technical Support, if you are working with them on an issue. You can then search precisely for the rule.

b) To change the action, do one of the following:

- Change one rule at a time—Click the **Action** column for the rule and select the required action:
  - **Alert**—Create an event when this rule matches traffic, but do not drop the connection.
  - **Drop**—Create an event when this rule matches traffic, and also drop the connection.
  - **Disabled**—Do not match traffic against this rule. No events are generated.

- Change multiple rules at one time—Click the check boxes for the rules you want to change, then click the **Bulk** drop-down above the table and pick the desired action. Click the check box in the GID:SID header to select all the rules in the list. You can change up to 5000 rules at a time.

Managing Custom Intrusion Rules and Rule Groups

The system comes with thousands of intrusion rules defined by Cisco Talos Intelligence Group (Talos). If you know of additional attacks, you can create and upload custom intrusion rules to screen for those attacks, and either alert or drop them. You can also create, edit, and delete rules one at a time.

For uploaded rules, you create the rules offline using a text editor. We recommend that you include a group of custom rules in each text file you upload. You can then easily upload changes to your rules, and either merge new rules into your custom rule groups, or replace your rules with new, edited copies.

Explaining how to create these rules is outside the scope of this document. For detailed information on how to write intrusion rules for Snort, including how to convert Snort 2 rules to Snort 3 format, see the guides on https://snort.org/documents. For example, *Rules Authors Introduction to Writing Snort 3 Rules* at https://snort.org/documents/rules-writers-guide-to-snort-3-rules.

Before you begin

You create custom rule groups during the process of uploading custom rules, as described in Uploading Custom Intrusion Rules, on page 531, or when creating individual rules or managing rule membership. After you create the group, you can manage the group and its contents.

Note that custom groups are available for all intrusion policies, not just the policy you were editing when you created the group. Thus, changes you make to a group are made for all policies. For example, if you delete a custom rule group, it is deleted from all policies and is no longer available for any of them.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Click the edit icon ( ) for a policy.

We recommend that you add custom rules to a custom intrusion policy rather than one of the built-in policies.

Step 3 Do any of the following:

- To create a group, click + > **Upload Custom Rules**. See Uploading Custom Intrusion Rules, on page 531.
- To edit a group's name or description, select the group in the group table of contents in the User Defined Groups folder. You can then click **Edit** and make your changes.
- To exclude the group and its rules from the policy, select the group in the group table of contents in the User Defined Groups folder. You can then click **Exclude** to remove the group.
- To delete the group from the system and all policies that use it, select the group in the group table of contents in the User Defined Groups folder. Then click **Delete**. Note that if a rule exists only in the deleted

group, it is also deleted from the system. However, if a rule also exists in other custom rule groups that you are not deleting, the rule will remain in those groups.

- To replace or update the rules in a group in bulk, select the group in the group table of contents in the User Defined Groups folder. Then, click **Upload Rule File** next to the Action drop-down list above the group's rule table. The process is the same as that described in Uploading Custom Intrusion Rules, on page 531.
- To create and manage individual rules, and their assignment to rule groups, see Configuring Individual Custom Intrusion Rules, on page 533.

Uploading Custom Intrusion Rules

If you know of attacks that are not currently covered by other rules, you can create and upload custom intrusion rules to screen for those attacks, and either alert or drop them. The action of the imported rules must be either alert or drop, and the default action for the rule is defined by the action in the imported file. Once imported, you can change a rule action and disable a rule if necessary.

You must create these rules offline. In the Firewall Device Manager, you are simply uploading a rules file, you are not configuring the rules directly. The rules file should be a text file. You can use line returns to format the rules to be readable, or put a rule on a single line, and empty lines are allowed. Rule format is explained at snort.org.

For example, an upload file of three rules could look like the following:

```
alert tcp $HOME_NET any -> $EXTERNAL_NET $HTTP_PORTS (
   msg: "My Custom Rule: EXPLOIT-KIT Styx exploit kit landing page request";
   flow:to_server,established;
   http_raw_uri;
   bufferlen:>100;
   http_uri;
   content:"/i.html?",depth 8; pcre:"/\/i\.html\?[a-z0-9]+\=[a-zA-Z0-9]{25}/";
   flowbits:set,styx_landing;
   metadata: copied from talos sid 29452;
   service:http;
   classtype:trojan-activity;
   gid:1;
   sid:1000000;
   rev:1;
)

alert tcp $HOME_NET 8811 -> $EXTERNAL_NET any (
   msg:"My Custom rule: MALWARE-BACKDOOR fear1.5/aciddrop1.0 runtime detection - initial
connection";
   flow:to_client,established;
   flowbits:isset,Fear15_conn.2;
   content:"Drive",nocase;
   metadata:copied from talos sid 7710;
   classtype:trojan-activity;
   gid:1;
   sid:1000001;
   rev:1;
)

alert tcp $EXTERNAL_NET $FILE_DATA_PORTS -> $HOME_NET any (
   msg:"My Custom Rule: INDICATOR-COMPROMISE download of a Office document with embedded
PowerShell";
```

```
    flow:to_client,established;
    flowbits:isset,file.doc;
    file_data;
    content:"powershell.exe",fast_pattern,nocase;
    metadata:copied from talos sid 37244;
    classtype:trojan-activity;
    gid:1;
    sid:1000002;
    rev:1;
)
```

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Click the edit icon ( ) for a policy.

We recommend that you add custom rules to a custom intrusion policy rather than one of the built-in policies.

Step 3 Do one of the following:

- Above the list of groups, click + > **Upload Custom Rules**.
- If you are uploading rules into a custom rule group that you already created, you can select the custom rule group, then click **Upload Rule File** next to the **Action** drop-down list above the group's rule table.

Step 4 Click **Browse** and select your custom rule file, or drag and drop the file into the Upload File dialog box.

Wait for the upload to complete.

Step 5 Select how you want to handle conflicts:

A conflict occurs when a rule you are adding is the same as one already in the system. This should happen only if you are uploading the same rules or edited versions of rules that you previously uploaded.

Select one of these options:

Note
Merge and **Replace** are basically the same thing. Uploaded rules must have higher revision numbers than the ones you already uploaded for any changes to be made to the existing rules. The only difference is that if the upload file is missing rules that are in the targeted custom rule group, the **Replace** option will delete those rules from the rule group. The **Merge** option will leave those "missing" rules in place.

- **Merge**—Any changed rules in the uploaded file that also exist in the selected group will have those changes merged, if the rule in the uploaded file has a higher revision number. Any unchanged rules, or rules in the group that do not have corresponding rules in the upload, will be left unchanged. Any new rules in the upload will be added. This is the default option.
- **Replace**—The rules in the uploaded file will replace the rules in the selected group, if the revision number on the uploaded rule is higher. Any existing rules not in the uploaded file will be deleted from the group. Existing rules whose uploaded version has a revision number that is the same or lower will be left unchanged. Any new rules in the upload will be added.

Step 6 Click + and select the custom rule group for the uploaded rules.

If the custom rule group you want to use does not yet exist, click **Create New Group** and create it now. The new group needs a name and optionally, a description. Then, you can select the new group.

If you are replacing rules, you can select a single group only. If you are merging them, you can select multiple groups.

Step 7 Click **OK**.

The files are uploaded and placed in the new group. You should see a summary of how many rules were uploaded, and how many rules were updated, deleted, or ignored.

If there are errors in the file, the upload will fail. You can click the **Download Error File** link to get more information on the errors.

The group is automatically activated in this intrusion policy. The group and new rules are available to be added to other policies, but the group and rules are not automatically enabled in any other policy. For information on adding groups to other policies, see Adding or Removing Rule Groups in an Intrusion Policy (Snort 3), on page 526.

Configuring Individual Custom Intrusion Rules

You can configure custom intrusion rules one at a time rather than in bulk through file uploads. This method works well when you need to make a quick adjustment to a rule, or you need to create or modify just a few rules at a time.

When configuring intrusion rules, keep the following in mind:

- The GID for all custom rules should be 1.
- The SID for a rule must be unique across all rules in the system. It also must be one million (1000000) or higher.
- If you edit a rule, you must change the rule version. Normally, you increment the version number by 1.
- You can duplicate a Cisco Talos Intelligence Group (Talos) rule to create your own version of the rule, but you still must change the SID of the duplicate to make it unique.

The system will do some validity checking to ensure the rule is well formed, and you will see error messages about any problems. However, the system cannot determine whether the rule is sensible.

For detailed information on how to write intrusion rules for Snort, including how to convert Snort 2 rules to Snort 3 format, see the guides on https://snort.org/documents. For example, *Rules Authors Introduction to Writing Snort 3 Rules* at https://snort.org/documents/rules-writers-guide-to-snort-3-rules.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Click the edit icon () for a policy.

We recommend that you add custom rules to a custom intrusion policy rather than one of the built-in policies.

Step 3 Do one of the following:

- To add an intrusion rule, click the **Add New Intrusion Rule** button (+) above the rule table. When adding a rule, you must select one or more custom rule groups to contain the new rule. You can create new groups while adding the rule if necessary.
- To add a rule by duplicating and editing an existing rule, mouse over the right end of the rule and click the Duplicate ( ) button. The button appears only on mouse-over. For custom rules, the **Duplicate** command is under the more options (**...**) button.
- To edit a custom rule, find the rule in a custom rule group and click the edit ( ) button for the rule. Your edits apply across all groups in which the rule resides. Make sure you increment the rule version number by at least 1 when making changes.
- To delete a custom rule, click the delete ( ) button for the rule. The rule is deleted from all rule groups that contain it. If you just want to remove a rule from a group, use the **Manage Group Assignments** option instead of deleting the rule.
- To change which groups contain the rule, click the more options (**...**) button and select **Manage Group Assignments**. You can then add or remove groups. Your changes simply affect group membership, it does not change the rule or delete it.

Step 4 For new rules and groups, add the rule to the policy.

When you create a new group when creating a new rule, or editing an existing rule, that group is not added to your policy automatically, nor is the rule enabled automatically. You are prompted to add the group to the policy you are editing. If you do not add the group while adding or editing the rule, you can add the group later using the following process:

a) Click **+** > **Add Existing Rule Group** above the group table of contents.
b) Find the group under the User Defined Groups folder, select it, and click **OK**.
c) Select the group in the table of contents and verify that the new rule is in the group and has the desired action.

Managing Intrusion Policies (Snort 2)

You can apply any of the pre-defined intrusion policies. Each of these policies includes the same list of intrusion rules (also known as signatures), but they differ in the actions taken for each rule. For example, a rule might be active in one policy, but disabled in another policy.

If you find that a particular rule is giving you too many false positives, where the rule is blocking traffic that you do not want blocked, you can disable the rule without needing to switch to a less-secure intrusion policy. You could alternatively change it to alert on matches without dropping traffic.

Conversely, if you know you need to protect against a specific attack, but the related rule is disabled in your chosen intrusion policy, you can enable the rule without changing to a more secure policy.

Use the intrusion related dashboards, and the Event Viewer (both on the **Monitoring** page), to evaluate how intrusion rules are impacting traffic. Keep in mind that you will see intrusion events and intrusion data only for traffic that matches intrusion rules set to alert or drop; disabled rules are not evaluated.

The following topics explain intrusion policies and rule tuning in more detail.

Configure the Inspection Mode for an Intrusion Policy (Snort 2)

By default, all intrusion policies operate in Prevention mode to implement an Intrusion Prevention System (IPS). In the Prevention inspection mode, if a connection matches an intrusion rule whose action is to drop traffic, the connection is actively blocked.

If you instead want to test the effect of the intrusion policy on your network, you can change the mode to Detection, which implements an Intrusion Detection System (IDS). In this inspection mode, drop rules are treat like alert rules, where you are notified of matching connections, but the action result becomes Would Have Blocked, and connections are never in fact blocked.

You change the inspection mode per intrusion policy, so you can have a mix of prevention and detection.

Procedure

Step 1 Select **Policies** > **Intrusion**.

Step 2 Click the tab for the intrusion policy whose inspection mode you want to change.

The **Inspection Mode** is indicated above the rules table.

Step 3 Click the **Edit** link next to the inspection mode, change the mode for the policy, and click **OK**.

The options are:

- **Prevention**—Intrusion rule actions are always applied. Connections that match a drop rule are blocked.
- **Detection**—Intrusion rules generate alerts only. A connection that matches a drop rule will generate alert messages, but the connection will not be blocked.

Changing Intrusion Rule Actions (Snort 2)

Each pre-defined intrusion policy has the same rules. The difference is the action taken for each rule can be different from policy to policy.

By changing the rule action, you can disable rules that are giving you too many false positives, or you can change whether the rule alerts on or drops matching traffic. You can also enable disabled rules to alert or drop matching traffic.

Procedure

Step 1 Select **Policies** > **Intrusion**.

Step 2 Click the tab for the intrusion policy whose rule actions you want to change.

The pre-defined policies are:

- Connectivity over Security
- Balanced Security and Connectivity

- Security over Connectivity
- Maximum Detection

Step 3 Find the rule whose action you want to change.

The rules are sorted with the overridden ones listed first, and sorted by action within the group of overridden rules. Otherwise, the rules are sorted by GID, then SID.

Use the search box to locate the rule you want to change. Ideally, you can get the Snort identifier (SID) and generator identifier (GID) from an event or from Cisco Technical Support, if you are working with them on an issue.

For detailed information about the elements of each rule, see Intrusion Rule Attributes, on page 511.

To search the list:

a) Click in the **Search** box to open the search attributes dialog box.
b) Enter a combination of Generator ID (**GID**), Snort ID (**SID**), or rule **Action**, and click **Search**.

For example, you could select **Action = Drop** to see all the rules in the policy that will drop matching connections. The text beside the search box indicates how many rules match your criteria, for example, "8937 of 9416 rules found."

To clear a search criteria, click the x for the criteria in the search box.

Step 4 Click the **Action** column for the rule and select the required action:

- **Alert**—Create an event when this rule matches traffic, but do not drop the connection.
- **Drop**—Create an event when this rule matches traffic, and also drop the connection.
- **Disabled**—Do not match traffic against this rule. No events are generated.

The default action for the rule is indicated by "(Default)" added to the action. If you change the default, the status column indicates "Overridden" for that rule.

Monitoring Intrusion Policies

You can find intrusion policy statistics on the **Attackers** and **Targets** dashboards on the **Monitoring** page. You must apply an intrusion policy to at least one access control rule to see any information on these dashboards. See Monitoring Traffic and System Dashboards, on page 102.

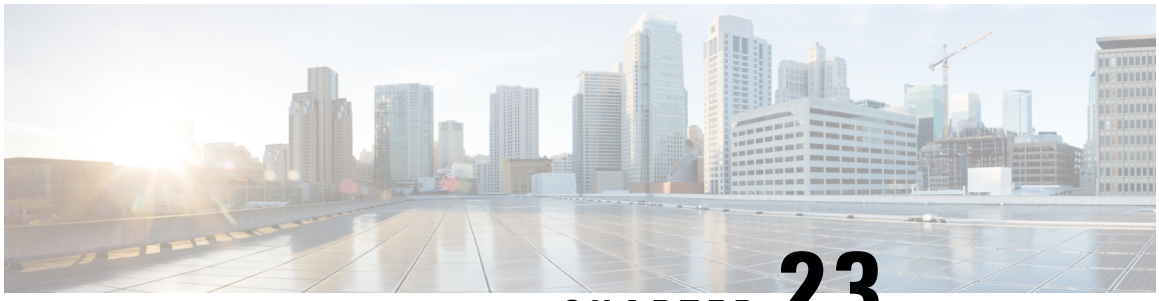
To see intrusion events, select **Monitoring** > **Events**, then click the **Intrusion** tab. You can hover over an event and click the **View Details** link to get more information. From the details page, you can click the **View IPS Rule** to go to the rule in the relevant intrusion policy, where you can change the rule action. This can help you reduce the impact of false positives, where a rule is blocking too many good connections, by changing the action from drop to alert. Conversely, you can change an alert rule into a drop rule if you are seeing a lot of attack traffic for a rule.

If you configure a syslog server for the intrusion policy, intrusion events have the message ID 430001.

Examples for Intrusion Policies

The use case chapter includes the following examples of implementing intrusion policies.

- How to Block Threats, on page 51
- How to Passively Monitor the Traffic on a Network, on page 72



CHAPTER 23

Network Address Translation (NAT)

The following topics explain Network Address Translation (NAT) and how to configure it.

- Why use NAT?, on page 539
- NAT basics, on page 540
- Guidelines for NAT, on page 546
- Configure NAT, on page 552
- Translating IPv6 networks, on page 578
- Monitoring NAT, on page 592
- Examples for NAT, on page 593

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:

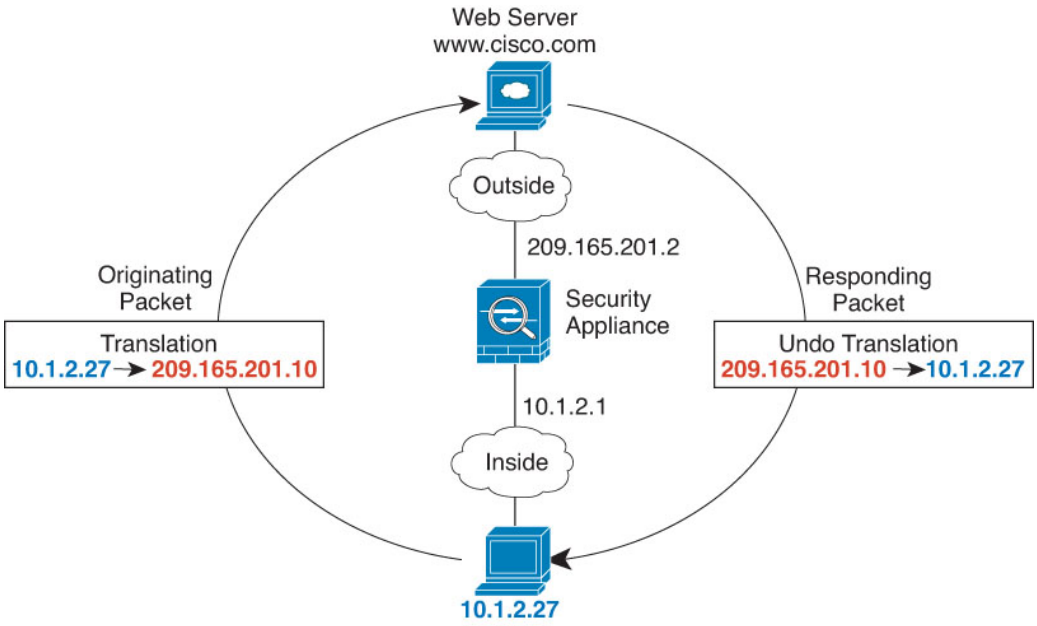
- Dynamic NAT—A group of real IP addresses are mapped to a (usually smaller) group of mapped IP addresses, on a first come, first served basis. Only the real host can initiate traffic. See Dynamic NAT, on page 552.
- Dynamic Port Address Translation (PAT)—A group of real IP addresses are mapped to a single IP address using a unique source port of that IP address. See Dynamic PAT, on page 557.

- Static NAT—A consistent mapping between a real and mapped IP address. Allows bidirectional traffic initiation. See Static NAT, on page 562.
- 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 570.

NAT in routed mode

The following figure shows a typical NAT example in routed mode, with a private network on the inside.

Figure 28: 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 Firewall Threat Defense receives the packet because the Firewall Threat Defense performs proxy ARP to claim the packet.
3. The Firewall Threat Defense 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.

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.

Note There is also a Section 0, which contains any NAT rules that the system creates for its own use. These rules have priority over all others. The system automatically creates these rules and clears xlates as needed. You cannot add, edit, or modify rules in Section 0.

Table 13: 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. If you do not follow the above recommendation, ensure that accurate routes are configured for the mapped addresses to match the correct NAT rule.

Table Section	Rule Type	Order of Rules within the Section
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:

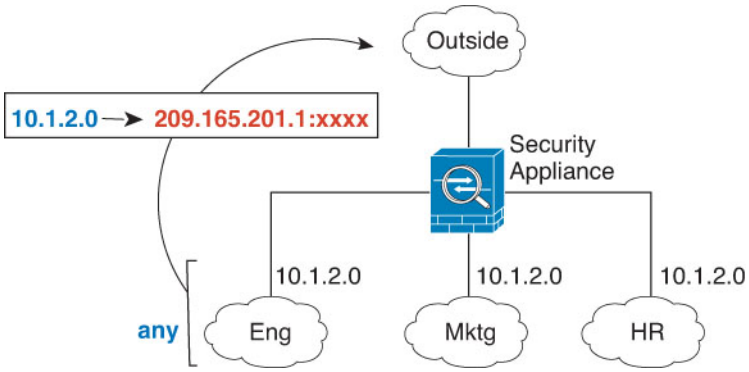
- 192.168.1.1/32 (static)
- 10.1.1.0/24 (static)
- 192.168.1.0/24 (static)
- 172.16.1.0/24 (dynamic) (object abc)
- 172.16.1.0/24 (dynamic) (object def)
- 192.168.1.0/24 (dynamic)

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 29: 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 Firewall 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 Firewall Threat Defense 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 Firewall Threat Defense 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.

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 Firewall Threat Defense.

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 Firewall Threat Defense will then proxy ARP for the address, even though the packet is not actually destined for the Firewall Threat Defense. (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 Firewall Threat Defense ARP response is received before the actual host ARP response, then traffic will be mistakenly sent to the Firewall Threat Defense.

Guidelines for NAT

The following topics provide detailed guidelines for implementing NAT.

Interface Guidelines

NAT is supported for standard routed physical or subinterfaces.

However, configuring NAT on bridge group member interfaces (interfaces that are part of a Bridge 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 source and destination interfaces. You cannot specify “any” as the interface.
- You cannot configure interface PAT when the destination interface 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).

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.
- 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. If you use these protocols on non-standard ports, do not use NAT on the connections.

Table 14: 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	No limitations.	Yes
H.323 H.225 (Call signaling) H.323 RAS	TCP/1720 UDP/1718 For RAS, UDP/1718-1719	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 NAT64.	No
RSH	TCP/514	No PAT. No NAT64.	Yes

Application	Inspected Protocol, Port	NAT Limitations	Pinholes Created
RTSP	TCP/554 (No handling for HTTP cloaking.)	No NAT64.	Yes
SIP	TCP/5060 UDP/5060	No extended PAT. No NAT64 or NAT46.	Yes
Skinny (SCCP)	TCP/2000	No NAT64, NAT46, or NAT66.	Yes
SQL*Net (versions 1, 2)	TCP/1521	No NAT64.	Yes
Sun RPC	TCP/111 UDP/111	No NAT64.	Yes
TFTP	UDP/69	No NAT64. Payload IP addresses are not translated.	Yes
XDMCP	UDP/177	No NAT64.	Yes

FQDN destination guidelines

You can specify the translated (mapped) destination in a manual NAT rule using a fully-qualified domain name (FQDN) network object instead of an IP address. For example, you can create a rule based on traffic that is destined for the www.example.com web server.

When using an FQDN, the system obtains the DNS resolution and writes the NAT rule based on the returned address. If more than one address is obtained from the DNS server, the address used is based on the following:

- If there is an address on the same subnet as the specified interface, that address is used. If there isn't one on the same subnet, the first address returned is used.
- The IP type for the translated source and translated destination must match. For example, if the translated source address is IPv6, the FQDN object must specify IPv6 as the address type. If the translated source is IPv4, the FQDN object can specify IPv4 or both IPv4 and IPv6. In this case, an IPv4 address is selected.

You cannot include an FQDN object in a network group that is used for manual NAT destination. In NAT, an FQDN object must be used alone, as only a single destination host makes sense for this type of NAT rule.

If the FQDN cannot be resolved to an IP address, the rule is not functional until a DNS resolution is obtained.

Additional guidelines for NAT

- NAT rules apply to through the device traffic only, they do not apply to traffic initiated by the device, such as a RADIUS authentication.
- 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.

- You cannot write NAT rules for a Virtual Tunnel Interface (VTI), which are used in site-to-site VPN. Writing rules for the VTI's source interface will not apply NAT to the VPN tunnel. To write NAT rules that will apply to VPN traffic tunneled on a VTI, you must use "any" as the interface; you cannot explicitly specify interface names.
- (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.
- A network object used in NAT cannot include more than 131,838 IP addresses, either explicitly or implied in a range of addresses or a subnet. Break up the address space into smaller ranges and write separate rules for the smaller objects.
- (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 Firewall Threat Defense performs NAT on a packet, the packet must be IPv6-to-IPv6 or IPv4-to-IPv4; with this prerequisite, the Firewall Threat Defense 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.
  - (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.
- 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.
- 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.
- DHCP and BOOTP share ports UDP/67-68. Because BOOTP is obsolete, writing NAT rules for the bootps port can cause port allocation problems when also running DHCP. Consider using DHCP relay instead for transmitting DHCP requests between network segments.
- In rare cases, return traffic (server to client) with an existing translation (xlate) might be logged as a new flow in Connection Events. This can occur when the client has already terminated the connection and the server sends another packet that reaches the device during the brief interval between connection closure and xlate removal, often due to application behavior or TCP stack cleanup. Because the device removes the xlate only after deleting the connection, a server packet can arrive while the xlate still exists. If no valid connection entry is found, the device logs a separate connection event based on the matched Access Control Policy rule.

Configure NAT

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.

Procedure

Step 1 Select **Policies** > **NAT**.

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 540.

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 541.

Step 4 Create the rules as explained in the following sections.

- Dynamic NAT, on page 552
- Dynamic PAT, on page 557
- Static NAT, on page 562
- Identity NAT, on page 570

Step 5 Manage the NAT policy and rules.

You can do the following to manage the policy and its rules.

- To edit a rule, click the edit icon ( ) for the rule.
- To delete a rule, click the delete icon ( ) for the rule.

Dynamic NAT

The following topics explain dynamic NAT and how to configure it.

About dynamic NAT

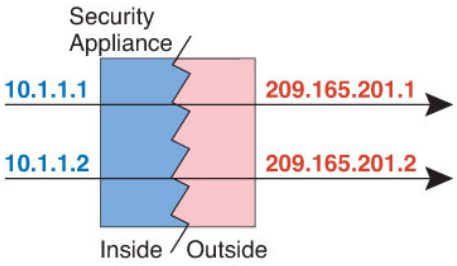
Dynamic NAT translates a group of real addresses to a pool of mapped addresses that are routable on the destination network. The mapped pool typically includes fewer addresses than the real group. When a host you want to translate accesses the destination network, NAT assigns the host an IP address from 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. A successful connection from a remote host can reset the idle timer for the connection.

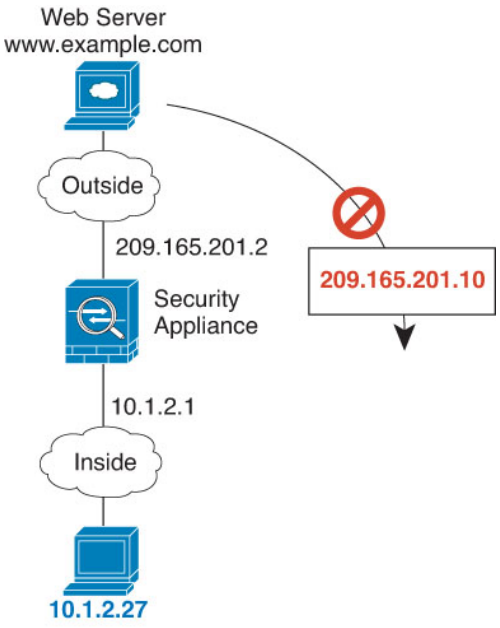
The following figure shows a typical dynamic NAT scenario. Only real hosts can create a NAT session, and responding traffic is allowed back.

Figure 30: 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 31: 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** 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 Address**—This must be a network object (not a group), and it can be a host, range, or subnet.
- **Translated Address**—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. If the object contains one host address only, it is used for PAT.

Procedure

Step 1 Select **Policies** > **NAT**.

Step 2 Do one of the following:

- To create a new rule, click the + button.
- To edit an existing rule, click the edit icon () for the rule.

(To delete a rule you no longer need, click the trash can icon for the rule.)

Step 3 Configure the basic rule options:

- **Title**—Enter a name for the rule.
- **Create Rule For**—Select **Auto NAT**.
- **Type**—Select **Dynamic**.

Step 4 Configure the following packet translation options:

- **Source Interface**, **Destination Interface**—(Required for bridge group member interfaces.) The interfaces where this NAT rule applies. **Source** is the real interface, the one through which the traffic enters the

device. **Destination** is 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.

- **Original Address**—The network object that contains the addresses you are translating.
- **Translated Address**—The network object or group that contains the mapped addresses.

Step 5 (Optional.) Click the **Advanced Options** link and 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 613.
- **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.

Step 6 Click **OK**.

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** 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 Address**—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 Address**—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. If the object contains one host address only, it is used for PAT.

You can also create network objects for the **Original Destination Address** and **Translated Destination Address** 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 **Policies** > **NAT**.

Step 2 Do one of the following:

- To create a new rule, click the + button.
- To edit an existing rule, click the edit icon ( ) for the rule.

(To delete a rule you no longer need, click the trash can icon for the rule.)

Step 3 Configure the basic rule options:

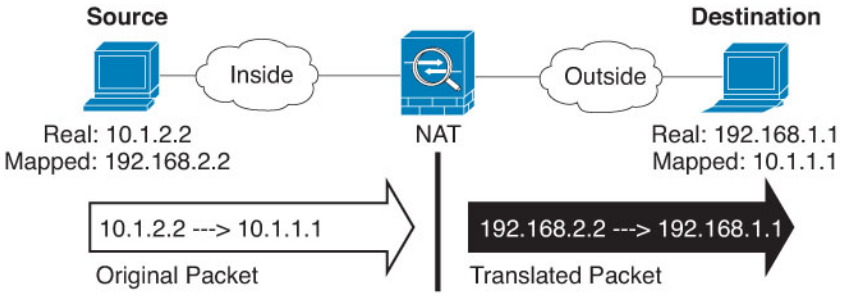
- **Title**—Enter a name for the rule.
- **Create Rule For**—Select **Manual NAT**.
- **Rule Placement**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule you select.
- **Type**—Select **Dynamic**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.

Step 4 Configure the following interface options:

- **Source Interface**, **Destination Interface**—(Required for bridge group member interfaces.) The interfaces where this NAT rule applies. **Source** is the real interface, the one through which the traffic enters the device. **Destination** is 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.



- **Original Source Address**—The network object or group that contains the addresses you are translating.
- **Original Destination Address**—(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** 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 Address**—The network object or group that contains the mapped addresses.
- **Translated Destination Address**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination Address**, 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.) Click the **Advanced Options** link and 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 613.
- **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.

Step 9 Click **OK**.

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 32: 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 Firewall Threat Defense interface IP address as the PAT address. However, you cannot use interface PAT for the IPv6 addresses on the interface.

You cannot use dynamic PAT for IPv6 (NAT66) when translating between interfaces in the same bridge group. This restriction does not apply 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 547.

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.

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.

Before you begin

Select **Objects** 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 Address**—This must be a network object (not a group), and it can be a host, range, or subnet.
- **Translated Address**—You have the following options to specify the PAT address:
  - **Destination Interface**—To use the destination interface IPv4 address, you do not need a network object. You cannot use interface PAT for IPv6.

- **Single PAT address**—Create a network object containing a single host.

Procedure

Step 1 Select **Policies** > **NAT**.

Step 2 Do one of the following:

- To create a new rule, click the + button.
- To edit an existing rule, click the edit icon ( ) for the rule.

(To delete a rule you no longer need, click the trash can icon for the rule.)

Step 3 Configure the basic rule options:

- **Title**—Enter a name for the rule.
- **Create Rule For**—Select **Auto NAT**.
- **Type**—Select **Dynamic**.

Step 4 Configure the following packet translation options:

- **Source Interface**, **Destination Interface**—(Required for bridge group member interfaces.) The interfaces where this NAT rule applies. **Source** is the real interface, the one through which the traffic enters the device. **Destination** is 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.
- **Original Address**—The network object that contains the addresses you are translating.
- **Translated Address**—One of the following:
  - (Interface PAT.) To use the IPv4 address of the destination interface, select **Interface**. You must also select a specific destination interface, which cannot be a bridge group member interface. You cannot use interface PAT for IPv6.
  - To use a single address other than the destination interface address, select the host network object you created for this purpose.

Step 5 (Optional.) Click the **Advanced Options** link and 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. You cannot select this option if you already configured interface PAT as the translated address. You also cannot use this option with IPv6 networks.

Step 6 Click **OK**.

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.

Before you begin

Select **Objects** 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 Address**—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 Address**—You have the following options to specify the PAT address:
  - **Destination Interface**—To use the destination interface IPv4 address, you do not need a network object. You cannot use interface PAT for IPv6.
  - **Single PAT address**—Create a network object containing a single host.

You can also create network objects for the **Original Destination Address** and **Translated Destination Address** if you are configuring a static translation for those addresses in the rule.

For dynamic PAT, 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 **Policies** > **NAT**.

Step 2 Do one of the following:

- To create a new rule, click the + button.
- To edit an existing rule, click the edit icon () for the rule.

(To delete a rule you no longer need, click the trash can icon for the rule.)

Step 3 Configure the basic rule options:

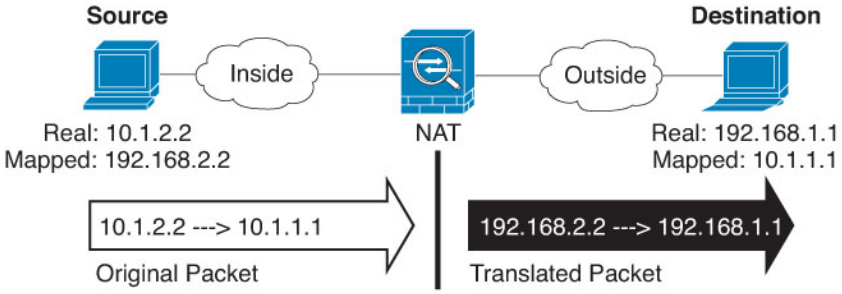
- **Title**—Enter a name for the rule.
- **Create Rule For**—Select **Manual NAT**.
- **Rule Placement**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule you select.
- **Type**—Select **Dynamic**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.

Step 4 Configure the following interface options:

- **Source Interface**, **Destination Interface**—(Required for bridge group member interfaces.) The interfaces where this NAT rule applies. **Source** is the real interface, the one through which the traffic enters the device. **Destination** is 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.



- **Original Source Address**—The network object or group that contains the addresses you are translating.
- **Original Destination Address**—(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** 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 Address**—One of the following:
  - (Interface PAT.) To use the IPv4 address of the destination interface, select **Interface**. You must also select a specific destination interface, which cannot be a bridge group member interface. You cannot use interface PAT for IPv6.
  - To use a single address other than the destination interface address, select the host network object you created for this purpose.
- **Translated Destination Address**—(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.) Click the **Advanced Options** link and 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. You cannot select this option if you already configured interface PAT as the translated address. You also cannot use this option with IPv6 networks.

Step 9 Click **OK**.

Static NAT

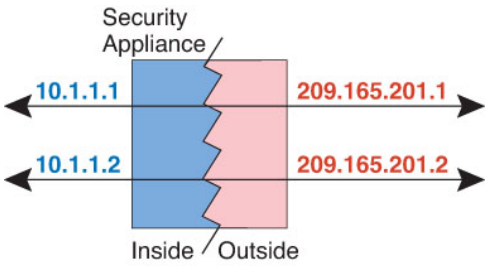
The following topics explain static NAT and how to implement it.

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 33: Static NAT



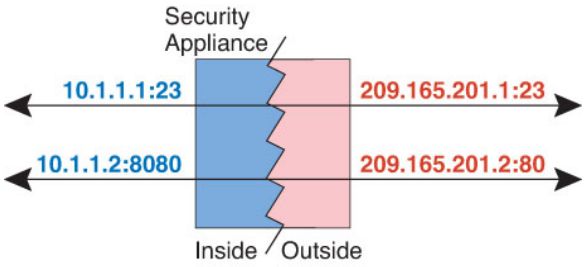
Static NAT with port translation

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 34: 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

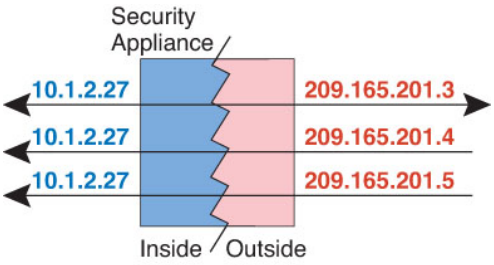
You can configure static NAT to map a real address to an interface address/port combination. For example, if you want to redirect Telnet access for the device's outside interface to an inside host, then you can map the inside host IP address/port 23 to the outside interface address/port 23.

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.

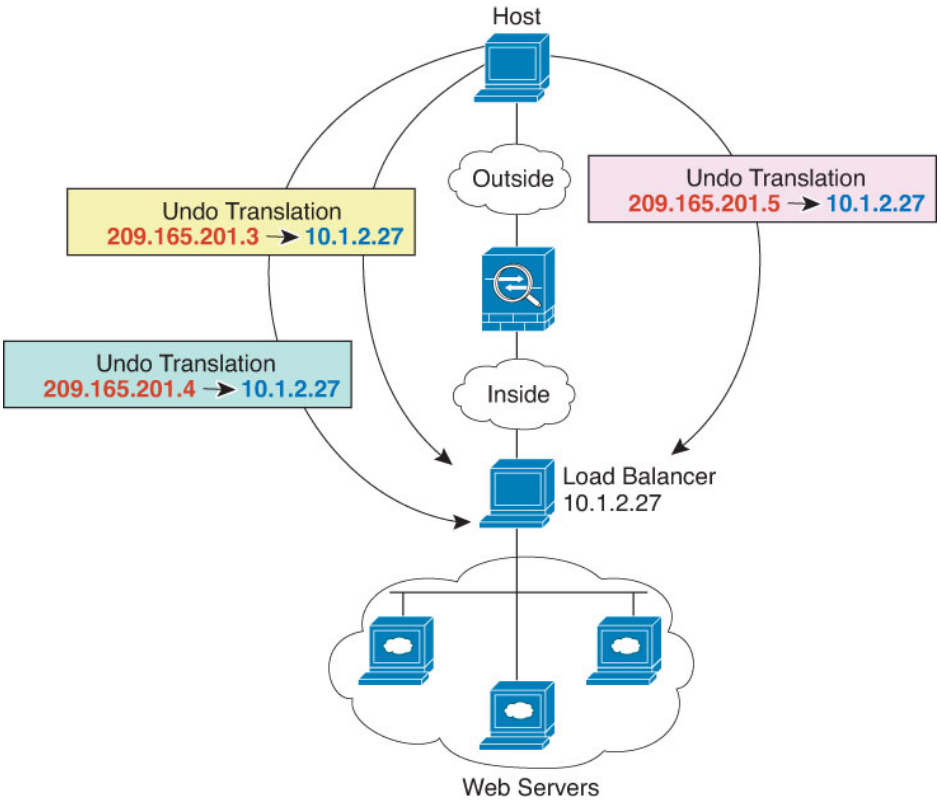
The following figure shows a typical one-to-many static NAT scenario. Because initiation by the real host always uses the first mapped address, the translation of real host IP/first mapped IP is technically the only bidirectional translation.

Figure 35: 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 36: One-to-Many Static NAT Example



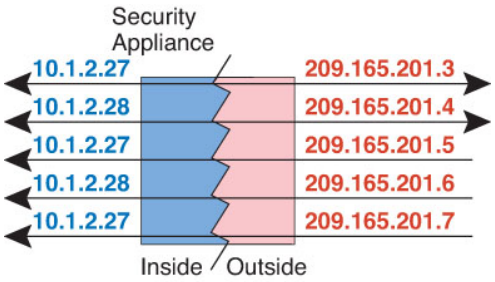
Other mapping scenarios (not recommended)

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 37: Few-to-Many Static NAT

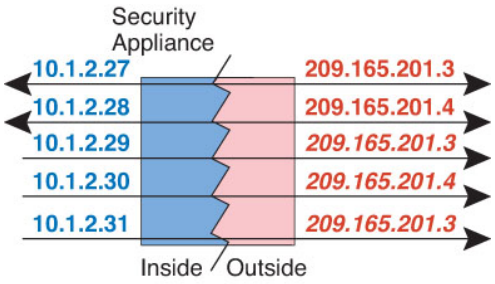


For a many-to-few or many-to-one configuration, where you have more real addresses than mapped addresses, you run out of mapped addresses before you run out of real addresses. Only the mappings between the lowest real IP addresses and the mapped pool result in bidirectional initiation. The remaining 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 38: 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** 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 Address**—This must be a network object (not a group), and it can be a host, range, or subnet.
- **Translated Address**—You have the following options to specify the translated address:
  - **Destination Interface**—To use the destination interface IPv4 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. You cannot use interface PAT for IPv6.
  - **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 **Policies** > **NAT**.

Step 2 Do one of the following:

- To create a new rule, click the + button.
- To edit an existing rule, click the edit icon () for the rule.

(To delete a rule you no longer need, click the trash can icon for the rule.)

Step 3 Configure the basic rule options:

- **Title**—Enter a name for the rule.
- **Create Rule For**—Select **Auto NAT**.
- **Type**—Select **Static**.

Step 4 Configure the following packet translation options:

- **Source Interface**, **Destination Interface**—(Required for bridge group member interfaces.) The interfaces where this NAT rule applies. **Source** is the real interface, the one through which the traffic enters the device. **Destination** is 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.
- **Original Address**—The network object that contains the addresses you are translating.
- **Translated Address**—One of the following:
  - To use a set group of addresses, select 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 **Interface**. You must also select a specific destination interface, which cannot be a bridge group member interface. You cannot use interface PAT for IPv6. 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 port objects that define the original and translated ports. The objects must be for the same protocol. Click the **Create New Object** link if the objects do not already exist. For example, you can translate TCP/80 to TCP/8080 if necessary.

Step 5 (Optional.) Click the **Advanced Options** link and 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 613. This option is not available if you are doing port translation.
- **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 6 Click **OK**.

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** 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 Address**—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 Address**—You have the following options to specify the translated address:

  - **Destination Interface**—To use the destination interface IPv4 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. You cannot use interface PAT for IPv6.

- **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.

You can also create network objects for the **Original Destination Address** and **Translated Destination Address** 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 **Policies** > **NAT**.

Step 2 Do one of the following:

- To create a new rule, click the + button.
- To edit an existing rule, click the edit icon ( ) for the rule.

(To delete a rule you no longer need, click the trash can icon for the rule.)

Step 3 Configure the basic rule options:

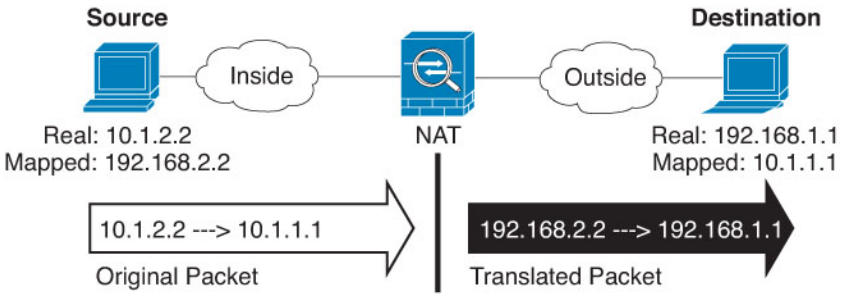
- **Title**—Enter a name for the rule.
- **Create Rule For**—Select **Manual NAT**.
- **Rule Placement**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule you select.
- **Type**—Select **Static**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.

Step 4 Configure the following interface options:

- **Source Interface**, **Destination Interface**—(Required for bridge group member interfaces.) The interfaces where this NAT rule applies. **Source** is the real interface, the one through which the traffic enters the device. **Destination** is 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.



- **Original Source Address**—The network object or group that contains the addresses you are translating.
- **Original Destination Address**—(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** 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 Address**—One of the following:
  - To use a set group of addresses, select 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 IPv4 address of the destination interface, select **Interface**. You must also select a specific destination interface, which cannot be a bridge group member interface. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number. You cannot use interface PAT for IPv6.
- **Translated Destination Address**—(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.) Click the **Advanced Options** link and 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 613. This option is not available if you are doing port translation.
- **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.

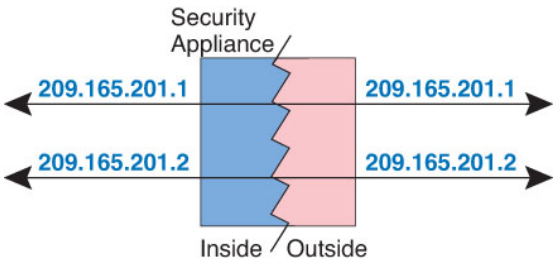
Step 9 Click **OK**.

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 that network to itself.

The following figure shows a typical identity NAT scenario.

Figure 39: 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** 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 Address**—This must be a network object (not a group), and it can be a host, range, or subnet.

- **Translated Address**—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 **Policies** > **NAT**.

Step 2 Do one of the following:

- To create a new rule, click the + button.
- To edit an existing rule, click the edit icon () for the rule.

(To delete a rule you no longer need, click the trash can icon for the rule.)

Step 3 Configure the basic rule options:

- **Title**—Enter a name for the rule.
- **Create Rule For**—Select **Auto NAT**.
- **Type**—Select **Static**.

Step 4 Configure the following packet translation options:

- **Source Interface**, **Destination Interface**—(Required for bridge group member interfaces.) The interfaces where this NAT rule applies. **Source** is the real interface, the one through which the traffic enters the device. **Destination** is 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.
- **Original Address**—The network object that contains the addresses you are translating.
- **Translated Address**—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 5 (Optional.) Click the **Advanced Options** link and select the desired options:

- **Translate DNS replies that match this rule**—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 6 Click **OK**.

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** 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 Address**—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 Address**—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 Address** and **Translated Destination Address** 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 **Policies** > **NAT**.

Step 2 Do one of the following:

- To create a new rule, click the + button.
- To edit an existing rule, click the edit icon () for the rule.

(To delete a rule you no longer need, click the trash can icon for the rule.)

Step 3 Configure the basic rule options:

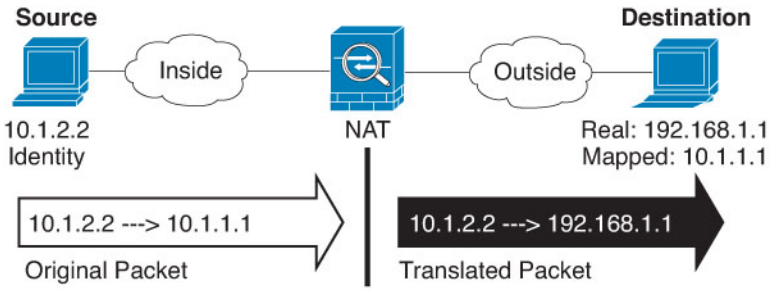
- **Title**—Enter a name for the rule.
- **Create Rule For**—Select **Manual NAT**.
- **Rule Placement**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule you select.
- **Type**—Select **Static**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.

Step 4 Configure the following interface options:

- **Source Interface**, **Destination Interface**—(Required for bridge group member interfaces.) The interfaces where this NAT rule applies. **Source** is the real interface, the one through which the traffic enters the device. **Destination** is 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 Address**—The network object or group that contains the addresses you are translating.
- **Original Destination Address**—(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** 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 Address**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.
- **Translated Destination Address**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination Address**, 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.) Click the **Advanced Options** link and select the desired options:

- **Translate DNS replies that match this rule**—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 9 Click **OK**.

NAT Rule Properties

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, 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.

Title

Enter a name for the rule. The name cannot include spaces.

Create Rule For

Whether the translation rule is **Auto NAT** or **Manual NAT**. Auto NAT is simpler than manual NAT, but manual NAT allows you to create separate translations for a source address based on the destination address.

Status

Whether you want the rule to be active or disabled.

Placement (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 you select.

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.

The following topics describe the remaining NAT rules properties.

Packet Translation Properties for Auto NAT

Use the **Packet Translation** options to define the source addresses and the mapped translated addresses. The following properties apply to auto NAT only.

Source Interface, Destination Interface

(Required for bridge group member interfaces.) The interfaces where this NAT rule applies. **Source** is the real interface, the one through which the traffic enters the device. **Destination** is 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.

Original Address (Always required.)

The network object that contains the source addresses you are translating. This must be a network object (not a group), and it can be a host, range, or subnet.

Translated Address (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. If the object contains one host address only, it is used for PAT.
- **Dynamic PAT**—One of the following:
  - (Interface PAT.) To use the IPv4 address of the destination interface, select **Interface**. You must also select a specific destination interface, which cannot be a bridge group member interface. You cannot use interface PAT for IPv6.
  - To use a single address other than the destination interface address, select the host network object you created for this purpose.
- **Static NAT**—One of the following:
  - To use a set group of addresses, select 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 **Interface**. You must also select a specific destination interface, which cannot be a bridge group member interface. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number. You cannot use interface PAT for IPv6.
- **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 port objects that define the original and translated ports. The objects must be for the same protocol. For example, you can translate TCP/80 to TCP/8080 if necessary.

Packet Translation Properties for Manual NAT

Use the **Packet Translation** options to define the source addresses and the mapped translated addresses. The following properties apply to manual NAT only. All are optional except as indicated.

Source Interface, Destination Interface

(Required for bridge group member interfaces.) The interfaces where this NAT rule applies. **Source** is the real interface, the one through which the traffic enters the device. **Destination** is 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.

Original Source Address (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 Address (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. If the object contains one host address only, it is used for PAT.
- **Dynamic PAT**—One of the following:
  - (Interface PAT.) To use the address of the destination interface, select **Interface**. You must also select a specific destination interface, which cannot be a bridge group member interface. You cannot use interface PAT for IPv6.
  - To use a single address other than the destination interface address, select the host network object you created for this purpose.
- **Static NAT**—One of the following:
  - To use a set group of addresses, select 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 **Interface**. You must also select a specific destination interface, which cannot be a bridge group member interface. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number. You cannot use interface PAT for IPv6.
- **Identity NAT**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.

Original Destination Address

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** 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 Address

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.

You can use a network object that specifies a fully-qualified domain name as the translated destination; for more information, see FQDN destination guidelines, on page 549.

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 object for both the real and mapped ports.

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 613. 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. You cannot select this option if you already configured interface PAT as the translated address. You cannot use this option with IPv6 networks.

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.

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.

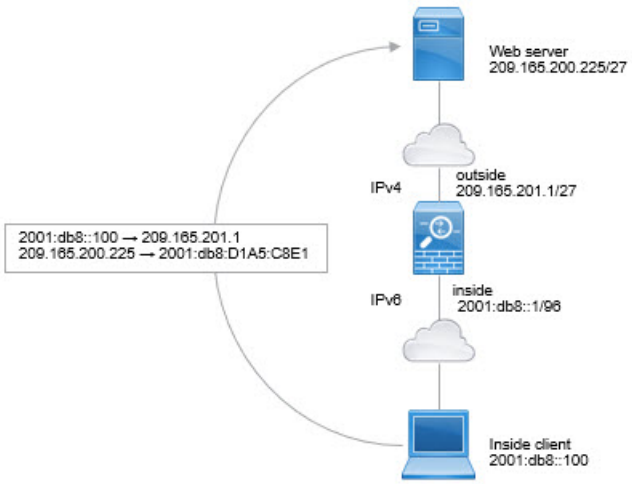
- The IPv4 address pool for the NAT64 rule is normally small and typically might not have enough addresses to map one-to-one with the IPv6 client addresses. Dynamic PAT might more easily meet the possible large number of IPv6 client addresses compared to dynamic or static NAT.

- 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

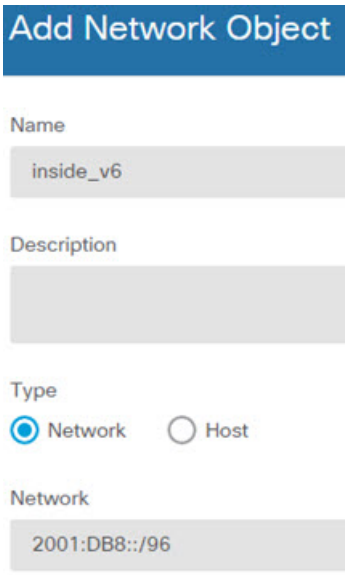
Step 1 Create a network object for the inside IPv6 network.

a) Choose **Objects**.

b) Select **Network** from the table of contents and click +.

c) Define the inside IPv6 network.

Name the network object (for example, inside_v6), select **Network**, and enter the network address, 2001:db8::/96.

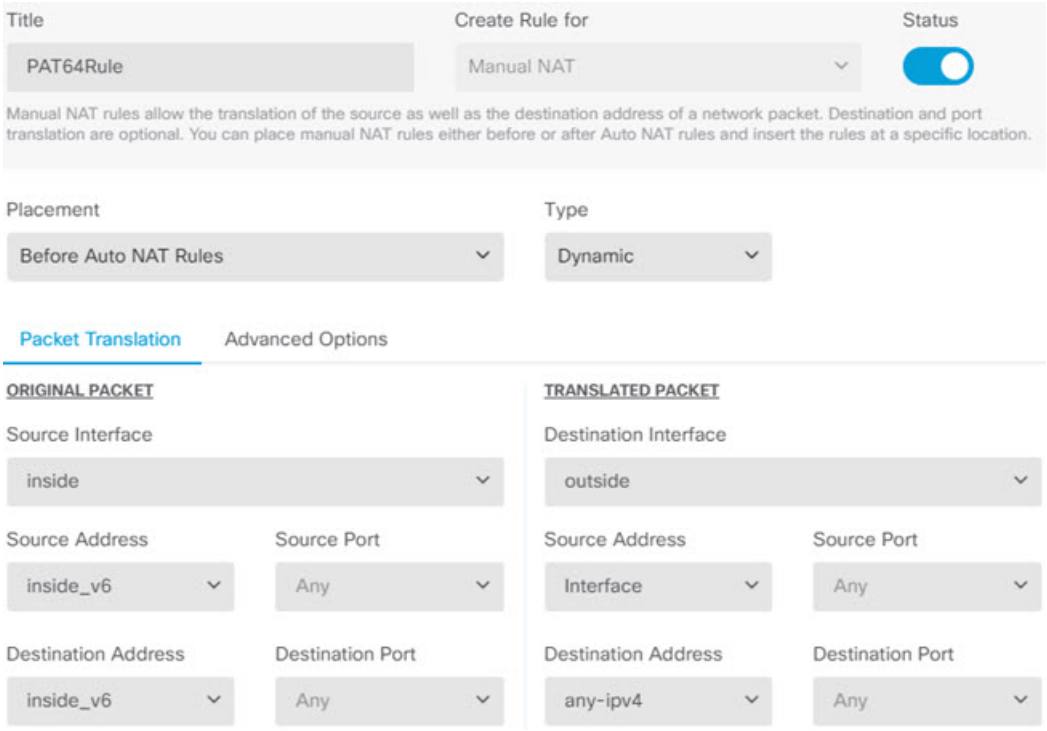


d) Click **OK**.

Step 2 Create the manual NAT rule to translate the IPv6 network to IPv4 and back again.

a) Select **Policies** > **NAT**.
b) Click the **+** button.
c) Configure the following properties:

- **Title** = PAT64Rule (or another name of your choosing).
- **Create Rule For** = **Manual NAT**.
- **Placement** = **Before Auto NAT Rules**
- **Type** = **Dynamic**.
- **Source Interface** = inside.
- **Destination Interface** = outside.
- **Original Packet Source Address** = inside_v6 network object.
- **Translated Packet Source Address** = Interface. This option uses the IPv4 address of the destination interface as the PAT address.
- **Original Packet Destination Address** = inside_v6 network object.
- **Translated Packet Destination Address** = any-ipv4 network object.

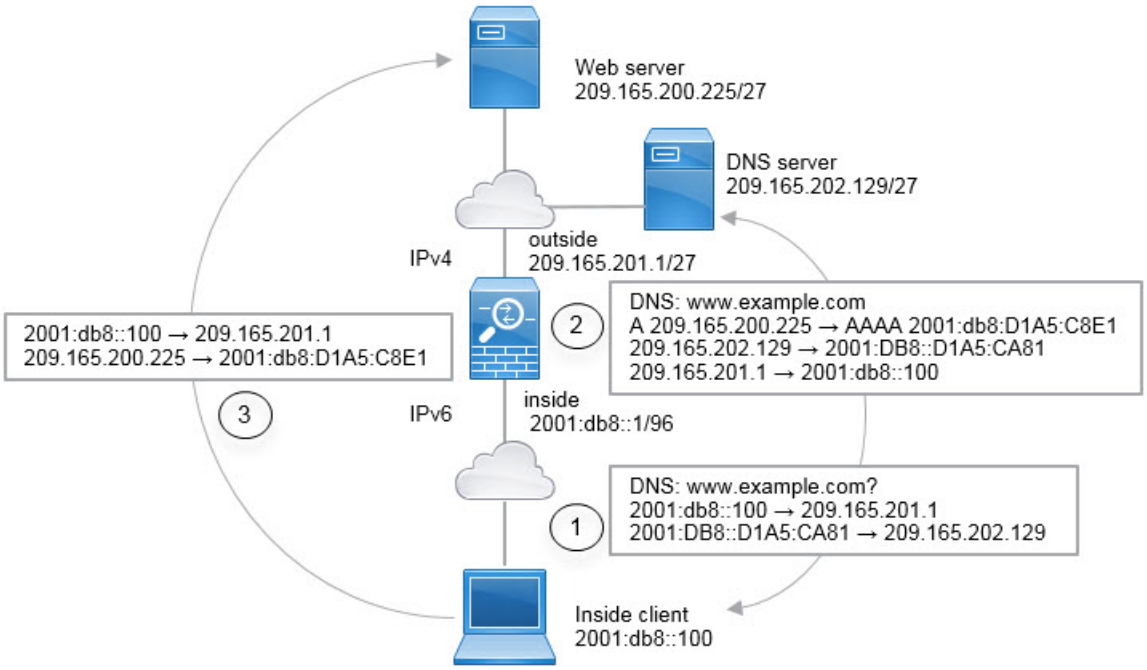


d) Click **OK**.

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.

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.

Procedure

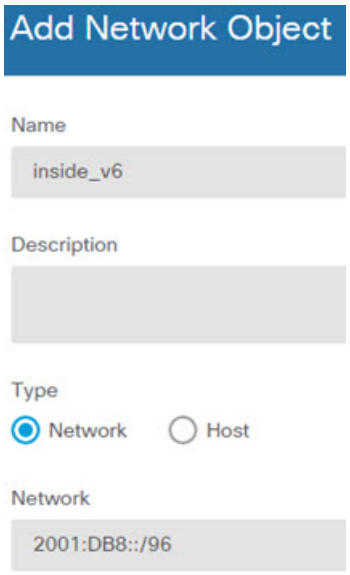
Step 1 Create the network objects that define the inside IPv6 and outside IPv4 networks.

a) Choose **Objects**.

b) Select **Network** from the table of contents and click +.

c) Define the inside IPv6 network.

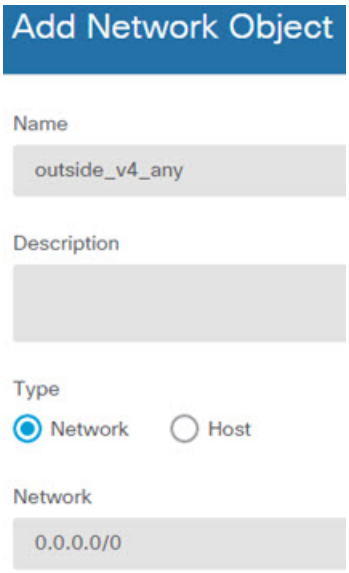
Name the network object (for example, inside_v6), select **Network**, and enter the network address, 2001:db8::/96.



d) Click **OK**.

e) Click + and define the outside IPv4 network.

Name the network object (for example, outside_v4_any), select **Network**, and enter the network address 0.0.0.0/0.



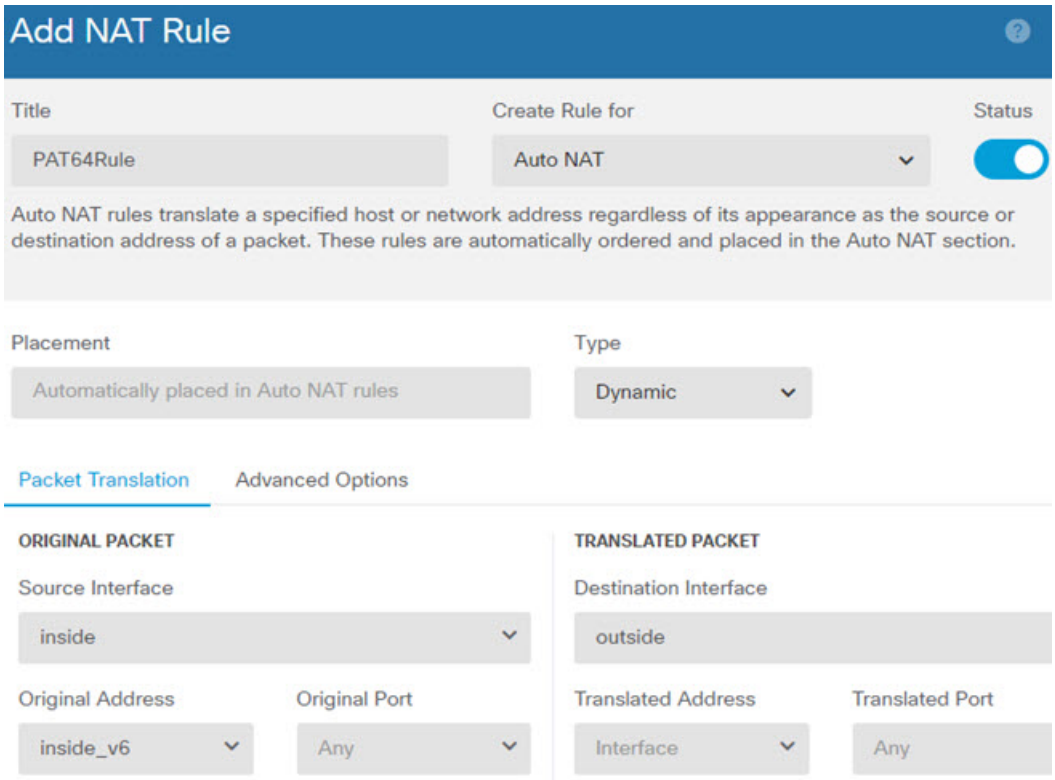
Step 2 Configure the NAT64 dynamic PAT rule for the inside IPv6 network.

a) Select **Policies** > **NAT**.

b) Click the + button.

c) Configure the following properties:

- **Title** = PAT64Rule (or another name of your choosing).
- **Create Rule For** = Auto NAT.
- **Type** = Dynamic.
- **Source Interface** = inside.
- **Destination Interface** = outside.
- **Original Address** = inside_v6 network object.
- **Translated Address** = **Interface**. This option uses the IPv4 address of the destination interface as the PAT address.



d) Click **OK**.

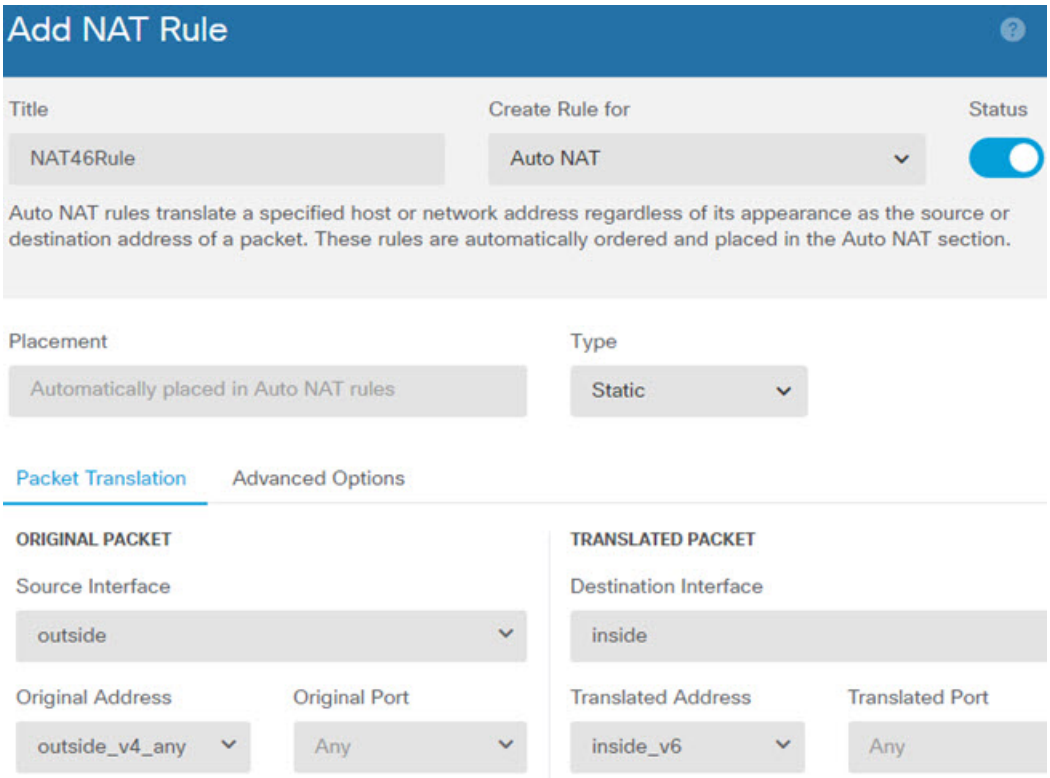
With this rule, any traffic from the 2001:db8::/96 subnet on the inside interface going to the outside interface gets a NAT64 PAT translation using the IPv4 address of the outside interface.

Step 3 Configure the static NAT46 rule for the outside IPv4 network.

a) Click the + button.

b) Configure the following properties:

- **Title** = NAT46Rule (or another name of your choosing).
- **Create Rule For** = Auto NAT.
- **Type** = Static.
- **Source Interface** = outside.
- **Destination Interface** = inside.
- **Original Address** = outside_v4_any network object.
- **Translated Address** = inside_v6 network object.
- On the **Advanced Options** tab, select **Translate DNS replies that match this rule**.



c) 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.

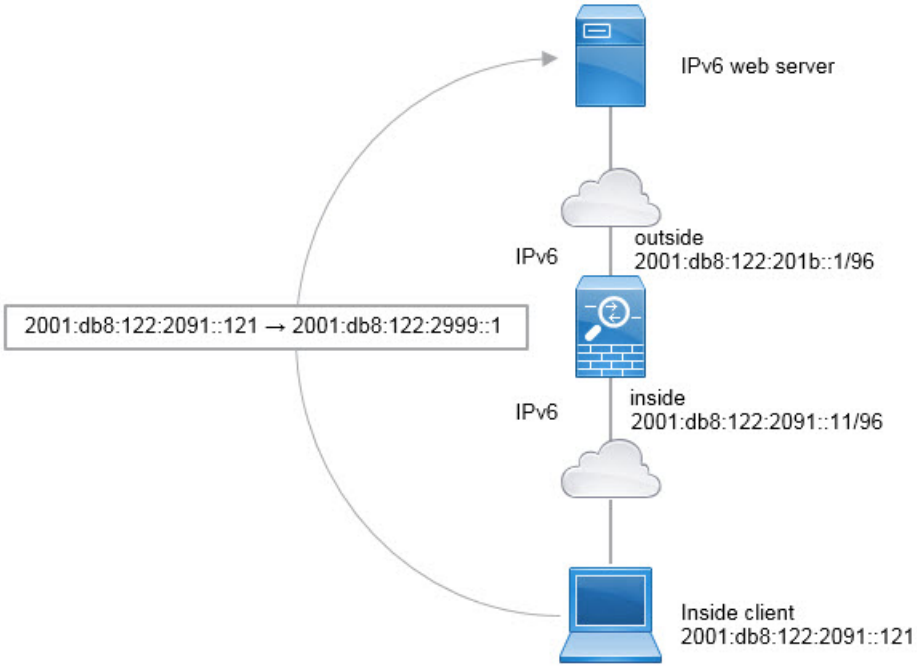
NAT66: translating IPv6 addresses to different IPv6 addresses

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.



Note This example assumes that the inside interface is not a bridge group interface (BVI) but a standard routed interface. If the inside interface is a BVI, you need to duplicate the rules for each member interface.

Procedure

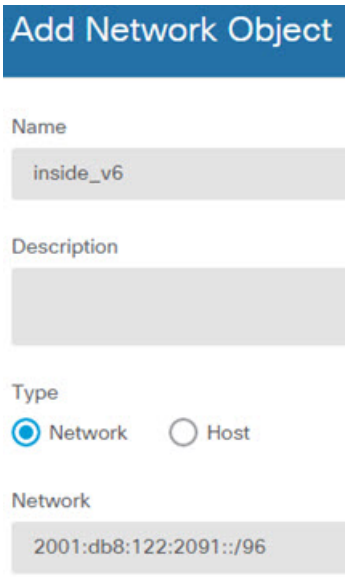
Step 1 Create the network objects that define the inside IPv6 and outside IPv6 NAT networks.

a) Choose **Objects**.

b) Select **Network** from the table of contents and click +.

c) Define the inside IPv6 network.

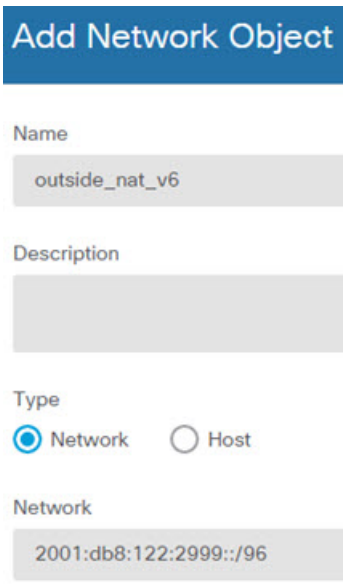
Name the network object (for example, inside_v6), select **Network**, and enter the network address, 2001:db8:122:2091::/96.



d) Click **OK**.

e) Click **+** and define the outside IPv6 NAT network.

Name the network object (for example, outside_nat_v6), select **Network**, and enter the network address 2001:db8:122:2999::/96.



Step 2 Configure the static NAT rule for the inside IPv6 network.

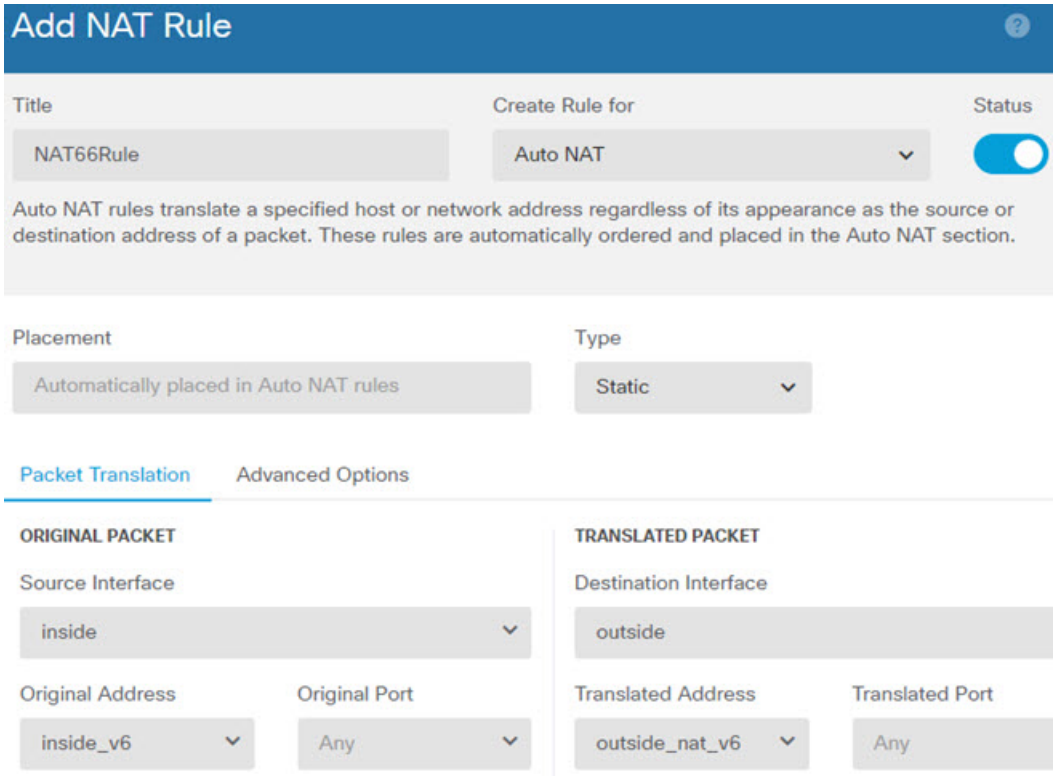
a) Select **Policies** > **NAT**.

b) Click the **+** button.

c) Configure the following properties:

- **Title** = NAT66Rule (or another name of your choosing).
- **Create Rule For** = Auto NAT.

- **Type** = Static.
- **Source Interface** = inside.
- **Destination Interface** = outside.
- **Original Address** = inside_v6 network object.
- **Translated Address** = outside_nat_v6 network object.



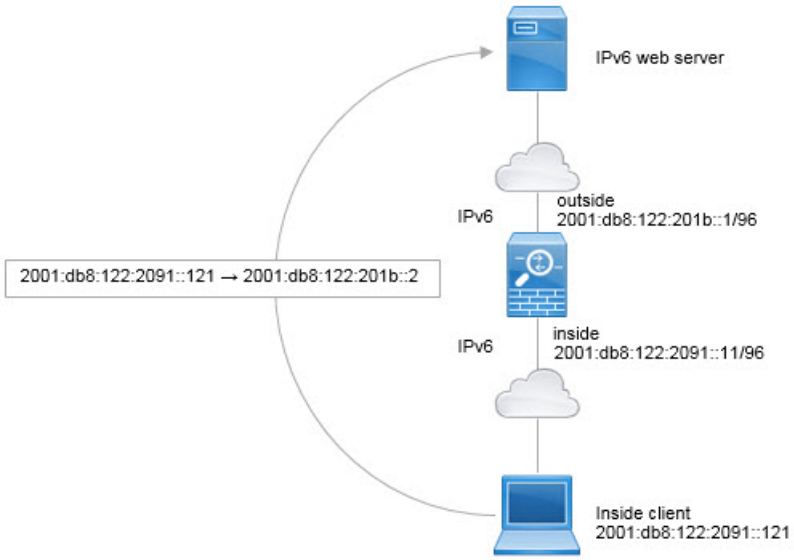
d) Click **OK**.

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.

However, you cannot configure interface PAT using the IPv6 address of an interface using the Firewall Device Manager. Instead, use a single free address on the same network as a dynamic PAT pool.



Note This example assumes that the inside interface is not a bridge group interface (BVI) but a standard routed interface. If the inside interface is a BVI, you need to duplicate the rules for each member interface.

Procedure

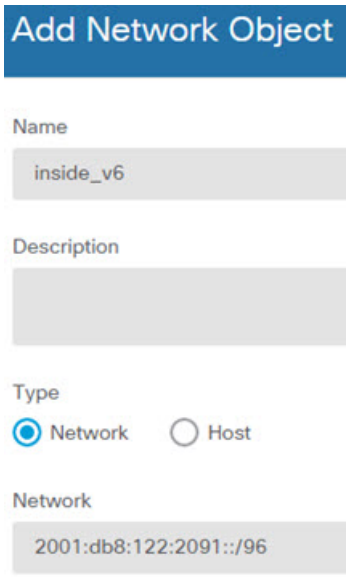
Step 1 Create the network objects that define the inside IPv6 network and the IPv6 PAT address.

a) Choose **Objects**.

b) Select **Network** from the table of contents and click +.

c) Define the inside IPv6 network.

Name the network object (for example, inside_v6), select **Network**, and enter the network address, 2001:db8:122:2091::/96.



d) Click **OK**.

e) Click + and define the outside IPv6 PAT address.

Name the network object (for example, ipv6_pat), select **Host**, and enter the host address 2001:db8:122:201b::2.



Step 2 Configure the dynamic PAT rule for the inside IPv6 network.

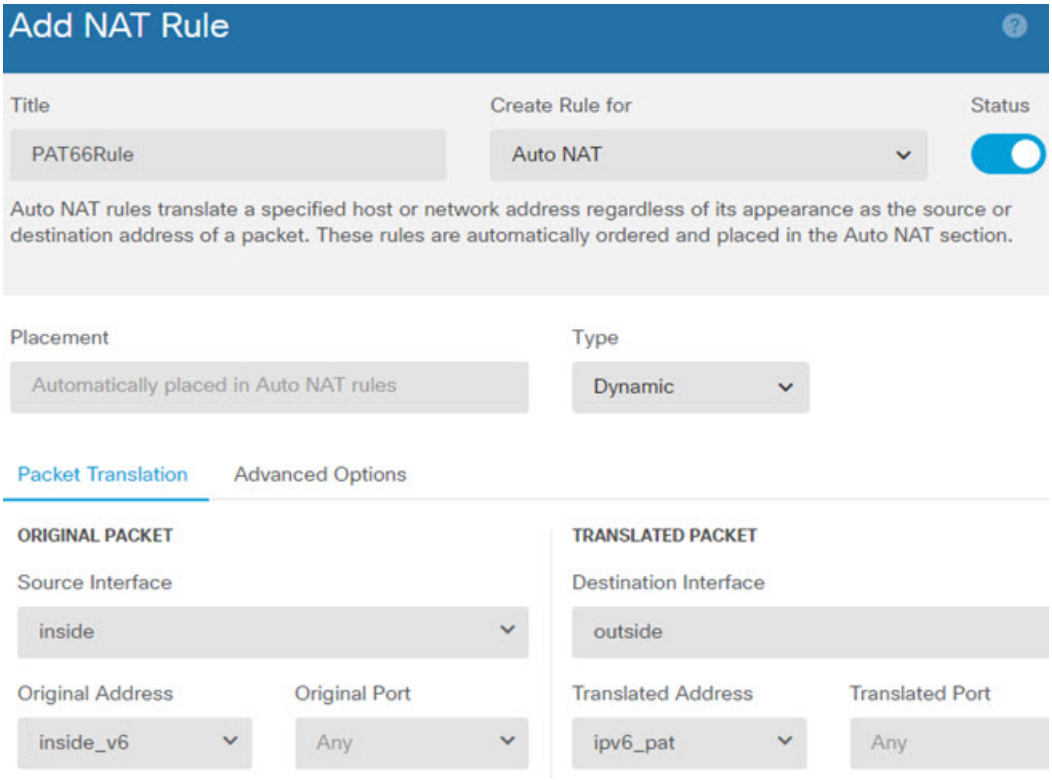
a) Select **Policies** > **NAT**.

b) Click the + button.

c) Configure the following properties:

- **Title** = PAT66Rule (or another name of your choosing).
- **Create Rule For** = Auto NAT.

- **Type** = Dynamic.
- **Source Interface** = inside.
- **Destination Interface** = outside.
- **Original Address** = inside_v6 network object.
- **Translated Address** = ipv6_pat network object.



d) 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 dynamic PAT66 translation to a port on 2001:db8:122:201b::2.

Monitoring NAT

To monitor and troubleshoot NAT connections, open the CLI console or 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. (You cannot use this command in the CLI console.)

Examples for NAT

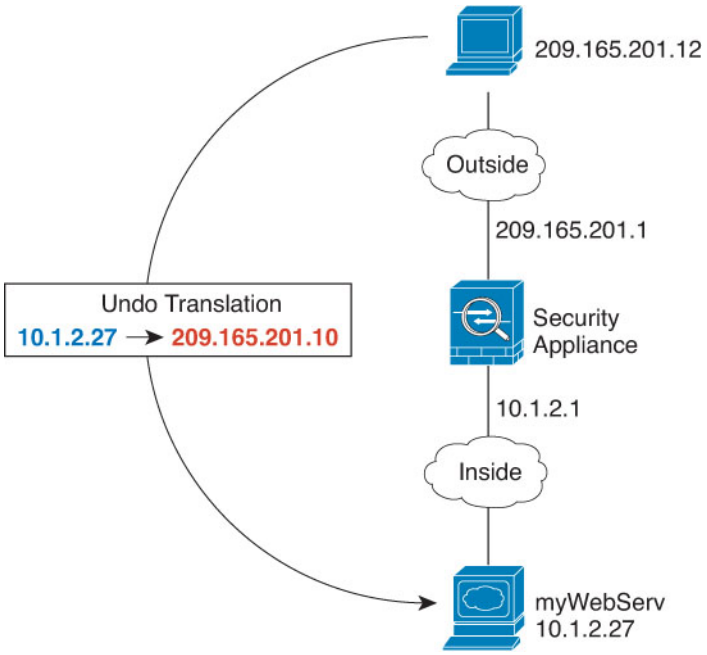
The following topics provide examples for configuring NAT on Threat Defense devices.

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.

Note This example assumes that the inside interface is not a bridge group interface (BVI) but a standard routed interface. If the inside interface is a BVI, select the specific bridge group member interface to which the web server is attached, for example, inside1_3.

Figure 40: Static NAT for an Inside Web Server



Procedure

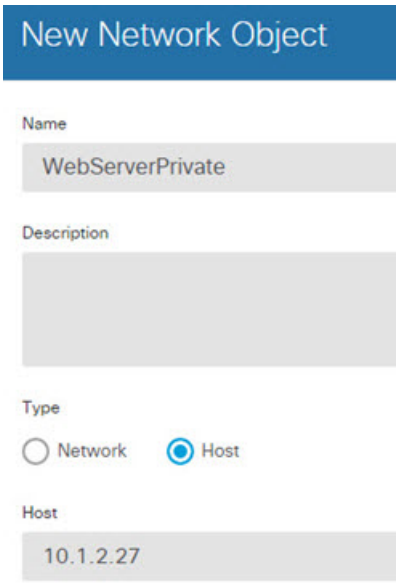
Step 1 Create the network objects that define the server's private and public host addresses.

a) Choose **Objects**.

b) Select **Network** from the table of contents and click +.

c) Define the web server's private address.

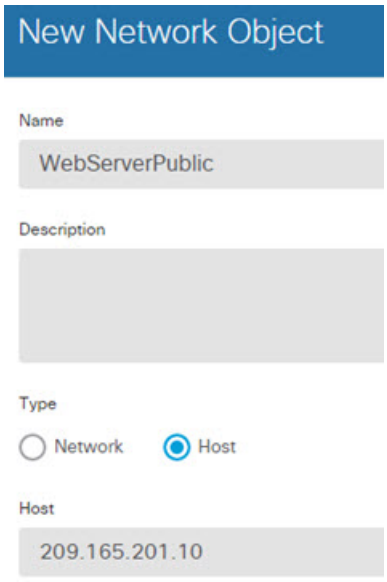
Name the network object (for example, WebServerPrivate), select **Host**, and enter the real host IP address, 10.1.2.27.



d) Click **OK**.

e) Click + and define the public address.

Name the network object (for example, WebServerPublic), select **Host**, and enter the host address 209.165.201.10.



f) Click **OK**.

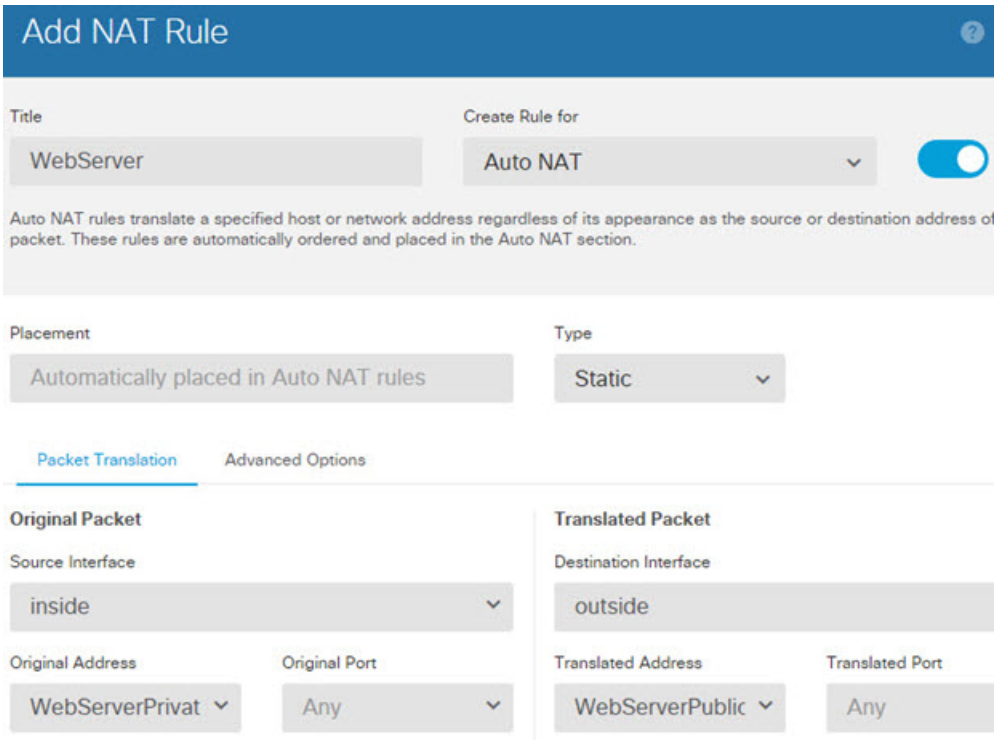
Step 2 Configure static NAT for the object.

a) Select **Policies** > **NAT**.

b) Click the + button.

c) Configure the following properties:

- **Title** = WebServer (or another name of your choosing).
- **Create Rule For** = Auto NAT.
- **Type** = Static.
- **Source Interface** = inside.
- **Destination Interface** = outside.
- **Original Address** = WebServerPrivate network object.
- **Translated Address** = WebServerPublic network object.



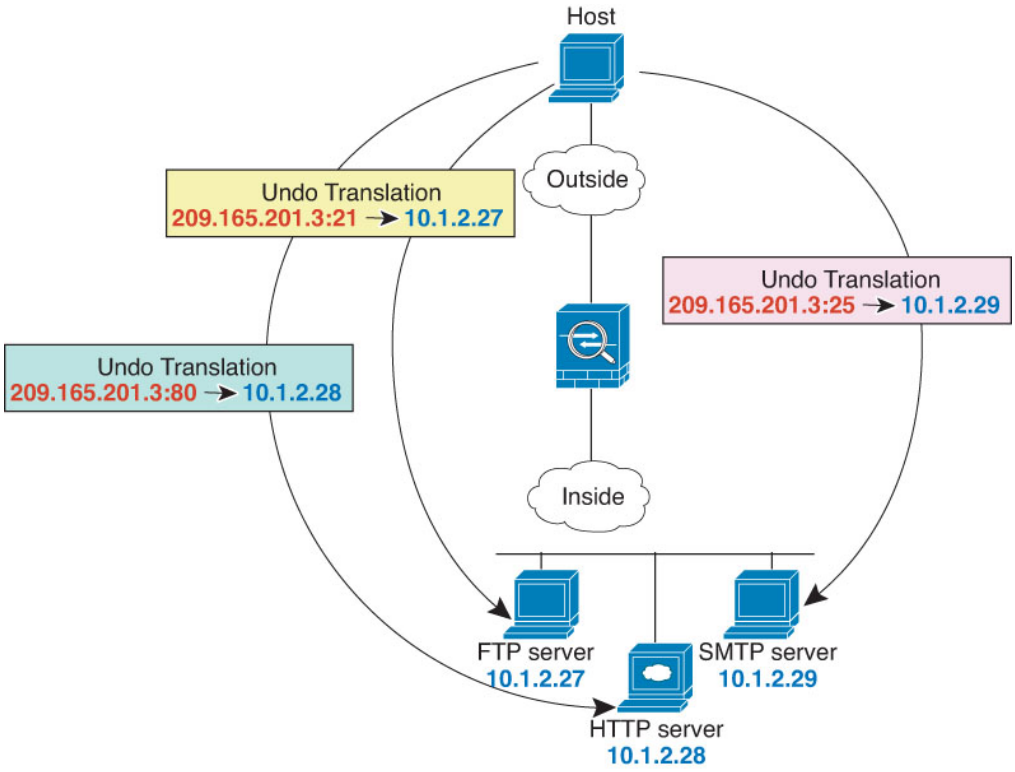
d) Click **OK**.

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.

Note This example assumes that the inside interface is a standard routed interface attached to a switch, with the servers attached to the switch. If your inside interface is a bridge group interface (BVI), and the servers are attached to separate bridge group member interfaces, select the specific member interface to which each server is attached for the corresponding rule. For example, the rules might have inside1_2, inside1_3, and inside1_4 for the source interface rather than inside.

Figure 41: Static NAT-with-Port-Translation



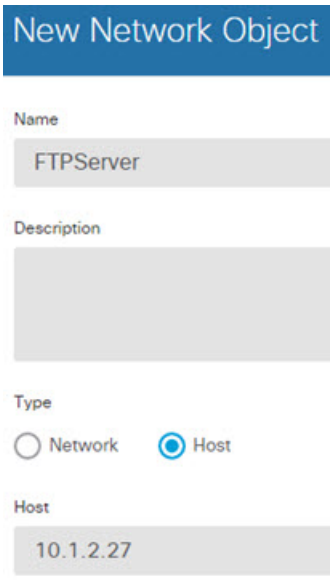
Procedure

Step 1 Create a network object for the FTP server.

a) Choose **Objects**.

b) Select **Network** from the table of contents and click +.

c) Name the network object (for example, FTPserver), select **Host**, and enter the real IP address for the FTP server, 10.1.2.27.

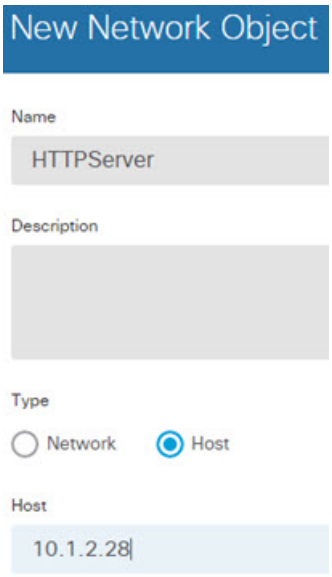


d) Click **OK**.

Step 2 Create a network object for the HTTP server.

a) Click +.

b) Name the network object (for example, HTTPserver), select **Host**, and enter the host address 10.1.2.28.

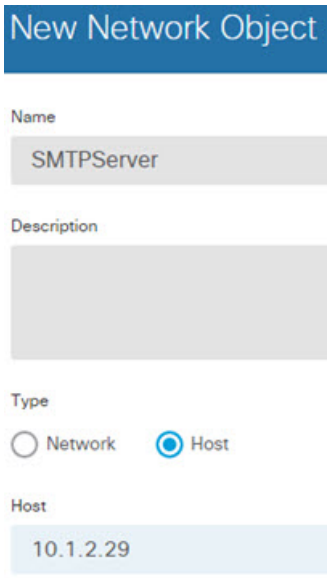


c) Click **OK**.

Step 3 Create a network object for the SMTP server.

a) Click +.

b) Name the network object (for example, SMTPserver), select **Host**, and enter the host address 10.1.2.29.

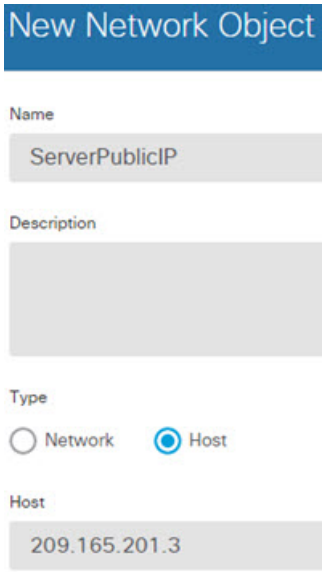


c) Click **OK**.

Step 4 Create a network object for the public IP address used for the three servers.

a) Click +.

b) Name the network object (for example, ServerPublicIP), select **Host**, and enter the host address 209.165.201.3.



c) Click **OK**.

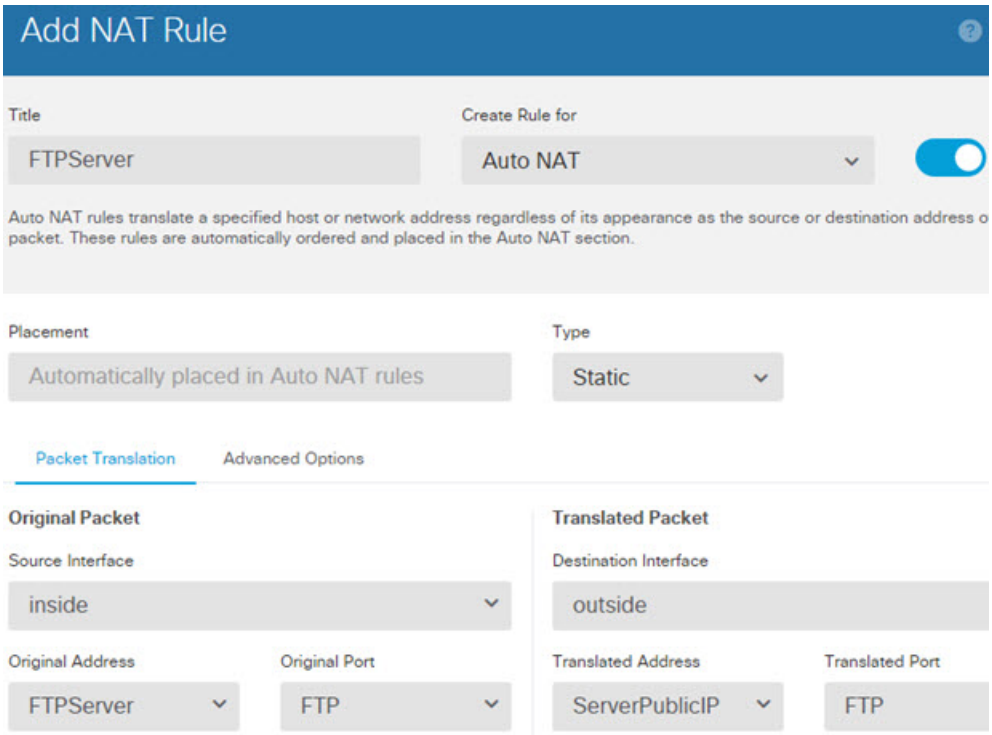
Step 5 Configure static NAT with port translation for the FTP server, mapping the FTP port to itself.

a) Select **Policies** > **NAT**.

b) Click the + button.

c) Configure the following properties:

- **Title** = FTPServer (or another name of your choosing).
- **Create Rule For** = Auto NAT.
- **Type** = Static.
- **Source Interface** = inside.
- **Destination Interface** = outside.
- **Original Address** = FTPserver network object.
- **Translated Address** = ServerPublicIP network object.
- **Original Port** = FTP port object.
- **Translated Port** = FTP port object.



d) Click **OK**.

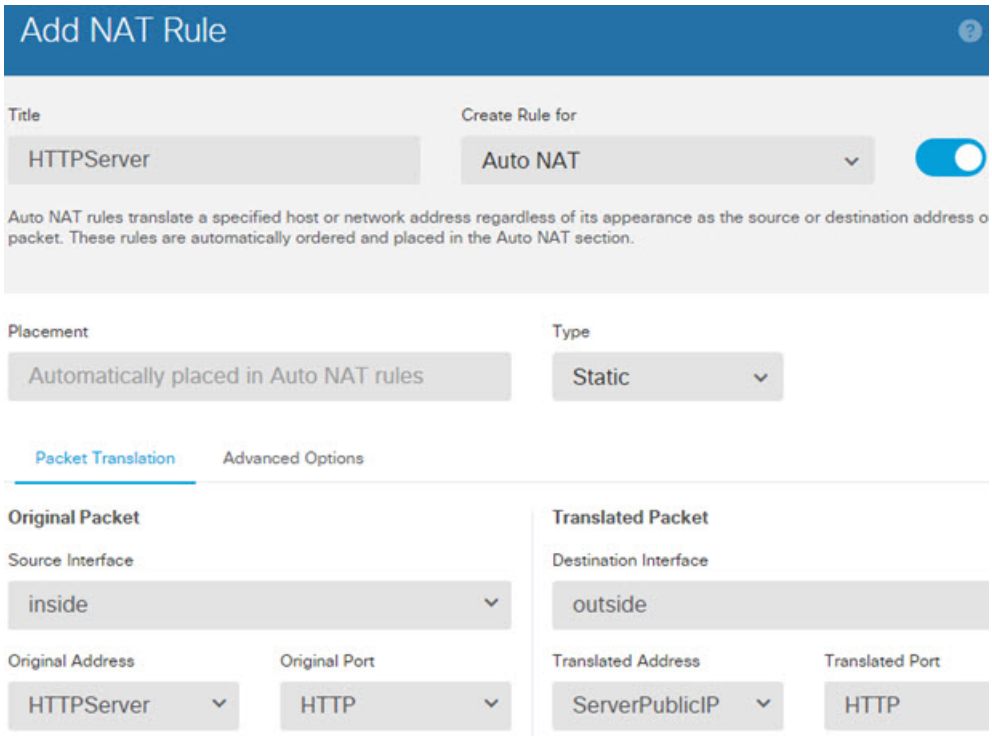
Step 6 Configure static NAT with port translation for the HTTP server, mapping the HTTP port to itself.

a) Click the + button.

b) Configure the following properties:

- **Title** = HTTPServer (or another name of your choosing).
- **Create Rule For** = Auto NAT.
- **Type** = Static.
- **Source Interface** = inside.
- **Destination Interface** = outside.

- **Original Address** = HTTPserver network object.
- **Translated Address** = ServerPublicIP network object.
- **Original Port** = HTTP port object.
- **Translated Port** = HTTP port object.



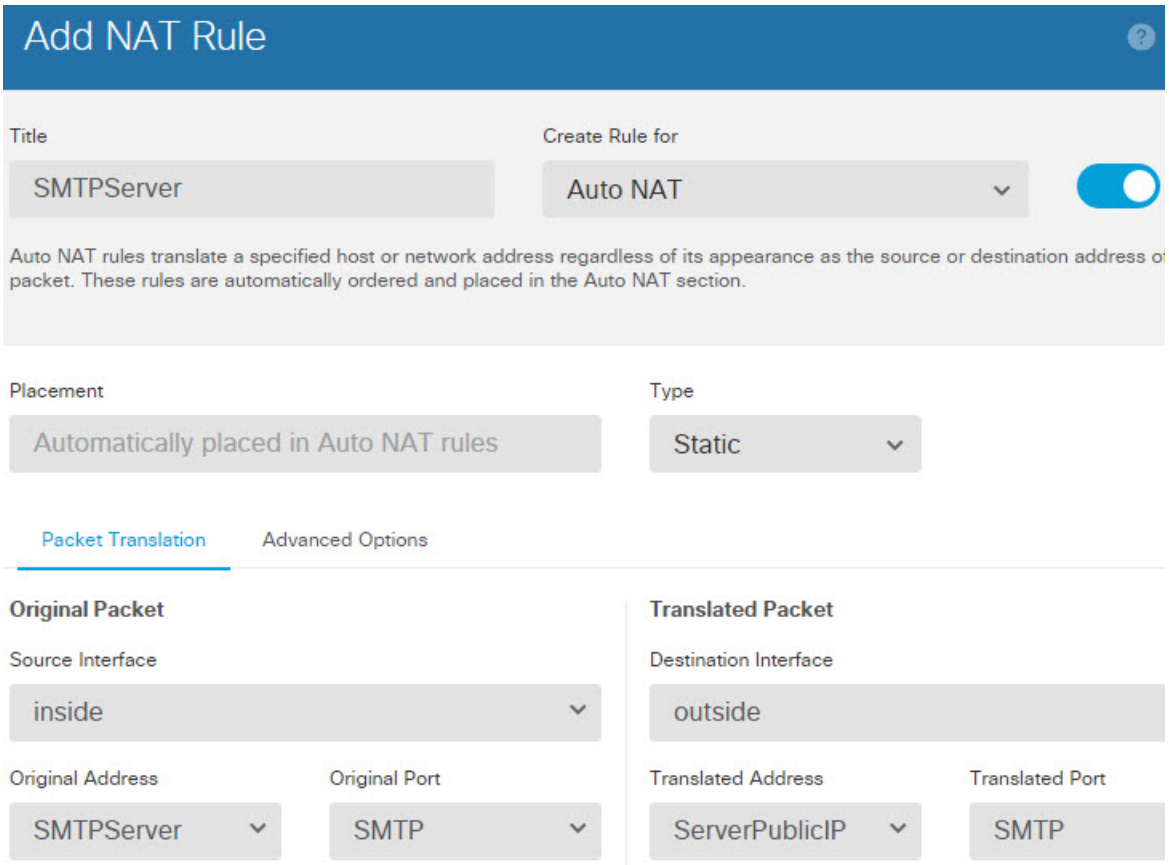
c) Click **OK**.

Step 7 Configure static NAT with port translation for the SMTP server, mapping the SMTP port to itself.

a) Click the + button.

b) Configure the following properties:

- **Title** = SMTPServer (or another name of your choosing).
- **Create Rule For** = Auto NAT.
- **Type** = Static.
- **Source Interface** = inside.
- **Destination Interface** = outside.
- **Original Address** = SMTPserver network object.
- **Translated Address** = ServerPublicIP network object.
- **Original Port** = SMTP port object.
- **Translated Port** = SMTP port object.



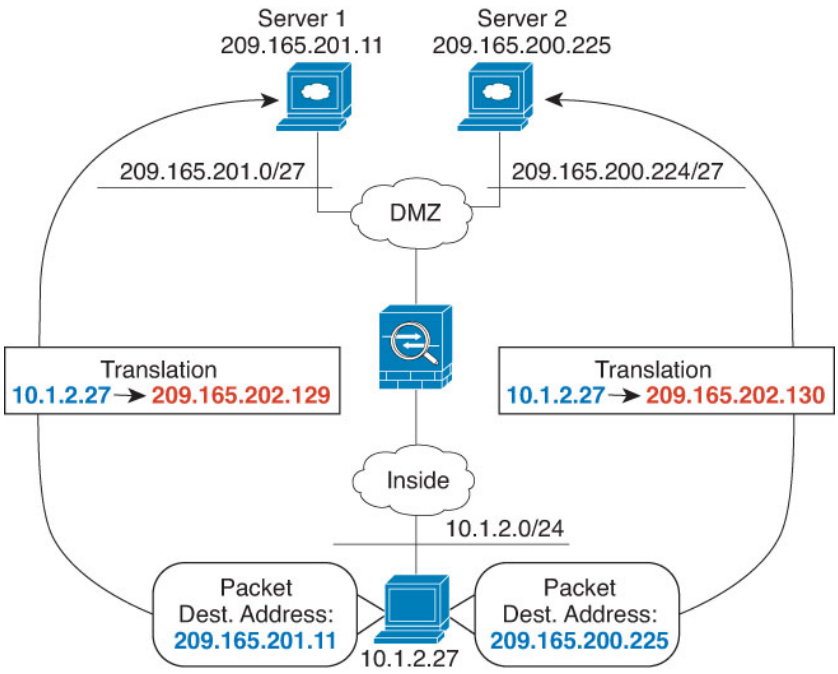
c) Click **OK**.

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*.

Note This example assumes that the inside interface is a standard routed interface attached to a switch, with the servers attached to the switch. If your inside interface is a bridge group interface (BVI), and the servers are attached to separate bridge group member interfaces, select the specific member interface to which each server is attached for the corresponding rule. For example, the rules might have inside1_2 and inside1_3 for the source interface rather than inside.

Figure 42: Manual NAT with Different Destination Addresses



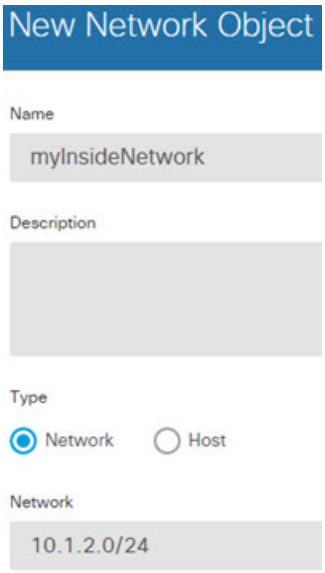
Procedure

Step 1 Create a network object for the inside network.

a) Choose **Objects**.

b) Select **Network** from the table of contents and click +.

c) Name the network object (for example, myInsideNetwork), select **Network**, and enter the real network address, 10.1.2.0/24.

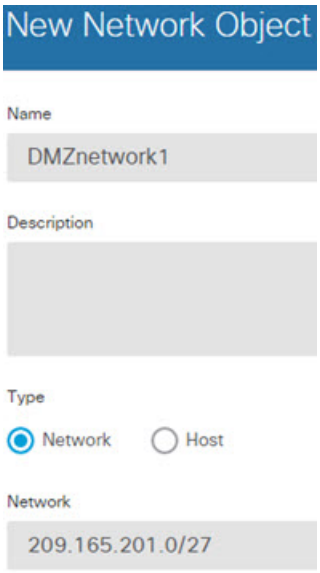


d) Click **OK**.

Step 2 Create a network object for the DMZ network 1.

a) Click +.

b) Name the network object (for example, DMZnetwork1), select **Network**, and enter the network address 209.165.201.0/27 (subnet mask of 255.255.255.224).

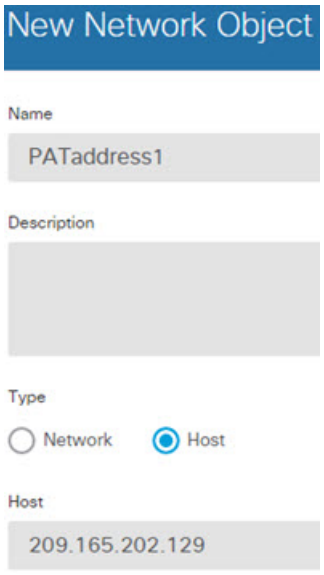


c) Click **OK**.

Step 3 Create a network object for the PAT address for DMZ network 1.

a) Click +.

b) Name the network object (for example, PATaddress1), select **Host**, and enter the host address 209.165.202.129.

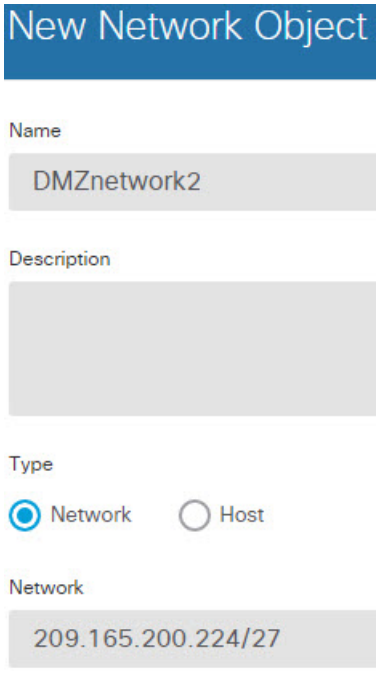


c) Click **OK**.

Step 4 Create a network object for the DMZ network 2.

a) Click +.

b) Name the network object (for example, DMZnetwork2), select **Network**, and enter the network address 209.165.200.224/27 (subnet mask of 255.255.255.224).

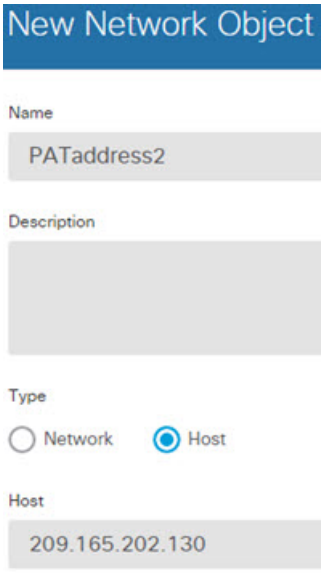


c) Click **OK**.

Step 5 Create a network object for the PAT address for DMZ network 2.

a) Click +.

b) Name the network object (for example, PATaddress2), select **Host**, and enter the host address 209.165.202.130.



c) Click **OK**.

Step 6 Configure dynamic manual PAT for DMZ network 1.

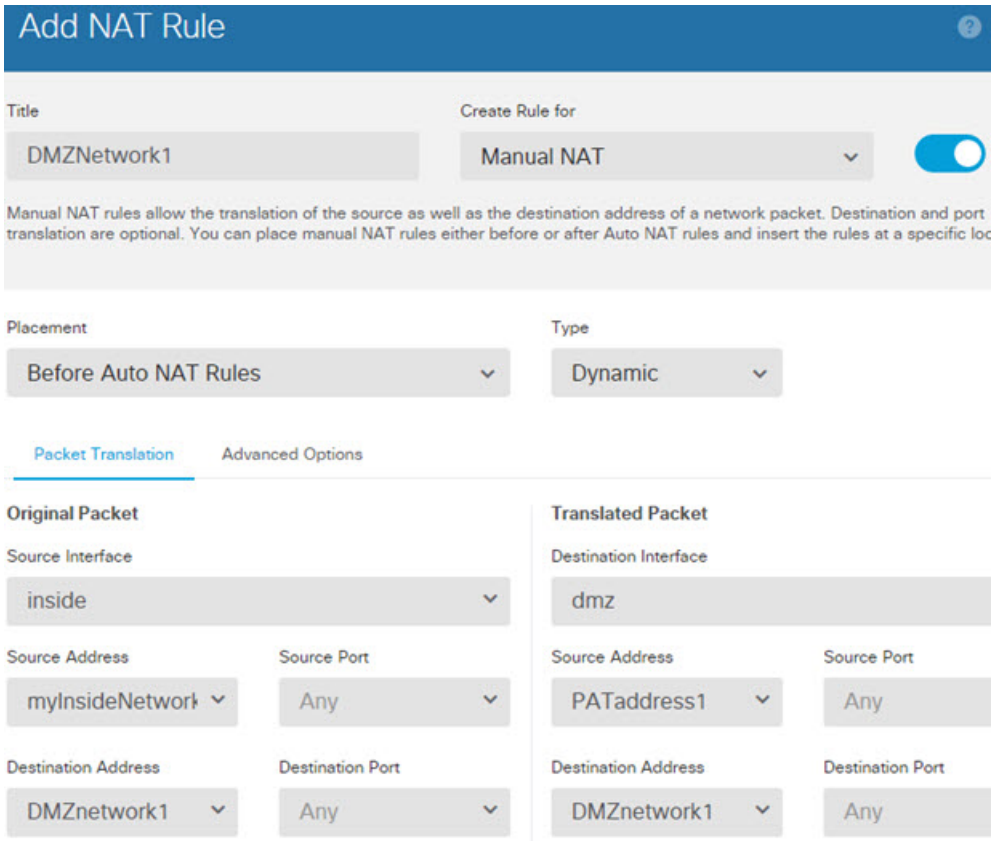
a) Select **Policies** > **NAT**.

b) Click the + button.

c) Configure the following properties:

- **Title** = DMZNetwork1 (or another name of your choosing).
- **Create Rule For** = Manual NAT.
- **Type** = Dynamic.
- **Source Interface** = inside.
- **Destination Interface** = dmz.
- **Original Source Address** = myInsideNetwork network object.
- **Translated Source Address** = PATaddress1 network object.
- **Original Destination Address** = DMZnetwork1 network object.
- **Translated Destination Address** = 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.



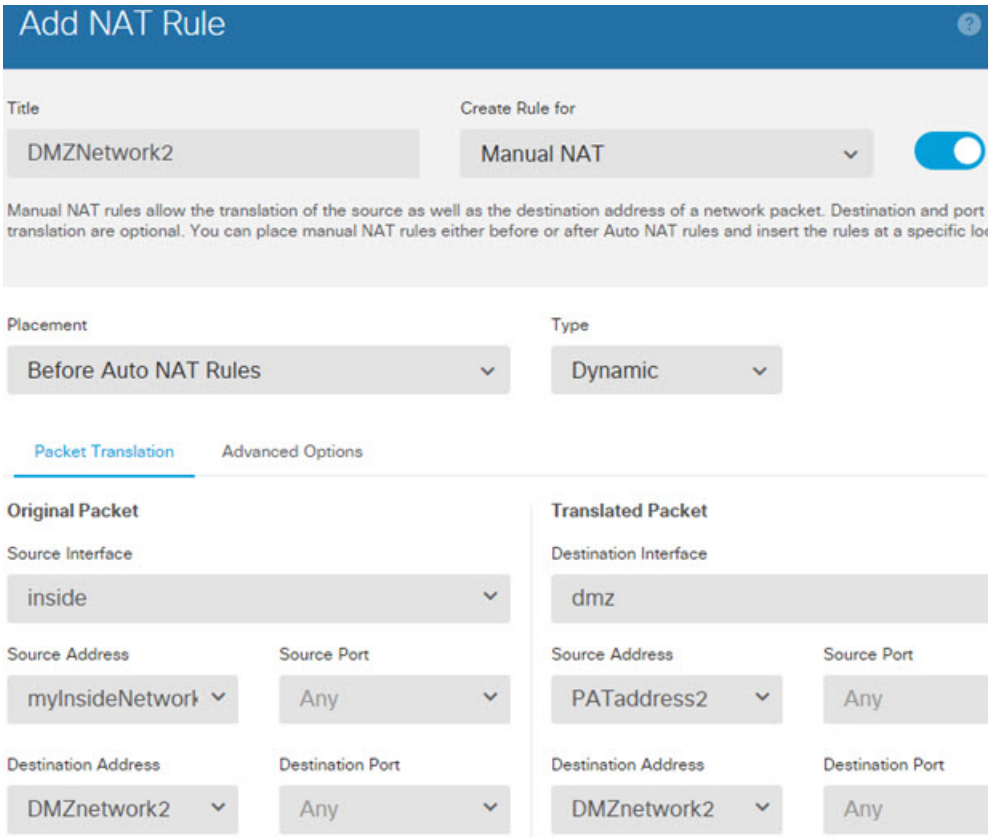
d) Click **OK**.

Step 7 Configure dynamic manual PAT for DMZ network 2.

a) Click the + button.

b) Configure the following properties:

- **Title** = DMZNetwork2 (or another name of your choosing).
- **Create Rule For** = Manual NAT.
- **Type** = Dynamic.
- **Source Interface** = inside.
- **Destination Interface** = dmz.
- **Original Source Address** = myInsideNetwork network object.
- **Translated Source Address** = PATaddress2 network object.
- **Original Destination Address** = DMZnetwork2 network object.
- **Translated Destination Address** = DMZnetwork2 network object.



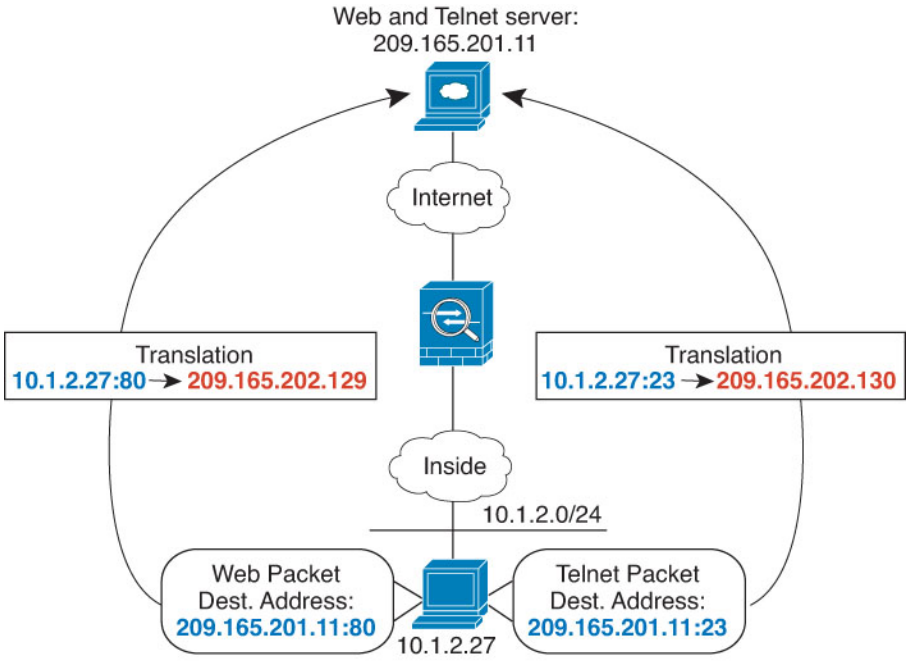
c) Click **OK**.

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*.

Note This example assumes that the inside interface is a standard routed interface attached to a switch, with the server attached to the switch. If your inside interface is a bridge group interface (BVI), and the server is attached to a bridge group member interface, select the specific member interface to which the server is attached. For example, the rule might have inside1_2 for the source interface rather than inside.

Figure 43: Manual NAT with Different Destination Ports



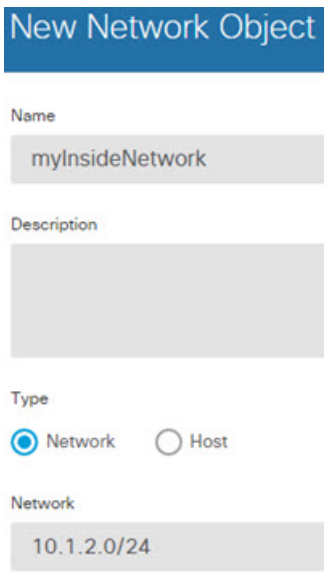
Procedure

Step 1 Create a network object for the inside network.

a) Choose **Objects**.

b) Select **Network** from the table of contents and click +.

c) Name the network object (for example, myInsideNetwork), select **Network**, and enter the real network address, 10.1.2.0/24.

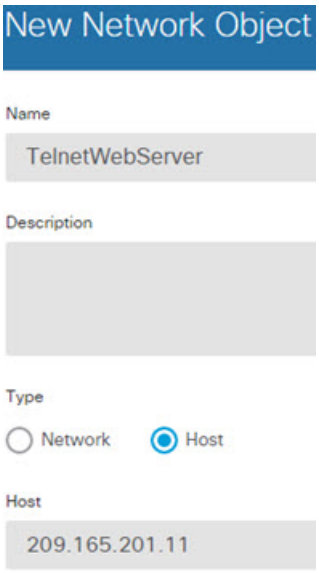


d) Click **OK**.

Step 2 Create a network object for the Telnet/Web server.

a) Click +.

b) Name the network object (for example, TelnetWebServer), select **Host**, and enter the host address 209.165.201.11.

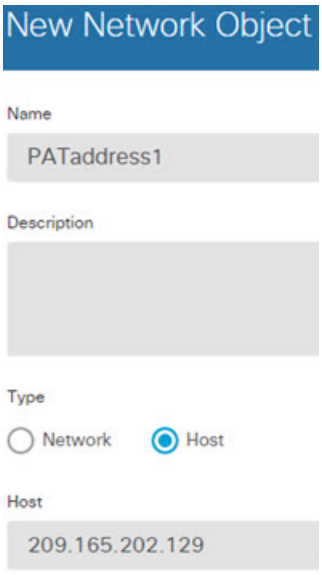


c) Click **OK**.

Step 3 Create a network object for the PAT address when using Telnet.

a) Click +.

b) Name the network object (for example, PATaddress1), select **Host**, and enter the host address 209.165.202.129.

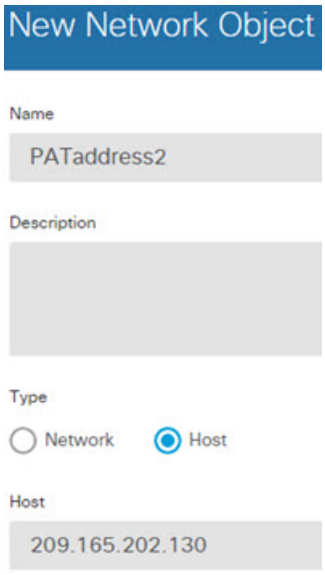


c) Click **OK**.

Step 4 Create a network object for the PAT address when using HTTP.

a) Click +.

b) Name the network object (for example, PATaddress2), select **Host**, and enter the host address 209.165.202.130.



c) Click **OK**.

Step 5 Configure dynamic manual PAT for Telnet access.

a) Select **Policies** > **NAT**.

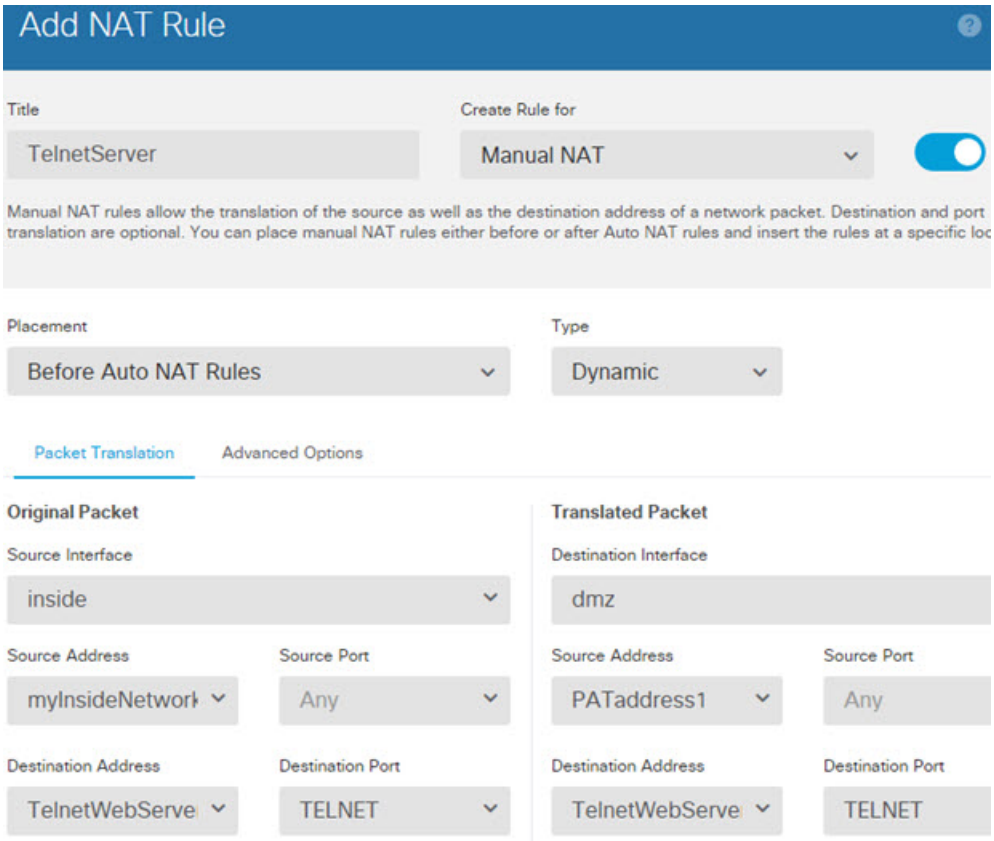
b) Click the + button.

c) Configure the following properties:

- **Title** = TelnetServer (or another name of your choosing).
- **Create Rule For** = Manual NAT.
- **Type** = Dynamic.
- **Source Interface** = inside.
- **Destination Interface** = dmz.
- **Original Source Address** = myInsideNetwork network object.
- **Translated Source Address**= PATaddress1 network object.
- **Original Destination Address** = TelnetWebServer network object.
- **Translated Destination Address** = TelnetWebServer network object.
- **Original Destination Port** = TELNET port object.
- **Translated Destination Port** = TELNET port object.

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.



d) Click **OK**.

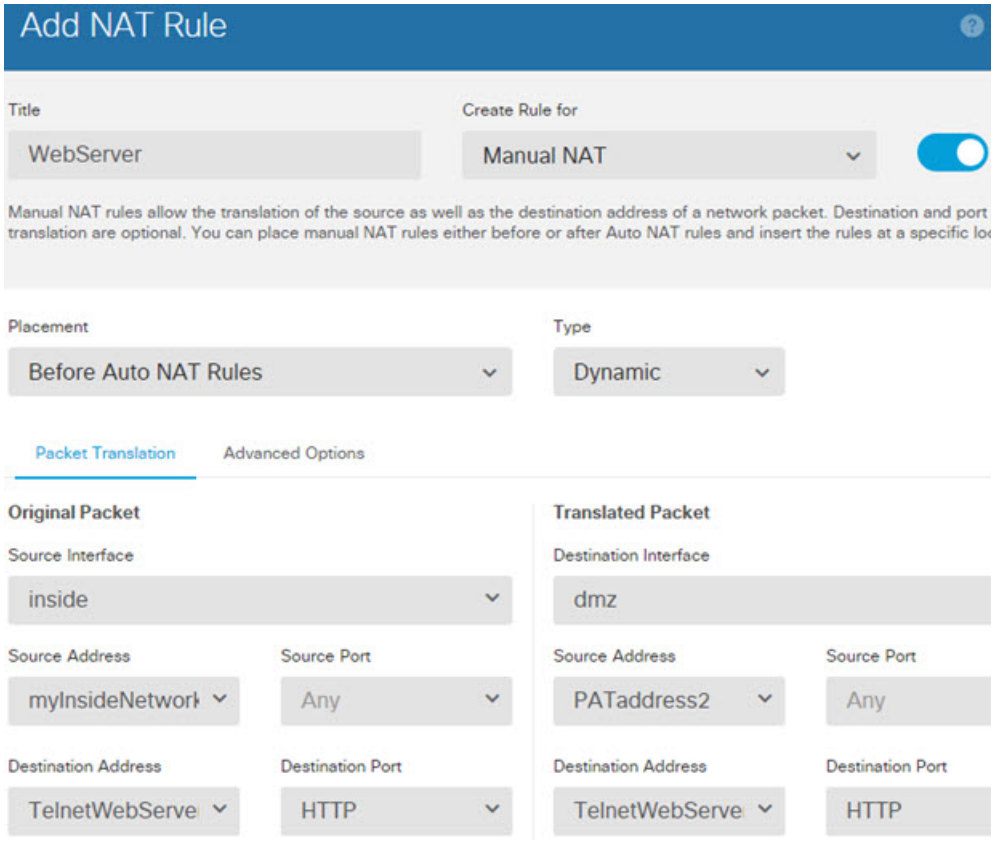
Step 6 Configure dynamic manual PAT for web access.

a) Click the + button.

b) Configure the following properties:

- **Title** = WebServer (or another name of your choosing).
- **Create Rule For** = Manual NAT.
- **Type** = Dynamic.
- **Source Interface** = inside.
- **Destination Interface** = dmz.
- **Original Source Address** = myInsideNetwork network object.
- **Translated Source Address** = PATaddress2 network object.
- **Original Destination Address** = TelnetWebServer network object.
- **Translated Destination Address** = TelnetWebServer network object.

- **Original Destination Port** = HTTP port object.
- **Translated Destination Port** = HTTP port object.



c) Click **OK**.

Rewriting DNS queries and responses using NAT

You might need to configure the Firewall Threat Defense 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:

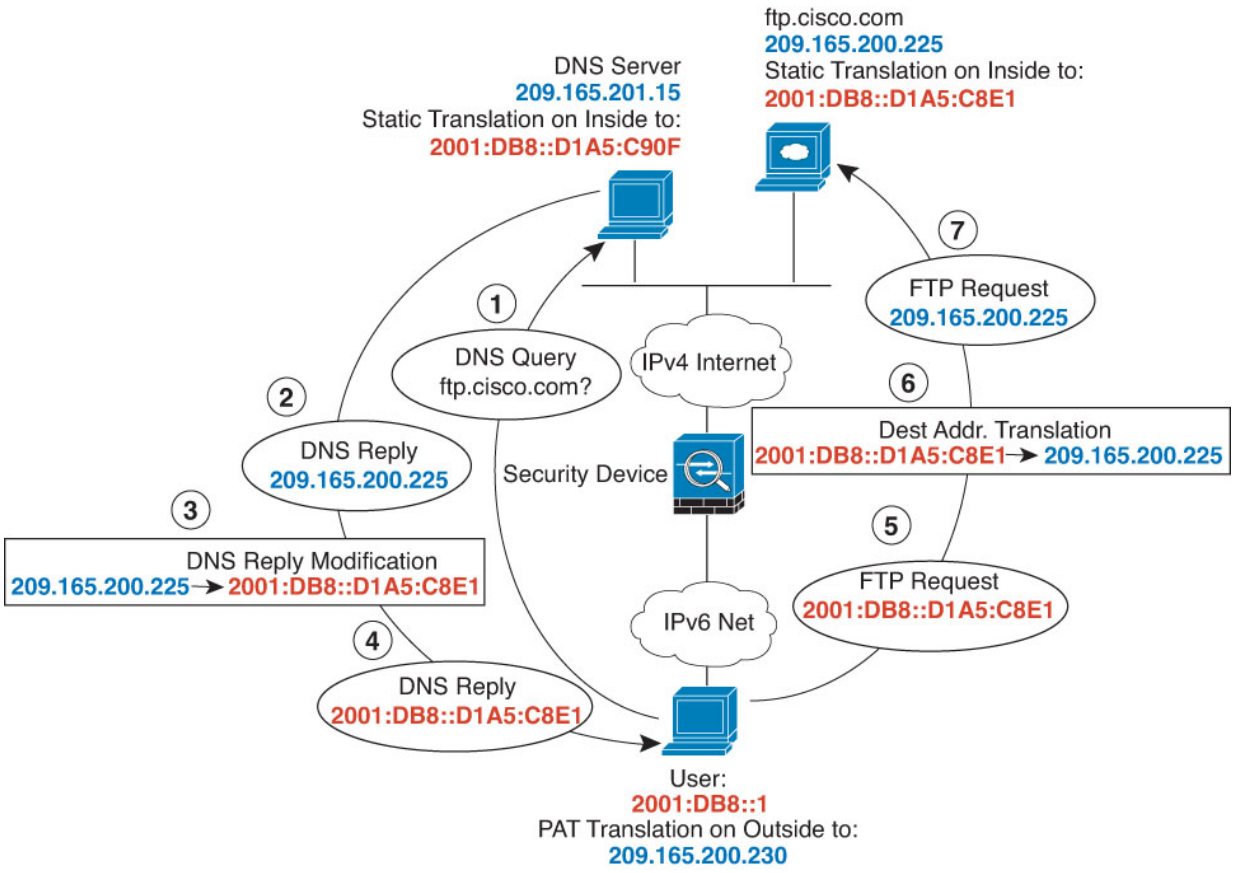
- DNS rewrite is not applicable for PAT because multiple PAT rules are applicable for each A or AAAA record, and the PAT rule to use is ambiguous.
- If you configure a manual NAT rule, you cannot configure DNS modification if you specify the destination address as well as the source address. These kinds 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.

DNS 64 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.



Note This example assumes that the inside interface is not a bridge group interface (BVI) but a standard routed interface. If the inside interface is a BVI, you need to duplicate the rules for each member interface.

Procedure

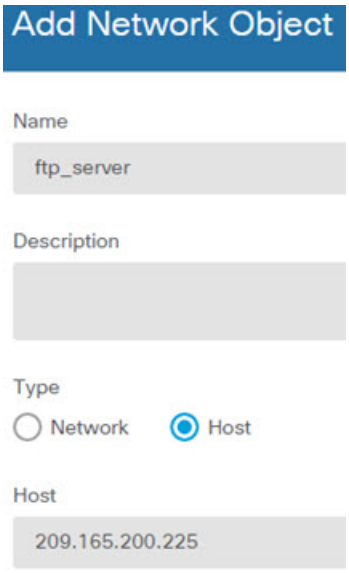
Step 1 Create the network objects for the FTP server, DNS server, inside network, and PAT pool.

a) Choose **Objects**.

b) Select **Network** from the table of contents and click +.

c) Define the real FTP server address.

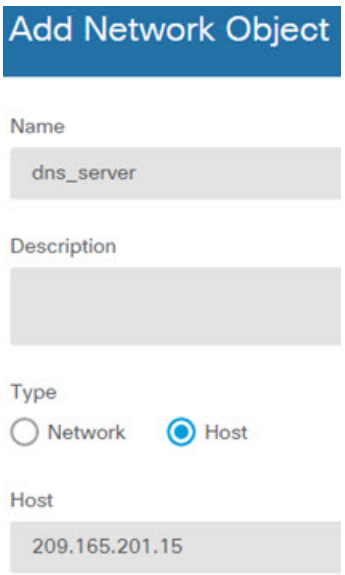
Name the network object (for example, ftp_server), select **Host**, and enter the real host IP address, 209.165.200.225.



d) Click **OK**.

e) Click + and define the DNS server's real address.

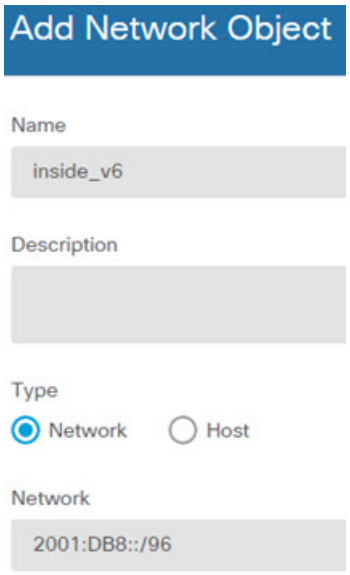
Name the network object (for example, dns_server), select **Host**, and enter the host address 209.165.201.15.



f) Click **OK**.

g) Click + and define the inside IPv6 network.

Name the network object (for example, inside_v6), select **Network**, and enter the network address, 2001:DB8::/96.



h) Click **OK**.

i) Click + and define the IPv4 PAT address for the inside IPv6 network.

Name the network object (for example, ipv4_pat), select **Host**, and enter the host address, 209.165.200.230.



j) Click **OK**.

Step 2 Configure the static NAT rule with DNS modification for the FTP server.

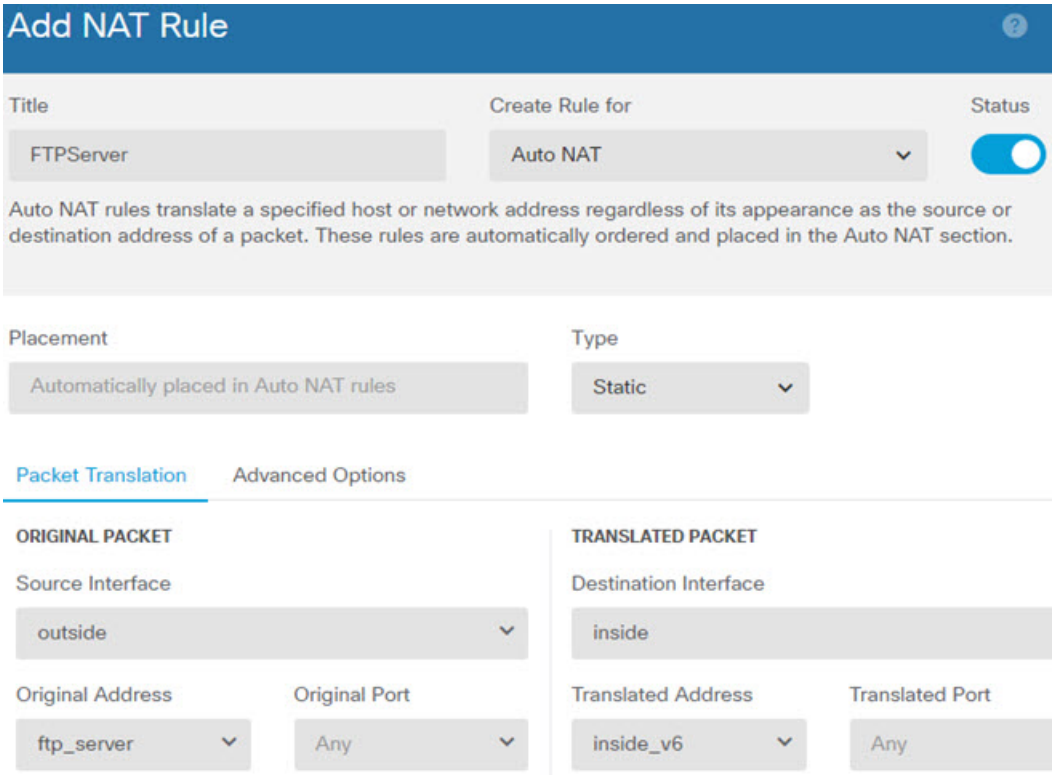
a) Select **Policies** > **NAT**.

b) Click the + button.

c) Configure the following properties:

- **Title** = FTPServer (or another name of your choosing).
- **Create Rule For** = Auto NAT.

- **Type** = Static.
- **Source Interface** = outside.
- **Destination Interface** = inside.
- **Original Address** = ftp_server network object.
- **Translated Address** = inside_v6 network object. Because the IPv4 embedded address method is used when converting IPv4 to IPv6 addresses, 209.165.200.225 is converted to the IPv6 equivalent D1A5:C8E1 and the network prefix is added to get the full address, 2001:DB8::D1A5:C8E1.
- On the **Advanced Options** tab, select **Translate DNS replies that match this rule**.



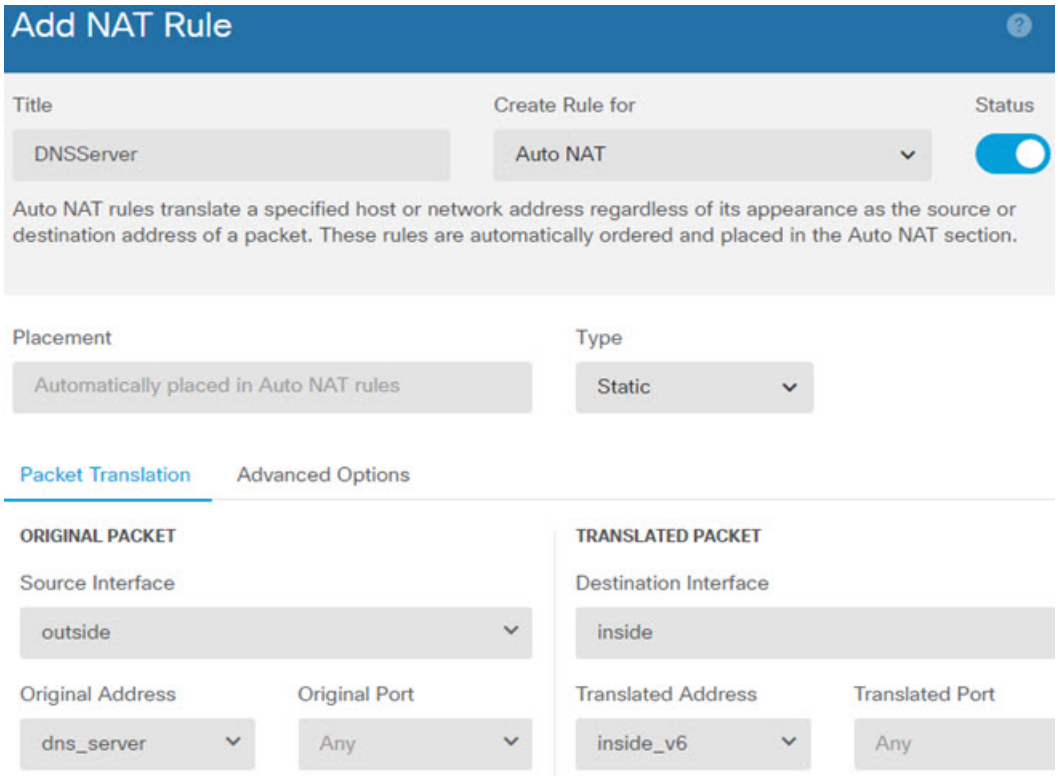
d) Click **OK**.

Step 3 Configure the static NAT rule for the DNS server.

a) Select **Policies** > **NAT**.
b) Click the + button.
c) Configure the following properties:

- **Title** = DNSServer (or another name of your choosing).
- **Create Rule For** = Auto NAT.
- **Type** = Static.
- **Source Interface** = outside.
- **Destination Interface** = inside.

- **Original Address** = dns_server network object.
- **Translated Address** = inside_v6 network object. Because the IPv4 embedded address method is used when converting IPv4 to IPv6 addresses, 209.165.201.15 is converted to the IPv6 equivalent D1A5:C90F and the network prefix is added to get the full address, 2001:DB8::D1A5:C90F.



d) Click **OK**.

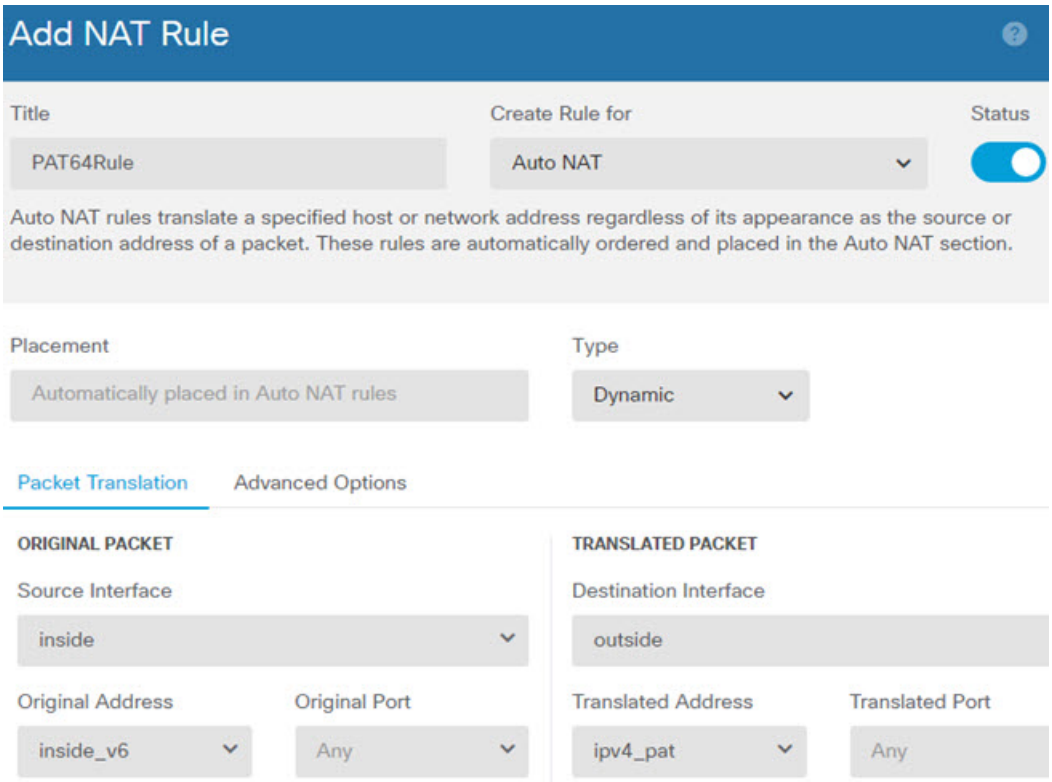
Step 4 Configure the dynamic PAT rule for the inside IPv6 network.

a) Select **Policies** > **NAT**.

b) Click the + button.

c) Configure the following properties:

- **Title** = PAT64Rule (or another name of your choosing).
- **Create Rule For** = Auto NAT.
- **Type** = Dynamic.
- **Source Interface** = inside.
- **Destination Interface** = outside.
- **Original Address** = inside_v6 network object.
- **Translated Address** = ipv4_pat network object.



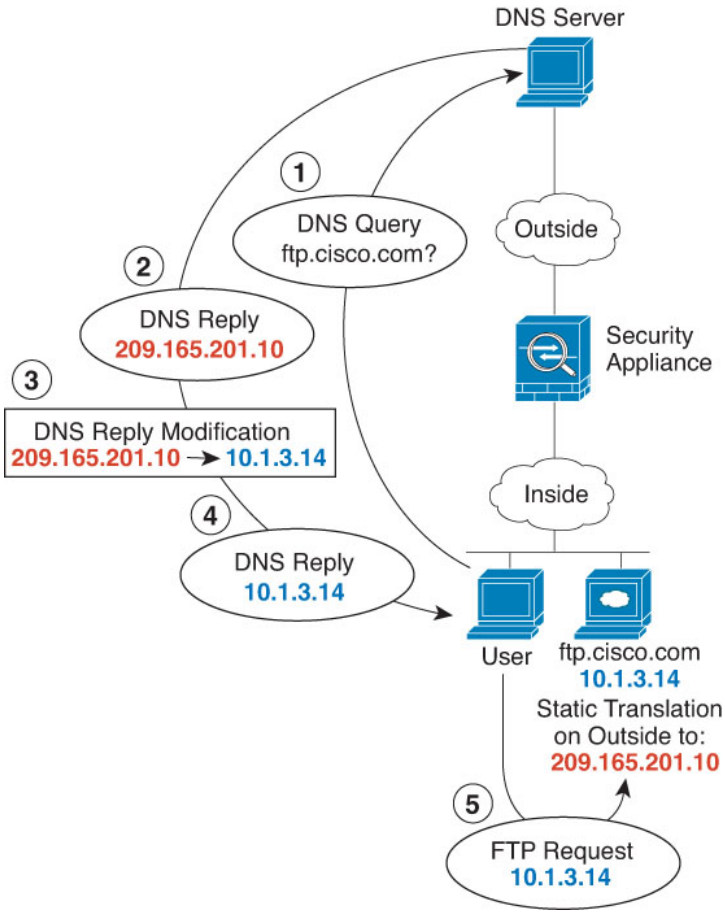
d) Click **OK**.

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.



Note This example assumes that the inside interface is not a bridge group interface (BVI) but a standard routed interface. If the inside interface is a BVI, you need to duplicate the rules for each member interface.

Procedure

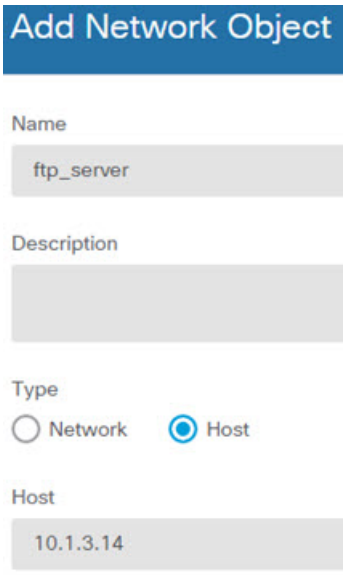
Step 1 Create the network objects for the FTP server.

a) Choose **Objects**.

b) Select **Network** from the table of contents and click +.

c) Define the real FTP server address.

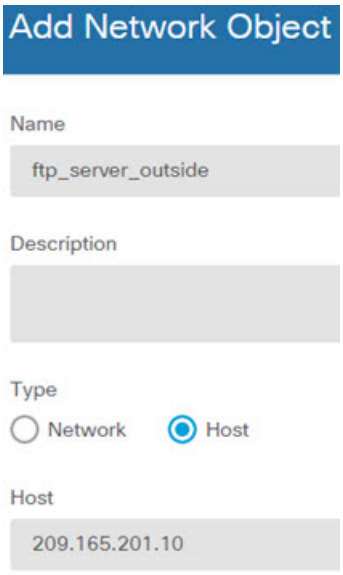
Name the network object (for example, ftp_server), select **Host**, and enter the real host IP address, 10.1.3.14.



d) Click **OK**.

e) Click + and define the FTP server's translated address.

Name the network object (for example, ftp_server_outside), select **Host**, and enter the host address 209.165.201.10.



Step 2 Configure the static NAT rule with DNS modification for the FTP server.

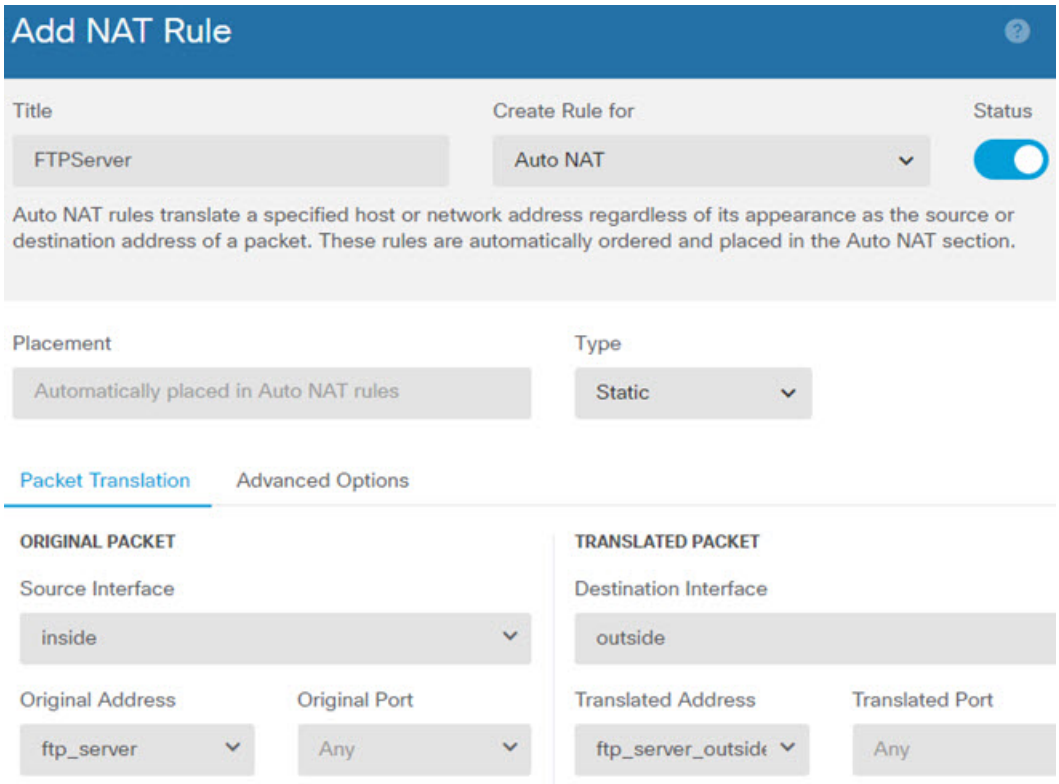
a) Select **Policies** > **NAT**.

b) Click the + button.

c) Configure the following properties:

- **Title** = FTPServer (or another name of your choosing).
- **Create Rule For** = Auto NAT.

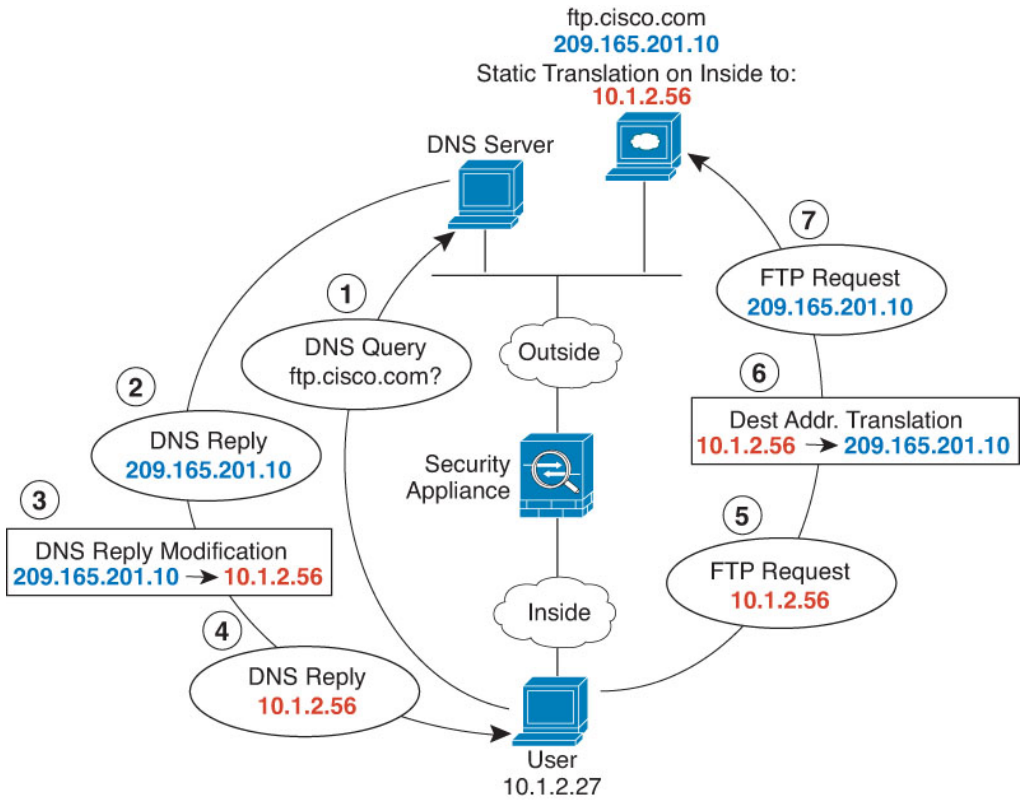
- **Type** = Static.
- **Source Interface** = inside.
- **Destination Interface** = outside.
- **Original Address** = ftp_server network object.
- **Translated Address** = ftp_server_outside network object.
- On the **Advanced Options** tab, select **Translate DNS replies that match this rule**.



d) 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.



Note This example assumes that the inside interface is not a bridge group interface (BVI) but a standard routed interface. If the inside interface is a BVI, you need to duplicate the rules for each member interface.

Procedure

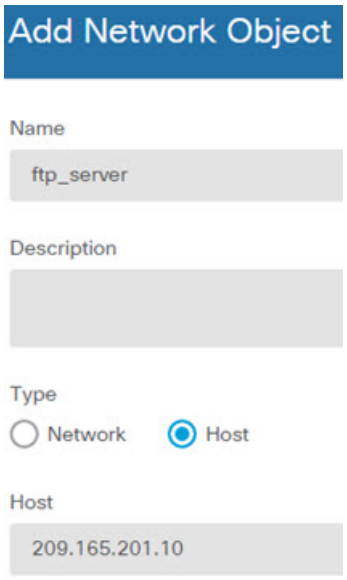
Step 1 Create the network objects for the FTP server.

a) Choose **Objects**.

b) Select **Network** from the table of contents and click +.

c) Define the real FTP server address.

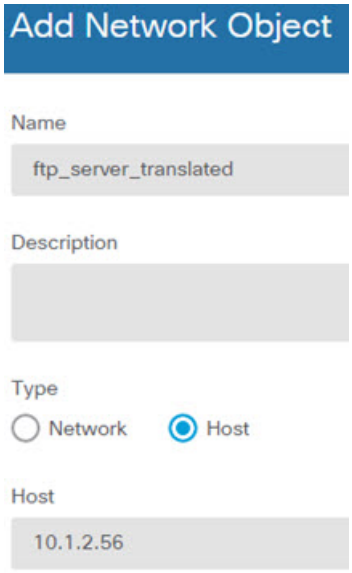
Name the network object (for example, ftp_server), select **Host**, and enter the real host IP address, 209.165.201.10.



d) Click **OK**.

e) Click + and define the FTP server's translated address.

Name the network object (for example, ftp_server_translated), select **Host**, and enter the host address 10.1.2.56.



Step 2 Configure the static NAT rule with DNS modification for the FTP server.

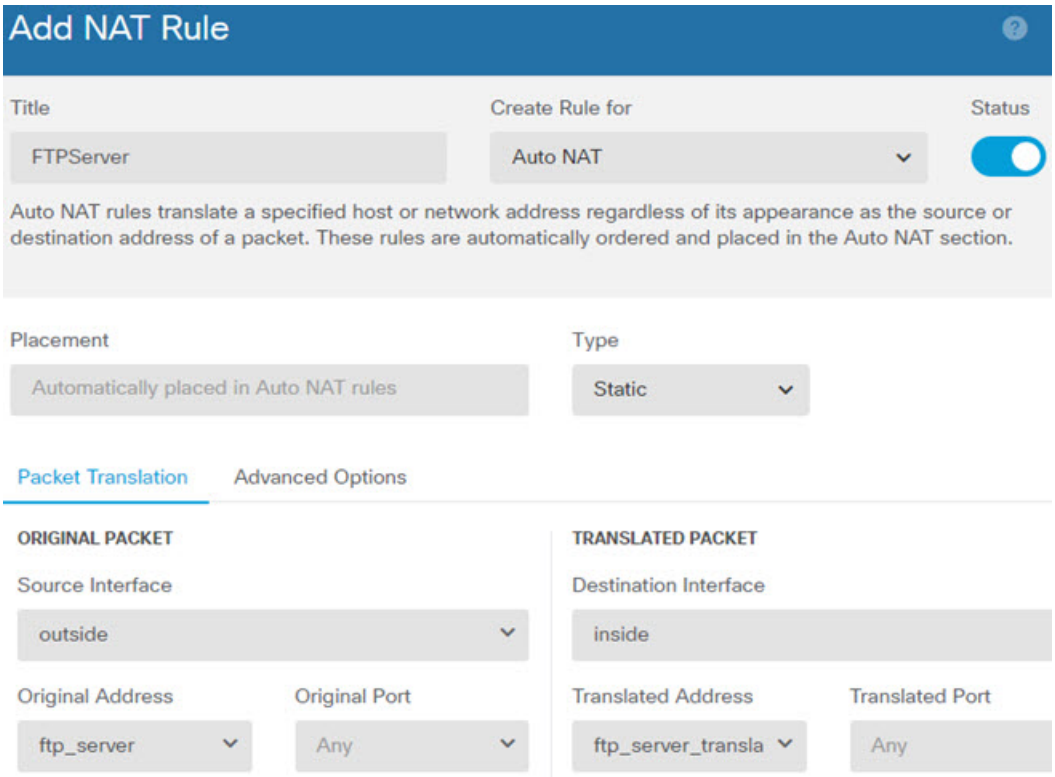
a) Select **Policies** > **NAT**.

b) Click the + button.

c) Configure the following properties:

- **Title** = FTPServer (or another name of your choosing).
- **Create Rule For** = Auto NAT.

- **Type** = Static.
- **Source Interface** = outside.
- **Destination Interface** = inside.
- **Original Address** = ftp_server network object.
- **Translated Address** = ftp_server_translated network object.
- On the **Advanced Options** tab, select **Translate DNS replies that match this rule**.



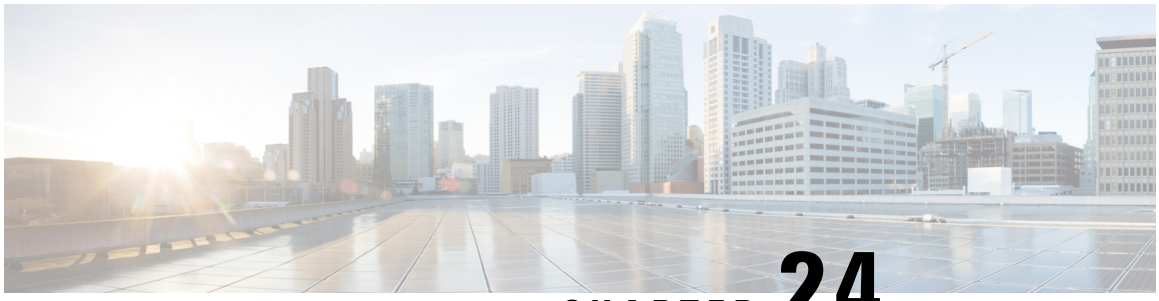
d) Click **OK**.



PART VI

Virtual Private Networks (VPN)

- Site-to-Site VPN, on page 629
- Remote Access VPN, on page 671



CHAPTER 24

Site-to-Site VPN

A virtual private network (VPN) is a network connection that establishes a secure tunnel between remote peers using a public source, such as the Internet or other network. VPNs use tunnels to encapsulate data packets within normal IP packets for forwarding over IP-based networks. They use encryption to ensure privacy and authentication to ensure the integrity of data.

- VPN Basics, on page 629
- Managing Site-to-Site VPNs, on page 636
- Monitoring Site-to-Site VPN, on page 652
- Examples for Site-to-Site VPN, on page 652

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 enabled 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, obtained from the remote peer, applies. By default, a simple IKE policy that uses DES is the only enabled policy. You can enable other IKE policies at higher priorities to negotiate stronger encryption standards, but the DES policy should ensure a successful negotiation.

VPN encryption and performance

When you configure VPN tunnel encryption, provide sufficient protection and maintain efficiency by balancing security and performance.

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 with IKE policies and IPsec proposals. Using stronger tunnel encryption may lower system performance.

If your device license allows strong encryption, you can choose from a range of encryption and hash algorithms and Diffie-Hellman groups. This document does not provide specific guidance on which options you should 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.

Decide encryption algorithms for VPN policies

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 available encryption algorithms. If you are not qualified for strong encryption, you can select DES only.

Strong encryption license considerations

Note If you are qualified for strong encryption, before upgrading from the evaluation license to a smart license, check and update your encryption algorithms for stronger encryption so that the VPN configuration works properly. Choose AES-based algorithms. DES is not supported if you are registered using an account that supports strong encryption. After registration, you cannot deploy changes until you remove all uses of DES.

Available encryption algorithms

- AES-GCM—(IKEv2 only) Advanced Encryption Standard in Galois or Counter Mode is a block cipher mode that provides confidentiality and data-origin authentication. It offers greater security than AES.

  AES-GCM offers three different key strengths: 128-bit, 192-bit, and 256-bit keys. Longer keys increase security but reduce performance. NSA Suite B, a set of cryptographic algorithms that devices must support to meet federal cryptographic strength standards, requires GCM.
- 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-bit, 192-bit, and 256-bit keys. Longer keys increase security but reduce performance.
- 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.
- Null, ESP-Null—A null encryption algorithm provides authentication without encryption. This method is not secure; use at your own discretion.

Decide which hash algorithms to use

In IKE policies, the hash algorithm creates a message digest to ensure your messages remain secure. In IKEv2, you can choose one hash algorithm for integrity and another for the pseudo-random function (PRF).

In IPsec proposals, the Encapsulating Security Protocol (ESP) uses the hash algorithm for authentication. In IKEv2 IPsec proposals, this is called the integrity hash. In IKEv1 IPsec proposals, the algorithm name has the ESP- prefix and an -HMAC suffix.

The system arranges your settings from the most secure to the least secure and negotiates with the peer in that order. For IKEv1, select only one option.

Select a hash algorithm that meets your security and performance needs:

- SHA (Secure Hash Algorithm)—The standard SHA (SHA1) produces a 160-bit digest.

  These SHA-2 options provide increased security and are available for IKEv2 configurations. Choose one if you require NSA Suite B cryptography compliance.

  - 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.

- Null or None (NULL, ESP-NONE)—(IPsec Proposals only) Use a null hash algorithm only for testing purposes. If you select one of the AES-GCM options as the encryption algorithm, choose the null integrity algorithm. For these encryption standards, the integrity hash is ignored even if you choose a non-null option.

Decide which Diffie-Hellman modulus group to use

You can use Diffie-Hellman key derivation algorithms to generate IPsec security association (SA) keys. Each group has a different size modulus. A larger modulus increases security but requires more processing time. You must have a matching modulus group on both peers.

For AES encryption, use Diffie-Hellman (DH) Group 5 or higher to support the large key sizes required by AES. IKEv1 policies do not support all of the groups.

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 a 2048-bit modulus reduce exposure 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, then negotiates with the peer using that order. For IKEv1, you can select only a single option.

- 14—Diffie-Hellman Group 14: 2048-bit modular exponential (MODP) group. Considered good protection for 192-bit keys.
- 15—Diffie-Hellman Group 15: 3072-bit MODP group.
- 16—Diffie-Hellman Group 16: 4096-bit MODP group.
- 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.

- 31—Diffie-Hellman Group 31: Curve25519 256-bit EC Group.

Deciding Which Authentication Method to Use

You can use the following methods to authenticate the peers in a site-to-site VPN connection.

Preshared Keys

Preshared keys are secret key strings configured on each peer in the connection. These keys are used by IKE during the authentication phase. For IKEv1, you must configure the same preshared key on each peer. For IKEv2, you can configure unique keys on each peer.

Preshared keys do not scale well compared to certificates. If you need to configure a large number of site-to-site VPN connections, use the certificate method instead of the preshared key method.

Certificates

Digital certificates use RSA key pairs to sign and encrypt IKE key management messages. When you configure each end of the site-to-site VPN connection, you select the local device's identity certificate, so the remote peer can authenticate the local peer.

To use the certificate method, you need to do the following:

1. Enroll your local peer with a Certificate Authority (CA) and obtain a device identity certificate. Upload this certificate to the device. For more information, see Uploading Internal and Internal CA Certificates, on page 151.

   If you also are responsible for the remote peer, also enroll that peer. Although using the same CA for the peers is convenient, it is not a requirement.

   You cannot use a self-signed certificate to establish a VPN connection. You must enroll the device with a Certificate Authority.

   If you use a Windows Certificate Authority (CA) to create certificates for site-to-site VPN endpoints, you must use a certificate that specifies IP security end system for the Application Policies extension. You can find this on the certificate's Properties dialog box on the Extensions tab (on the Windows CA server). The default for this extension is IP security IKE intermediate, which does not work for a site-to-site VPN configured using Firewall Device Manager.

2. Upload the trusted CA certificate that was used to sign the local peer's identity certificate. If you used an intermediate CA, upload the full chain, including the root and intermediate certificates. For more information, see Uploading Trusted CA Certificates, on page 154.

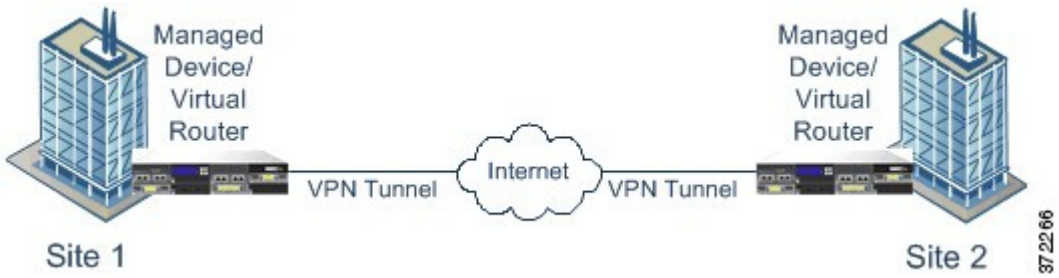
3. If the remote peer was enrolled with a different CA, also upload the trusted CA certificate used to sign the remote peer's identity certificate. Obtain the certificate from the organization that controls the remote peer. If they used an intermediate CA, upload the full chain, including the root and intermediate certificates.

4. When you configure the site-to-site VPN connection, select the certificate method, and then select the local peer's identity certificate. Each end of the connection specifies the certificate for the local end of the connection; you do not specify the certificate for the remote peer.

VPN Topologies

You can configure only point-to-point VPN connections using Firewall Device Manager. Although all connections are point-to-point, you can link into larger hub-and-spoke or meshed VPNs by defining each of the tunnels in which your device participates.

The following diagram displays a typical 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.



Establishing Site-to-Site VPN Connections with Dynamically-Addressed Peers

You can create site-to-site VPN connections to peers even when you do not know the peer's IP address. This can be useful in the following situations:

- If the peer obtains its address using DHCP, you cannot depend on the remote endpoint having a specific static IP address.
- When you want to allow an indeterminate number of remote peers to establish a connection with the device, which will serve as a hub in a hub-and-spoke topology.

When you need to establish a secure connection to a dynamically-addressed peer B, you need to ensure that your end of the connection, A, has a static IP address. Then, when you create the connection on A, specify that the peer's address is dynamic. However, when you configure the connection on the peer B, ensure that you enter the IP address for A as the remote-peer address.

When the system establishes site-to-site VPN connections, any connections where the peer has a dynamic address will be response-only. That is, the remote peer must be the one that initiates the connection. When the remote peer attempts to establish the connection, your device validates the connection using the preshared key or the certificate, whichever method you defined in the connection.

Because the VPN connection is established only after the remote peer initiates the connection, any outbound traffic that matches access control rules that allow traffic in the VPN tunnel will be dropped until that connection is established. This ensures that data does not leave your network without the appropriate encryption and VPN protection.

Virtual Tunnel Interfaces and Route-Based VPN

Traditionally, you configure a site-to-site VPN connection by defining the specific local and remote networks that will be encrypted over the VPN tunnel. These are defined in a crypto map that is part of the VPN connection profile. This type of site-to-site VPN is called policy-based.

Alternatively, you can configure a route-based site-to-site VPN. In this case, you create a Virtual Tunnel Interface (VTI), which is a virtual interface that is associated with a specific physical interface, typically the outside interface. Then, you use the routing table, with static and dynamic routes, to direct desirable traffic to the VTI. Any traffic routed through the VTI (egressing) is encrypted over the VPN tunnel that you configure for the VTI.

Thus, with route-based site-to-site VPN, you can manage the protected networks in a given VPN connection by simply changing the routing table, without altering the VPN connection profile at all. You do not need to keep track of remote networks and update the VPN connection profile to account for these changes. This simplifies VPN management for cloud service providers and large enterprises.

In addition, you can create access control rules for the VTI to fine-tune the types of traffic allowed in the tunnel. For example, you can apply intrusion inspection, and URL and application filtering.

Overview Process for Configuring Route-Based VPNs

As an overview, the process for setting up a route-based site-to-site VPN includes the following steps:

Procedure

Step 1 Create the IKEv1/2 policy and IPsec proposal for the local endpoint.

Step 2 Create a Virtual Tunnel Interface (VTI) associated with the physical interface that faces the remote peer.

Step 3 Create the site-to-site VPN connection profile that uses the VTI, IKE policy, and IPsec proposal.

Step 4 Create the same IKE and IPsec proposals on the remote peer, and a remote VTI, and the site-to-site VPN connection profile that specifies this local VTI as the remote endpoint (from the point of view of the remote peer).

Step 5 Create routes and access control rules on both peers to send the appropriate traffic through the tunnel.

Ensure that the routes and access control on each endpoint mirror each other, to allow traffic to flow in both directions.

Static routes would have these general characteristics:

- **Interface**—The virtual tunnel interface (VTI) name.
- **Networks**—Network objects that define the remote networks that are protected by the remote endpoint.
- **Gateway**—Network object that defines the IP address of the remote endpoint of the VPN tunnel.

Guidelines for Virtual Tunnel Interfaces and Route-Based VPN

IPv6 Guidelines

Virtual tunnel interfaces support IPv4 addresses only. You cannot configure an IPv6 address on a VTI.

Additional Guilelines

- You can create a maximum of 1024 VTIs.

- You cannot configure reverse route injection, either static or dynamic, on a VTI route-based VPN. (You can configure reverse route injection using the Firewall Threat Defense API only.)
- You cannot configure a dynamic peer address when you select a VTI as the local interface.
- You cannot configure remote backup peers when you select a VTI as the local interface.
- You cannot create a VTI for a source interface that is assigned to a custom virtual router. When using virtual routers, you can configure VTIs on interfaces in the global virtual router only.
- IKE and IPsec security associations will be re-keyed continuously regardless of data traffic in the tunnel. This ensures that VTI tunnels are always up.
- You cannot configure both IKEv1 and IKEv2 on a route-based connection profile: you must configure one IKE version only.
- You can configure different VTI and policy-based (crypto map) configurations on the same physical interface provided that the peer address configured in the crypto map and the tunnel destination for the VTI are different.
- Only the BGP routing protocol is supported over the VTI.
- If the system is terminating IOS IKEv2 VTI clients, disable the config-exchange request on IOS, because the system cannot retrieve the mode-CFG attributes for the session initiated by an IOS VTI client.
- Route-based site-to-site VPNs are configured as bidirectional, meaning that either endpoint of the VPN tunnel can initiate the connection. After you create the connection profile, you can change this endpoint to be either the sole initiator (INITIATE_ONLY) or exclusively the responder (RESPOND_ONLY). Ensure that you modify the remote endpoint to use the complementary connection type. To make this change, you must go to the API explorer and use GET /devices/default/s2sconnectionprofiles to find the connection profile. You can then copy/paste the body content to the PUT /devices/default/s2sconnectionprofiles/{objId} method, update **connectionType** to specify the desired type, and run the method.

Managing Site-to-Site VPNs

A virtual private network (VPN) is a network connection that establishes a secure tunnel between remote peers using a public source, such as the Internet or other network. VPNs use tunnels to encapsulate data packets within normal IP packets for forwarding over IP-based networks. They use encryption to ensure privacy and authentication to ensure the integrity of data.

You can create VPN connections to peer devices. All connections are point-to-point, but you can link the device into larger hub-and-spoke or meshed VPNs by configuring all relevant connections.

Before you begin

The following facts control the type and number of site-to-site VPN connections that you can recreate:

- VPN connections use encryption to secure network privacy. The encryption algorithms that you can use depend on whether your base license allows strong encryption. 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.

- You can create at most 20 unique IPsec profiles. Uniqueness is determined by the combination of IKEv1/v2 proposals and certificates, connection type, DH group and SA lifetime. You can reuse existing profiles. Thus, if you use the same settings for all your site-to-site VPN connections, you have one unique IPsec profile. Once you reach the limit of 20 unique IPsec profiles, you cannot create new site-to-site VPN connections unless you use the same combination of attributes that you used for an existing connection profile.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Site-to-Site VPN group.

This opens the Site-to-Site VPN page, which lists all of the connections that you have configured.

Step 2 Do any of the following.

- To create a new Site-to-Site VPN connection, click the + button. See Configuring a Site-to-Site VPN Connection, on page 637.

  If there are no connections yet, you can also click the **Create Site-to-Site Connection** button.

- To edit an existing connection, click the edit icon ( ) for the connection. See Configuring a Site-to-Site VPN Connection, on page 637.

- To copy a summary of the connection configuration to the clipboard, click the copy icon ( ) for the connection. You can paste this information in a document and send it to the administrator for the remote device to help configure that end of the connection.

- To delete a connection that you no longer need, click the delete icon ( ) for the connection.

Configuring a Site-to-Site VPN Connection

You can create a point-to-point VPN connection to link your device to another device, assuming that you have the cooperation and permission of the remote device owner. Although all connections are point-to-point, you can link into larger hub-and-spoke or meshed VPNs by defining each of the tunnels in which your device participates.

Before you begin

You can create a single VPN connection per local network/remote network combination. However, you can create multiple connections for a local network if the remote network is unique in each connection profile.

If the remote networks overlap, be careful that you create the more restrictive connection profile first. The system will create the tunnels in the order in which you create the connection profiles, not the order in which they are shown (which is simply alphabetical).

For example, if you want one tunnel from 192.16.0.0/16 to 10.91.0.0/16 to go to remote endpoint A, but tunnel 192.16.0.0/24 to the rest of 10.0.0.0/8 through remote endpoint B, you must create the connection profile for A before you create the one for B.

Note If you create more than one policy-based site-to-site VPN connection, you must specify the same interface as the Local VPN Access Interface in each profile. You cannot create policy-based site-to-site VPNs to more than one interface using FDM.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Site-to-Site VPN group.

Step 2 Do any of the following:

- To create a new Site-to-Site VPN connection, click the + button.

  If there are no connections yet, you can also click the **Create Site-to-Site Connection** button.

- To edit an existing connection, click the edit icon ( ) for the connection.

To delete a connection that you no longer need, click the delete icon ( ) for the connection.

Step 3 Define the endpoints of the point-to-point VPN connection.

- **Connection Profile Name**—The name for this connection, up to 64 characters without spaces. For example, MainOffice. You cannot use an IP address as the name.
- **Type**—How you will identify which traffic should be sent through the VPN tunnel. Select one of the following:
  - **Route Based (VTI)**—You will use the routing table, primarily static routes, to define the local and remote networks that should participate in the tunnel. If you select this option, you must select a Virtual Tunnel Interface (VTI) as the local VPN access interface. You must also use a static IP address for the remote end of the tunnel. Ensure that you configure the appropriate static routes and access control rules for the VTI after you create the VPN connection profile.
  - **Policy Based**—You will specify the local and remote networks directly in the site-to-site VPN connection profile. This is the classic approach to defining which traffic should be protected by the VPN tunnel.
- **Local Site**—These options define the local endpoint.
  - **Local VPN Access Interface**—Select the interface to which the remote peer can connect. This is typically the outside interface. The interface cannot be a member of a bridge group.
  - **Local Network**—(Policy-based only.) Click + and select the network objects that identify the local networks that should participate in the VPN connection. Users on these networks will be able to reach the remote networks through the connection.

Note
You can use IPv4 or IPv6 addresses for these networks, but you must have a matching address type on each side of the connection. For example, the VPN connection for a local IPv4 network must have at least one remote IPv4 network. You can combine IPv4 and IPv6 on both sides of a singe connection. The protected networks for the endpoints cannot overlap.

- **Remote Site**—These options define the remote endpoint.
  - **Static**/**Dynamic**—Whether the IP address of the remote peer is statically or dynamically defined (for example, through DHCP). If you select **Static**, also enter the remote peer's IP address. If you select **Dynamic**, only the remote peer will be able to initiate this VPN connection.

    For route-based VPN, you can select **Static** only.
  - **Remote IP Address** (Static addressing only.)—Enter the IP address of the remote VPN peer's interface that will host the VPN connection.
  - **Remote Backup Peers**—(Optional, policy-based connections only.) Click **Add Peer** to add a backup for the remote endpoint. If the primary endpoint becomes unavailable, the system tries to re-establish the VPN connection with one of the backup peers. You can add multiple backups.

    When you configure each backup peer, you can configure the preshared keys and certificates to use with that peer. Use the same technique you configure for the primary remote peer. Leave these settings blank to use the same values configured for the connection profile.

    After you configure the first backup peer, you can add another by clicking **Add Another Peer**, or delete a peer, or click **Edit** to change a peer's settings.

    If a backup peer is reachable through a different interface than the primary peer, ensure that you select the required interface under **Local VPN Access Interface**.
  - **Remote Network**—(Policy-based only.) Click **+** and select the network objects that identify the remote networks that should participate in the VPN connection. Users on these networks will be able to reach the local networks through the connection.

Step 4 Click **Next**.

Step 5 Define the privacy configuration for the VPN.

Note
Your license determines which encryption protocols you can select. You must qualify for strong encryption, i.e. satisfy export controls, to choose any but the most basic options.

- **IKE Version 2**, **IKE Version 1**—Choose the IKE versions to use during Internet Key Exchange (IKE) negotiations. For policy-based connections, you can select either or both; for route-based, you can select one only. When the device attempts to negotiate a connection with the other peer, it uses whichever versions you allow and that the other peer accepts. If you allow both versions, the device automatically falls back to the other version if negotiations are unsuccessful with the initially chosen version. IKEv2 is always tried first if it is configured. Both peers must support IKEv2 to use it in a negotiation.
- **IKE Policy**—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). This is a global policy: the objects you enable are applied to all VPNs. Click **Edit** to examine the current globally-enabled policies per IKE version, and to enable and create new policies. For more information, see Configuring the Global IKE Policy, on page 642.
- **IPsec Proposal**—The IPsec proposal defines the combination of security protocols and algorithms that secure traffic in an IPsec tunnel. Click **Edit** and select the proposals for each IKE version. Select all proposals that you want to allow. Click **Set Default** to simply select the system defaults, which differ based on your export compliance. The system negotiates with the peer, starting from the strongest to the weakest proposal, until a match is agreed upon. For more information, see Configuring IPsec Proposals, on page 647.

- **Authentication Type**—How you want to authenticate the peers in the VPN connection, either **Preshared Manual Key** or **Certificate**. You also need to fill in the following fields based on your selection. For IKEv1, your selection must match the authentication method selected in the IKEv1 policy object configured for the connection. For detailed information on the options, see Deciding Which Authentication Method to Use, on page 633.
  - (IKEv2) **Local Preshared Key**, **Remote Peer Preshared Key**—The keys defined on this device and on the remote device for the VPN connection. These keys can be different in IKEv2. The key can be 1-127 alphanumeric characters.
  - (IKEv1) **Preshared Key**—The key that is defined on both the local and remote device. The key can be 1-127 alphanumeric characters.
  - **Certificate**—The device identity certificate for the local peer. This must be a certificate obtained from a Certificate Authority (CA); you cannot use a self-signed certificate. If you have not uploaded the certificate, click the **Create New Object** link. You will also need to upload the root and any intermediate trusted CA certificates used to sign the identity certificate. Ensure that you set the **Validation Usage** for the uploaded certificate to include **IPsec Client**. If you have not already uploaded them, you can do so after completing this wizard.
- **IPsec Settings**—The lifetime for the security association. Once the lifetime is reached, the system re-negotiates the security association. When the system receives a negotiation request from the peer, it uses the smaller of the lifetime values proposed by the peer or the locally configured lifetime values as the lifetime of the new security associations. There are two lifetimes: a “timed” lifetime and a “traffic-volume” lifetime. The security association expires after the first of these lifetimes is reached.
  - **Lifetime Duration**—The number of seconds a security association can live before it expires. The range is 120 to 214783647 seconds. The global default is 28,800 seconds (eight hours).
  - **Lifetime Size**—The volume of traffic (in kilobytes) that can pass between peers using a given security association before that security association expires. The range is 10 to 2147483647 kilobytes, or blank. The global default is 4,608,000 kilobytes. Leave the field blank to remove the size-based limit and use duration as the only limit.
- **NAT Exempt**—(Policy-based only.) Whether to exempt the VPN traffic from NAT policies on the local VPN access interface. If you do not want NAT rules to apply to the local network, select the interface that hosts the local network. This option works only if the local network resides behind a single routed interface (not a bridge group member). If the local network is behind more than one routed interface, or one or more bridge group members, you must manually create the NAT exempt rules. For information on manually creating the required rules, see Exempting Site-to-Site VPN Traffic from NAT, on page 652.
- **Diffie-Helman Group for 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. To enable Perfect Forward Secrecy, select the Diffie-Hellman key derivation algorithm to use when generating the PFS session key in the Modulus Group list. If you enable both IKEv1 and IKEv2, the options are limited to those supported by IKEv1. For an explanation of the options, see Decide which Diffie-Hellman modulus group to use, on page 632.

Step 6 Click **Next**.

Step 7 Review the summary and click **Finish**.

The summary information is copied to the clipboard. You can paste the information in a document and use it to help you configure the remote peer, or to send it to the party responsible for configuring the peer.

You must take additional steps to allow traffic within the VPN tunnel, as explained in Allowing Traffic Through the Site-to-Site VPN, on page 642.

After you deploy the configuration, log into the device CLI and use the **show ipsec sa** command to verify that the endpoints establish a security association. See Verifying Site-to-Site VPN Connections, on page 649.

Configuring a Virtual Tunnel Interface

You can use a Virtual Tunnel Interface (VTI) in a route-based site-to-site VPN connection profile only. A VTI is associated with a physical interface, through which the VPN connection is made to the remote peer. Using a virtual interface, you can simplify the site-to-site VPN connection and control traffic using static and dynamic routes rather than specifying the local and remote networks for the VPN in the connection profile.

Procedure

Step 1 Click **Device**, click the link in the Interfaces summary, then click **Virtual Tunnel Interfaces**.

Step 2 Do any of the following:

- Click + or **Create Virtual Tunnel Interface** to create a new interface.
- Click the edit icon ( ) for an existing interface.

If you no longer need an interface, click the delete icon ( ) for it. You must first delete any site-to-site connection profile that uses the interface before you can delete it.

Step 3 Configure the following options:

- **Name**—The interface name, up to 48 characters. If you change the name of an existing interface, it is automatically changed in all policies and objects that include it. You cannot use uppercase letters in the name.
- **Status**—Click the slider to the Enabled position .
- **Description**—(Optional.) The description can be up to 200 characters on a single line, without carriage returns.
- **Tunnel ID**—A number from 0-10413. This number is appended to the word Tunnel to form the hardware name of the interface. You must chose a number that you have not already used for another VTI. For example, enter 1 to create interface Tunnel1.
- **Tunnel Source**—Select the interface that is associated with this VTI. Tunnel Source is the interface through which the site-to-site VPN defined on the virtual tunnel interface will connect to the remote endpoint. Select an interface that can reach the remote endpoint, such as the outside interface. The source interface can be a physical interface, subinterface, or Etherchannel, and it must have a name. The interface cannot be a member of a Bridge Virtual Interface (BVI).

- **IP Address and Subnet Mask**—The IPv4 address and associated subnet mask. For example, 192.168.1.1/24 or /255.255.255.0. This address does not need to be on the same subnet as the address of the Tunnel Source interface. However, if you configure remote access (RA) VPN on the source interface, the VTI IP address cannot be within the address pool configured for the RA VPN.

Step 4 Click **OK**.

Allowing Traffic Through the Site-to-Site VPN

You can use one of the following techniques to enable traffic flow in the site-to-site VPN tunnel.

- Configure the **sysopt connection permit-vpn** command, which exempts traffic that matches the VPN connection from the access control policy. The default for this command is **no sysopt connection permit-vpn**, which means VPN traffic must also be allowed by the access control policy.

  This is the more secure method to allow traffic in the VPN, because external users cannot spoof IP addresses in the remote protected network. The downside is that the VPN traffic will not be inspected, which means that intrusion and file protection, URL filtering, or other advanced features will not be applied to the traffic. This also means that no connection events will be generated for the traffic, and thus statistical dashboards will not reflect VPN connections.

  The preferred method to configure this command is to create a remote access VPN connection profile in which you select the **Bypass Access Control policy for decrypted traffic** option. If you do not want to configure RA VPN, or you cannot configure RA VPN, you can use FlexConfig to configure the command.



Note This method does not apply to route-based VPN connections configured on a Virtual Tunnel Interface (VTI). You must always configure access control rules for route-based VPNs.

- Create access control rules to allow connections from the remote network. This method ensures that VPN traffic is inspected and advanced services can be applied to the connections. The downside is that it opens the possibility for external users to spoof IP addresses and thus gain access to your internal network.

Configuring the Global IKE Policy

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.

IKE policy objects define the IKE proposals for these negotiations. The objects that you enable are the ones used when the peers negotiate a VPN connection: you cannot specify different IKE policies per connection.

The relative priority of each object determines which of these policies are tried first, with the lower number being higher priority. The connection is not established if the negotiation fails to find a policy that both peers can support.

To define the global IKE policy, you select which objects to enable for each IKE version. If the pre-defined objects do not satisfy your requirements, create new policies to enforce your security policy.

The following procedure explains how to configure the global policy through the Objects page. You can also enable, disable, and create policies when editing a VPN connection by clicking **Edit** for the IKE Policy settings.

Note You can enable up to 20 IKE policies.

Procedure

Step 1 Select **Objects**, then select **IKE Policies** from the table of contents.

Policies for IKEv1 and IKEv2 are shown in separate lists.

Step 2 Enable the IKE policies you want to allow for each IKE version.

a) Select **IKEv1** or **IKEv2** above the object table to show the policies for that version.

b) Click the **State** toggle to enable the appropriate objects and to disable objects that do not meet your requirements.

If some of your security requirements are not reflected in the existing objects, define new ones to implement your requirements. For details, see the following topics:

- Configuring IKEv1 Policies, on page 643
- Configuring IKEv2 Policies, on page 645

c) Verify that the relative priorities match your requirements.

If you need to change the priority of a policy, edit it. If the policy is a pre-defined system policy, you need to create your own version of the policy to change the priority.

The priority is relative, and not absolute. For example, priority 80 is higher than 160. If 80 is the highest priority object that you enable, that becomes your first-choice policy. If you then enable a policy with priority 25, that becomes your first-choice policy.

d) If you use both IKE versions, repeat the process for the other version.

Configuring IKEv1 Policies

Internet Key Exchange (IKE) version 1 policy objects contain the parameters required for IKEv1 policies when defining VPN connections. IKE is a key management protocol that facilitates the management of IPsec-based communications. It is used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and automatically establish IPsec security associations (SAs).

There are several pre-defined IKEv1 policies. If any suit your needs, simply enable them by clicking the **State** toggle. You can also create new policies to implement other combinations of security settings. You cannot edit or delete system-defined objects.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create IKEv1 Policy objects while editing the IKEv1 settings in a VPN connection by clicking the **Create New IKE Policy** link shown in the object list.

Procedure

Step 1 Select **Objects**, then select **IKE Policies** from the table of contents.

Step 2 Select **IKEv1** above the object table to show IKEv1 policies.

Step 3 If any of the system-defined policies meet your requirements, click the **State** toggle to enable them.

Also use the **State** toggle to disable unwanted policies. The relative priority determines which of these policies are tried first, with the lower number being higher priority.

Step 4 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 5 Configure the IKEv1 properties.

- **Priority**—The relative priority of the IKE policy, from 1 to 65,535. The priority 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 highest priority policy, it tries to use the parameters defined in the next lowest priority. The lower the number, the higher the priority.
- **Name**—The name of the object, up to 128 characters.
- **State**—Whether the IKE policy is enabled or disabled. Click the toggle to change the state. Only enabled policies are used during IKE negotiations.
- **Authentication**—The method of authentication to use between the two peers. For more information, see Deciding Which Authentication Method to Use, on page 633.
  - **Preshared Key**—Use the preshared key that is defined on each device. These keys allow for a secret key to be shared between two peers and to be used by IKE during the authentication phase. If the peer is not configured with the same preshared key, the IKE SA cannot be established.
  - **Certificate**—Use the device identity certificates for the peers to identify each other. You must obtain these certificates by enrolling each peer in a Certificate Authority. You must also upload the trusted CA root and intermediate CA certificates used to sign the identity certificates in each peer. The peers can be enrolled in the same or a different CA. You cannot use self-signed certificates for either peer.

- **Encryption**—The encryption algorithm used to establish the Phase 1 security association (SA) for protecting Phase 2 negotiations. For an explanation of the options, see Decide encryption algorithms for VPN policies, on page 631.
- **Diffie-Hellman 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 an explanation of the options, see Decide which Diffie-Hellman modulus group to use, on page 632.
- **Hash**—The hash algorithm for creating a message digest, which is used to ensure message integrity. For an explanation of the options, see Decide which hash algorithms to use, on page 632.
- **Lifetime**—The lifetime of the security association (SA), in seconds, from 120 to 2147483647 or blank. When the lifetime is exceeded, the SA expires and must be renegotiated between the two peers. As a general rule, the shorter the lifetime (up to a point), the more secure your IKE negotiations will be. However, with longer lifetimes, future IPsec security associations can be set up more quickly than with shorter lifetimes. The default is 86400. To specify an unlimited lifetime, enter no value (leave the field blank).

Step 6 Click **OK** to save your changes.

Configuring IKEv2 Policies

Internet Key Exchange (IKE) version 2 policy objects contain the parameters required for IKEv2 policies when defining VPN connections. IKE is a key management protocol that facilitates the management of IPsec-based communications. It is used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and automatically establish IPsec security associations (SAs).

There are several pre-defined IKEv2 policies. If any suit your needs, simply enable them by clicking the **State** toggle. You can also create new policies to implement other combinations of security settings. You cannot edit or delete system-defined objects.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create IKEv2 Policy objects while editing the IKEv2 settings in a VPN connection by clicking the **Create New IKE Policy** link shown in the object list.

Procedure

Step 1 Select **Objects**, then select **IKE Policies** from the table of contents.

Step 2 Select **IKEv2** above the object table to show IKEv2 policies.

Step 3 If any of the system-defined policies meet your requirements, click the **State** toggle to enable them.

Also use the **State** toggle to disable unwanted policies. The relative priority determines which of these policies are tried first, with the lower number being higher priority.

Step 4 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon () for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 5 Configure the IKEv2 properties.

- **Priority**—The relative priority of the IKE policy, from 1 to 65,535. The priority 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 highest priority policy, it tries to use the parameters defined in the next lowest priority. The lower the number, the higher the priority.
- **Name**—The name of the object, up to 128 characters.
- **State**—Whether the IKE policy is enabled or disabled. Click the toggle to change the state. Only enabled policies are used during IKE negotiations.
- **Encryption**—The encryption algorithm used to establish the Phase 1 security association (SA) for protecting Phase 2 negotiations. Select all algorithms that you want to allow, although you cannot include both mixed-mode (AES-GCM) and normal mode options in the same policy. (Normal mode requires that you select an integrity hash, whereas mixed mode prohibits a separate integrity hash selection.) The system negotiates with the peer, starting from the strongest to the weakest algorithm, until a match is agreed upon. For an explanation of the options, see Decide encryption algorithms for VPN policies, on page 631.
- **Diffie-Hellman 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. Select all algorithms that you want to allow. The system negotiates with the peer, starting from the strongest to the weakest group, until a match is agreed upon. For an explanation of the options, see Decide which Diffie-Hellman modulus group to use, on page 632.
- **Integrity Hash**—The integrity portion of the hash algorithm for creating a message digest, which is used to ensure message integrity. Select all algorithms that you want to allow. The system negotiates with the peer, starting from the strongest to the weakest algorithm, until a match is agreed upon. The integrity hash is not used with the AES-GCM encryption options. For an explanation of the options, see Decide which hash algorithms to use, on page 632.
- **Pseudo Random Function (PRF) Hash**—The pseudo-random function (PRF) portion of the hash algorithm, which is used as the algorithm to derive keying material and hashing operations required for the IKEv2 tunnel encryption. In IKEv1, the Integrity and PRF algorithms are not separated, but in IKEv2, you can specify different algorithms for these elements. Select all algorithms that you want to allow. The system negotiates with the peer, starting from the strongest to the weakest algorithm, until a match is agreed upon. For an explanation of the options, see Decide which hash algorithms to use, on page 632.
- **Lifetime**—The lifetime of the security association (SA), in seconds, from 120 to 2147483647 or blank. When the lifetime is exceeded, the SA expires and must be renegotiated between the two peers. As a general rule, the shorter the lifetime (up to a point), the more secure your IKE negotiations will be. However, with longer lifetimes, future IPsec security associations can be set up more quickly than with shorter lifetimes. The default is 86400. To specify an unlimited lifetime, enter no value (leave the field blank).

Step 6 Click **OK** to save your changes.

Configuring IPsec Proposals

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 called a transform set. During the IPsec security association (SA) negotiation, peers search for a transform set that is the same at both peers.

There are separate IPsec proposal objects based on the IKE version, IKEv1, or IKEv2:

- When you create an IKEv1 IPsec proposal, 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 and select multiple IKEv1 IPsec Proposal objects.
- When you create an IKEv2 IPsec proposal, you can select all of the encryption and hash algorithms allowed in a VPN. The system orders the settings from the most secure to the least secure and negotiates with the peer until a match is found. This allows you to potentially send a single proposal to convey all the allowed combinations instead of the need to send each allowed combination individually as with IKEv1.

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.

The following topics explain how to configure IPsec proposals for each IKE version.

Configuring IPsec Proposals for IKEv1

Use IKEv1 IPsec Proposal objects configure the IPsec proposal used during IKE Phase 2 negotiations. The IPsec proposal defines the combination of security protocols and algorithms that secure traffic in an IPsec tunnel.

There are several pre-defined IKEv1 IPsec proposals. You can also create new proposals to implement other combinations of security settings. You cannot edit or delete system-defined objects.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create IKEv1 IPsec Proposals objects while editing the IKEv1 IPsec settings in a VPN connection by clicking the **Create New IPsec Proposal** link shown in the object list.

Procedure

Step 1 Select **Objects**, then select **IPsec Proposals** from the table of contents.

Step 2 Select **IKEv1** above the object table to show IKEv1 IPsec proposals.

Step 3 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 4 Configure the IKEv1 IPsec proposal properties.

- **Name**—The name of the object, up to 128 characters.
- **Mode**—The mode in which the IPSec tunnel operates.
  - **Tunnel** mode encapsulates the entire IP packet. The IPSec header is added between the original IP header and a new IP header. This is the default. Use tunnel mode when the firewall is protecting traffic to and from hosts positioned behind the firewall. Tunnel mode is the normal way regular IPSec is implemented between two firewalls (or other security gateways) that are connected over an untrusted network, such as the Internet.
  - **Transport** mode encapsulates only the upper-layer protocols of an IP packet. The IPSec header is inserted between the IP header and the upper-layer protocol header (such as TCP). Transport mode requires that both the source and destination hosts support IPSec, and can only be used when the destination peer of the tunnel is the final destination of the IP packet. Transport mode is generally used only when protecting a Layer 2 or Layer 3 tunneling protocol such as GRE, L2TP, and DLSW.
- **ESP Encryption**—The Encapsulating Security Protocol (ESP) encryption algorithm for this proposal. For an explanation of the options, see Decide encryption algorithms for VPN policies, on page 631.
- **ESP Hash**—The hash or integrity algorithm to use for authentication. For an explanation of the options, see Decide which hash algorithms to use, on page 632.

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

Configuring IPsec Proposals for IKEv2

Use IKEv2 IPsec Proposal objects configure the IPsec proposal used during IKE Phase 2 negotiations. The IPsec proposal defines the combination of security protocols and algorithms that secure traffic in an IPsec tunnel.

There are several pre-defined IKEv2 IPsec proposals. You can also create new proposals to implement other combinations of security settings. You cannot edit or delete system-defined objects.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create IKEv2 IPsec Proposals objects while editing the IKEv2 IPsec settings in a VPN connection by clicking the **Create New IPsec Proposal** link shown in the object list.

Procedure

Step 1 Select **Objects**, then select **IPsec Proposals** from the table of contents.

Step 2 Select **IKEv2** above the object table to show IKEv2 IPsec proposals.

Step 3 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 4 Configure the IKEv2 IPsec proposal properties.

- **Name**—The name of the object, up to 128 characters.
- **Encryption**—The Encapsulating Security Protocol (ESP) encryption algorithm for this proposal. Select all algorithms that you want to allow. The system negotiates with the peer, starting from the strongest to the weakest algorithm, until a match is agreed upon. For an explanation of the options, see Decide encryption algorithms for VPN policies, on page 631.
- **Integrity Hash**—The hash or integrity algorithm to use for authentication. Select all algorithms that you want to allow. The system negotiates with the peer, starting from the strongest to the weakest algorithm, until a match is agreed upon. For an explanation of the options, see Decide which hash algorithms to use, on page 632.

  Note
  You should choose the null integrity algorithm if you select one of the AES-GCM/GMAC options as the encryption algorithm. These encryption standards do not use the integrity hash even if you select a non-null option.

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

Verifying Site-to-Site VPN Connections

After you configure a site-to-site VPN connection, and deploy the configuration to the device, verify that the system establishes the security association with the remote device.

If the connection cannot be established, use the **ping interface** *interface_name remote_ip_address* command from the device CLI to ensure there is a path through the VPN interface to the remote device. If there is no connection through the configured interface, you can leave off the **interface** *interface_name* keyword and determine if connectivity is through a different interface. You might have selected the wrong interface for the connection: you must select the interface that faces the remote device, not the interface that faces the protected network.

If there is a network path, check the IKE versions and keys configured and supported by both endpoints, and adjust the VPN connection as needed. Ensure that no access control or NAT rules are blocking the connection.

Procedure

Step 1 Log into the device CLI as explained in Logging Into the Command Line Interface (CLI), on page 15.

Step 2 Use the **show ipsec sa** command to verify that the IPsec security association is established.

You should see that the VPN connection is established between your device (the **local addr**) and the remote peer (**current_peer**). The packets (pkts) counts should increase as you send traffic through the connection. The access list should show the local and remote networks for the connection.

For example, the following output shows an IKEv2 connection.

```
> show ipsec sa
```

```
interface: site-a-outside
    Crypto map tag: s2sCryptoMap, seq num: 1, local addr: 192.168.2.15

      access-list |s2sAcl|0730e31c-1e5f-11e7-899f-27f6e1030344
extended permit ip 192.168.1.0 255.255.255.0 192.168.3.0 255.255.255.0
      local ident (addr/mask/prot/port): (192.168.1.0/255.255.255.0/0/0)
      remote ident (addr/mask/prot/port): (192.168.3.0/255.255.255.0/0/0)
      current_peer: 192.168.4.6


      #pkts encaps: 69, #pkts encrypt: 69, #pkts digest: 69
      #pkts decaps: 69, #pkts decrypt: 69, #pkts verify: 69
      #pkts compressed: 0, #pkts decompressed: 0
      #pkts not compressed: 69, #pkts comp failed: 0, #pkts decomp failed: 0
      #pre-frag successes: 0, #pre-frag failures: 0, #fragments created: 0
      #PMTUs sent: 0, #PMTUs rcvd: 0, #decapsulated frgs needing reassembly: 0
      #TFC rcvd: 0, #TFC sent: 0
      #Valid ICMP Errors rcvd: 0, #Invalid ICMP Errors rcvd: 0
      #send errors: 0, #recv errors: 0

      local crypto endpt.: 192.168.2.15/500, remote crypto endpt.: 192.168.4.6/500
      path mtu 1500, ipsec overhead 55(36), media mtu 1500
      PMTU time remaining (sec): 0, DF policy: copy-df
      ICMP error validation: disabled, TFC packets: disabled
      current outbound spi: CD22739C
      current inbound spi : 52D2F1E4

    inbound esp sas:
      spi: 0x52D2F1E4 (1389556196)
         SA State: active
         transform: esp-aes-gcm-256 esp-null-hmac no compression
         in use settings ={L2L, Tunnel, PFS Group 19, IKEv2, }
         slot: 0, conn_id: 62738432, crypto-map: s2sCryptoMap
         sa timing: remaining key lifetime (kB/sec): (4285434/28730)
         IV size: 8 bytes
         replay detection support: Y
         Anti replay bitmap:
          0xFFFFFFFF 0xFFFFFFFF
    outbound esp sas:
      spi: 0xCD22739C (3441587100)
         SA State: active
         transform: esp-aes-gcm-256 esp-null-hmac no compression
         in use settings ={L2L, Tunnel, PFS Group 19, IKEv2, }
         slot: 0, conn_id: 62738432, crypto-map: s2sCryptoMap
         sa timing: remaining key lifetime (kB/sec): (4055034/28730)
         IV size: 8 bytes
         replay detection support: Y
         Anti replay bitmap:
          0x00000000 0x00000001
```

The following output shows an IKEv1 connection.

```
> show ipsec sa
interface: site-a-outside
    Crypto map tag: s2sCryptoMap, seq num: 1, local addr: 192.168.2.15

      access-list |s2sAcl|0730e31c-1e5f-11e7-899f-27f6e1030344
extended permit ip 192.168.1.0 255.255.255.0 192.168.3.0 255.255.255.0
      local ident (addr/mask/prot/port): (192.168.1.0/255.255.255.0/0/0)
      remote ident (addr/mask/prot/port): (192.168.3.0/255.255.255.0/0/0)
      current_peer: 192.168.4.6
```

```
    #pkts encaps: 10, #pkts encrypt: 10, #pkts digest: 10
    #pkts decaps: 10, #pkts decrypt: 10, #pkts verify: 10
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 10, #pkts comp failed: 0, #pkts decomp failed: 0
    #pre-frag successes: 0, #pre-frag failures: 0, #fragments created: 0
    #PMTUs sent: 0, #PMTUs rcvd: 0, #decapsulated frgs needing reassembly: 0
    #TFC rcvd: 0, #TFC sent: 0
    #Valid ICMP Errors rcvd: 0, #Invalid ICMP Errors rcvd: 0
    #send errors: 0, #recv errors: 0

    local crypto endpt.: 192.168.2.15/0, remote crypto endpt.: 192.168.4.6/0
    path mtu 1500, ipsec overhead 74(44), media mtu 1500
    PMTU time remaining (sec): 0, DF policy: copy-df
    ICMP error validation: disabled, TFC packets: disabled
    current outbound spi: 077D72C9
    current inbound spi : AC146DEC

  inbound esp sas:
    spi: 0xAC146DEC (2887020012)
       SA State: active
       transform: esp-aes-256 esp-sha-hmac no compression
       in use settings ={L2L, Tunnel, PFS Group 5, IKEv1, }
       slot: 0, conn_id: 143065088, crypto-map: s2sCryptoMap
       sa timing: remaining key lifetime (kB/sec): (3914999/28567)
       IV size: 16 bytes
       replay detection support: Y
       Anti replay bitmap:
        0x00000000 0x000007FF
  outbound esp sas:
    spi: 0x077D72C9 (125661897)
       SA State: active
       transform: esp-aes-256 esp-sha-hmac no compression
       in use settings ={L2L, Tunnel, PFS Group 5, IKEv1, }
       slot: 0, conn_id: 143065088, crypto-map: s2sCryptoMap
       sa timing: remaining key lifetime (kB/sec): (3914999/28567)
       IV size: 16 bytes
       replay detection support: Y
       Anti replay bitmap:
        0x00000000 0x00000001
```

Step 3 Use the **show isakmp sa** command to verify the IKE security associations.

You can use the command without the **sa** keyword (or use the **stats** keyword instead) to view IKE statistics.

For example, the following output shows an IKEv2 security association.

```
> show isakmp sa

There are no IKEv1 SAs

IKEv2 SAs:

Session-id:15317, Status:UP-ACTIVE, IKE count:1, CHILD count:1

Tunnel-id Local             Remote            Status   Role
592216161 192.168.2.15/500  192.168.4.6/500   READY    INITIATOR
      Encr: AES-GCM, keysize: 256, Hash: N/A, DH Grp:21, Auth sign: PSK, Auth verify: PSK
      Life/Active Time: 86400/12 sec
Child sa: local selector  192.168.1.0/0 - 192.168.1.255/65535
          remote selector 192.168.3.0/0 - 192.168.3.255/65535
          ESP spi in/out: 0x52d2f1e4/0xcd22739c
```

The following output shows an IKEv1 security association.

```
> show isakmp sa

IKEv1 SAs:

   Active SA: 1
    Rekey SA: 0 (A tunnel will report 1 Active and 1 Rekey SA during rekey)
Total IKE SA: 1

1   IKE Peer: 192.168.4.6
    Type    : L2L             Role    : initiator
    Rekey   : no              State   : MM_ACTIVE

There are no IKEv2 SAs
```

Monitoring Site-to-Site VPN

To monitor and troubleshoot site-to-site VPN connections, open the CLI console or log into the device CLI and use the following commands.

- **show ipsec sa** displays the VPN sessions (security associations). You can reset these statistics using the **clear ipsec sa counters** command.
- **show ipsec** *keyword* displays IPsec operational data and statistics. Enter **show ipsec ?** to see the available keywords.
- **show isakmp** displays ISAKMP operational data and statistics.

Examples for Site-to-Site VPN

The following are examples of configuring site-to-site VPN.

Exempting Site-to-Site VPN Traffic from NAT

When you have a site-to-site VPN connection defined on an interface, and you also have NAT rules for that interface, you can optionally exempt the traffic on the VPN from the NAT rules. You might want to do this if the remote end of the VPN connection can handle your internal addresses.

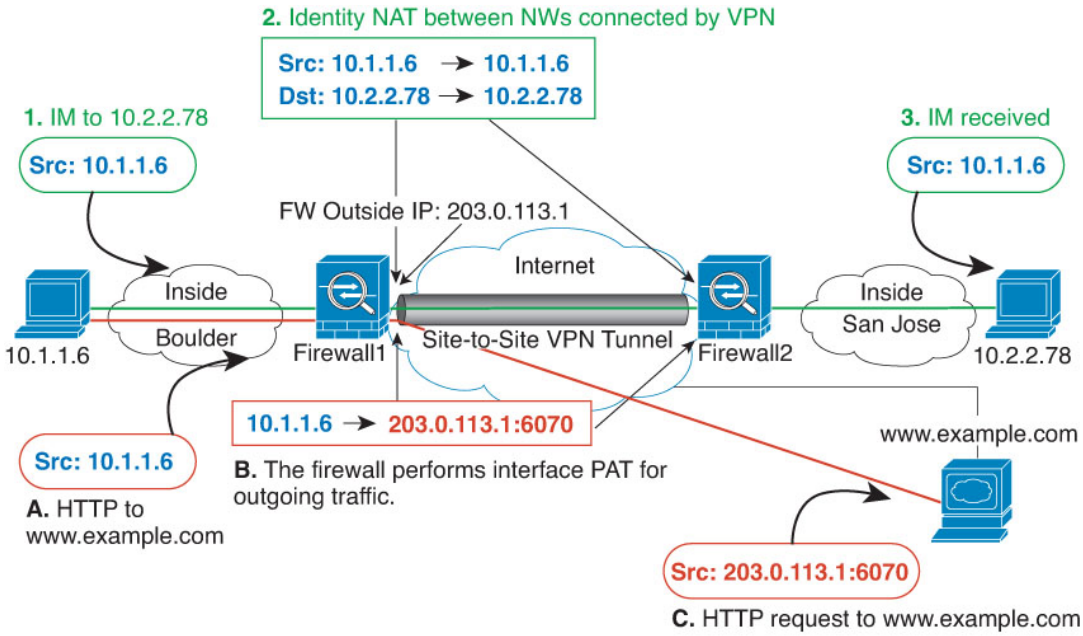
When you create the VPN connection, you can select the **NAT Exempt** option to create the rules automatically. However, this works only if your local protected network is connected through a single routed interface (not a bridge group member). If instead, the local networks in the connection reside behind two or more routed interfaces, or one or more bridge group members, you need to manually configure the NAT exempt rules.

To exempt VPN traffic from NAT rules, you create an identity manual NAT rule for the local traffic when the destination is the remote network. Then, apply NAT to the traffic when the destination is anything else (for example, the Internet). If you have more than one interface for the local network, create rules for each interface. Also, consider the following suggestions:

- If there is more than one local network in the connection, create a network object group to hold the objects that define the networks.
- If you are including both IPv4 and IPv6 networks in the VPN, create separate identity NAT rules for each.

Consider the following example, which 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 44: Interface PAT and Identity NAT for Site-to-Site VPN



The following example explains the configuration for Firewall1 (Boulder). The example assumes that the inside interface is a bridge group, so you need to write the rules for each member interface. The process is the same if you have a single or multiple routed inside interfaces.

Note This example assumes IPv4 only. If the VPN also includes IPv6 networks, create parallel rules for IPv6. Note that you cannot implement IPv6 interface PAT, so you need to create a host object with a unique IPv6 address to use for PAT.

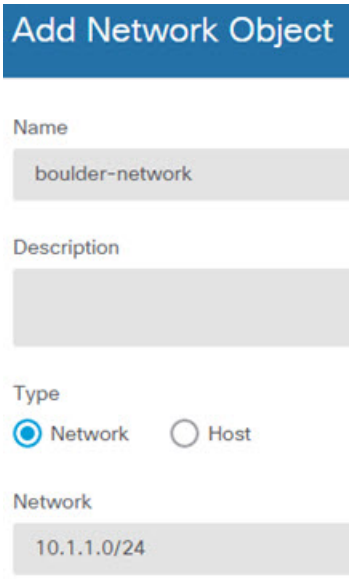
Procedure

Step 1 Create the objects to define the various networks.

a) Choose **Objects**.

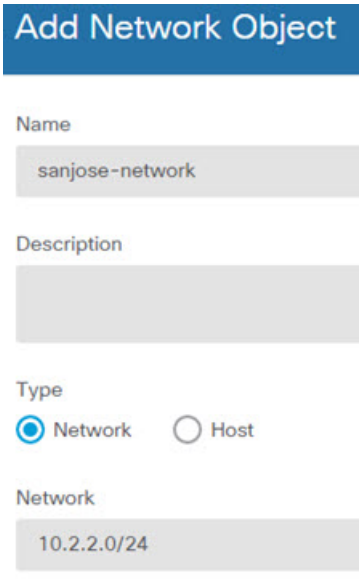
b) Select **Network** from the table of contents and click +.
c) Identify the Boulder inside network.

Name the network object (for example, boulder-network), select **Network**, and enter the network address, 10.1.1.0/24.



d) Click **OK**.
e) Click + and define the inside San Jose network.

Name the network object (for example, sanjose-network), select **Network**, and enter the network address 10.2.2.0/24.



f) Click **OK**.

Step 2 Configure manual identity NAT for the Boulder network when going over the VPN to San Jose on Firewall1 (Boulder).

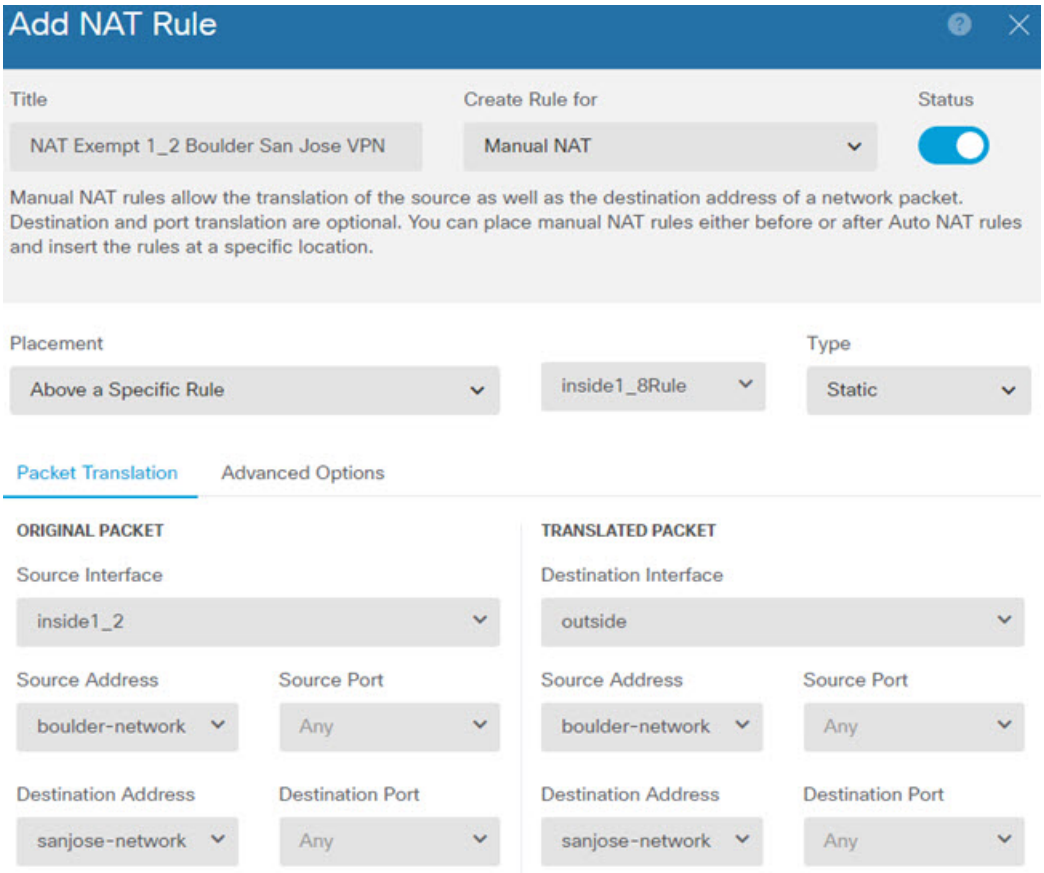
a) Select **Policies** > **NAT**.

b) Click the + button.

c) Configure the following properties:

- **Title** = NAT Exempt 1_2 Boulder San Jose VPN (or another name of your choosing).
- **Create Rule For** = Manual NAT.
- **Placement** = **Above a Specific Rule**, and select the first rule in the Manual NAT Before Auto NAT section. You want to ensure that this rule comes before any general interface PAT rules for the destination interface. Otherwise, the rule might not be applied to the right traffic.
- **Type** = Static.
- **Source Interface** = inside1_2.
- **Destination Interface** = outside.
- **Original Source Address** = boulder-network network object.
- **Translated Source Address** = boulder-network network object.
- **Original Destination Address** = sanjose-network network object.
- **Translated Destination Address** = sanjose-network 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.



d) On the **Advanced** tab, select **Do not proxy ARP on Destination interface**.
e) Click **OK**.
f) Repeat the process to create equivalent rules for each of the other inside interfaces.

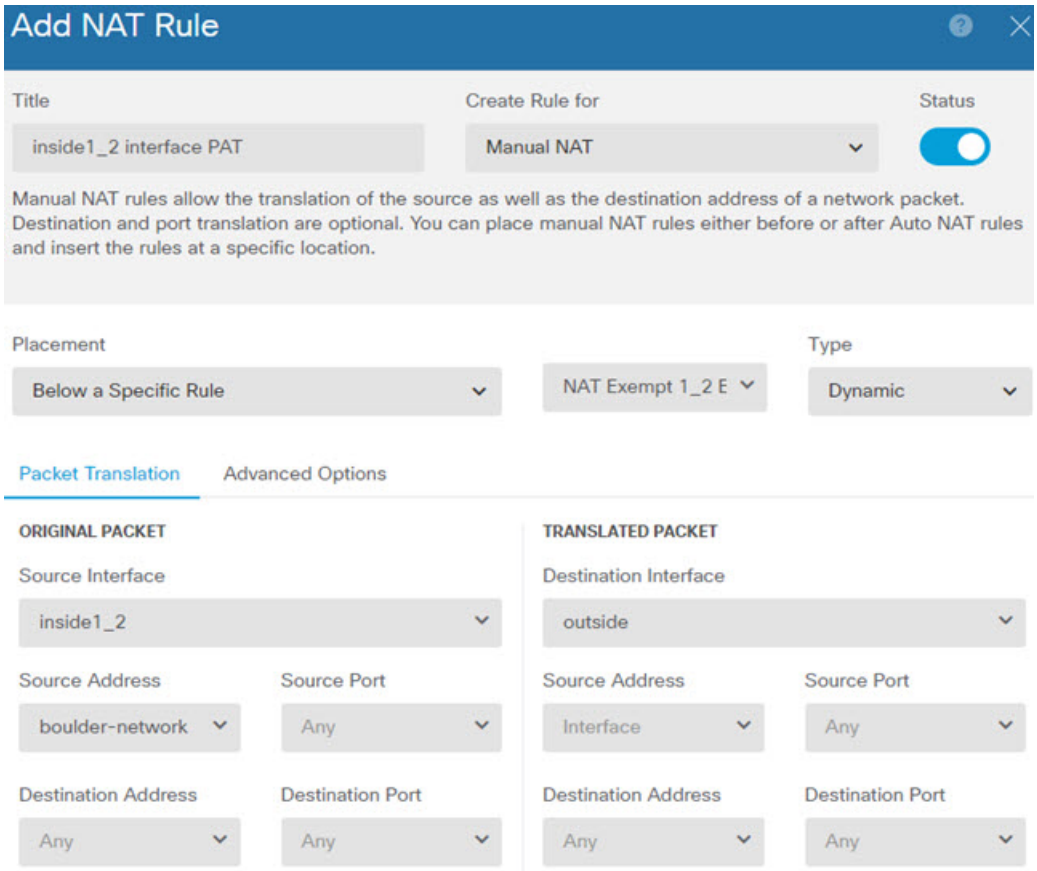
Step 3 Configure manual dynamic interface PAT when going to the Internet for the inside Boulder network on Firewall1 (Boulder).

Note
There might already be dynamic interface PAT rules for the inside interfaces, covering any IPv4 traffic, as these are created by default during initial configuration. However, the configuration is shown here for completeness. Before completing these steps, check whether a rule already exists that covers the inside interface and network, and skip this step if it does.

a) Click the + button.
b) Configure the following properties:

- **Title** = inside1_2 interface PAT (or another name of your choosing).
- **Create Rule For** = Manual NAT.
- **Placement** = **Below a Specific Rule**, and select the rule you created above for this interface in the Manual NAT Before Auto NAT section. 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, which is also sufficient.

- **Type** = Dynamic.
- **Source Interface** = inside1_2.
- **Destination Interface** = outside.
- **Original Source Address** = boulder-network network object.
- **Translated Source Address** = **Interface**. This option configures interface PAT using the destination interface.
- **Original Destination Address** = any.
- **Translated Destination Address** = any.



c) Click **OK**.

d) Repeat the process to create equivalent rules for each of the other inside interfaces.

Step 4 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.



b) Click the **Deploy Now** button.

You can wait until deployment completes, or click **OK** and check the task list or deployment history later.

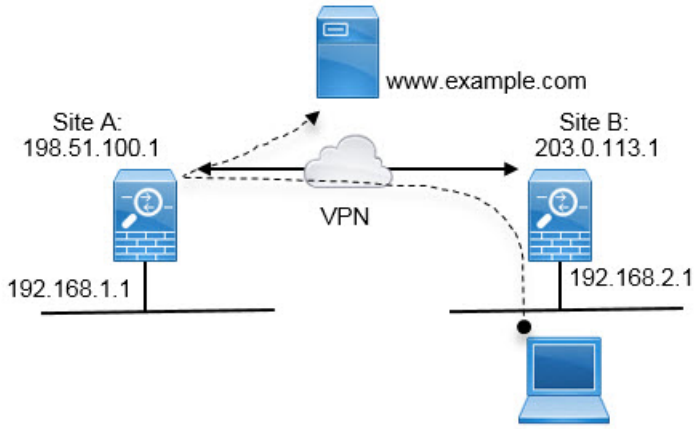
Step 5 If you are also managing Firewall2 (San Jose), you can configure similar rules for that device.

- 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."

How to Provide Internet Access on the Outside Interface for External Site-to-Site VPN Users (Hair Pinning)

In a site-to-site VPN, you might want users on the remote networks to access the Internet through your device. However, because the remote users are entering your device on the same interface that faces the Internet (the outside interface), you need to bounce Internet traffic right back out of the outside interface. This technique is sometimes called hair pinning.

The following graphic shows an example. There is a site-to-site VPN tunnel configured between 198.51.100.1 (on the main site, Site A) and 203.0.113.1 (the remote site, Site B). All user traffic from the remote site inside network, 192.168.2.0/24, goes through the VPN. Thus, when a user on that network wants to go to a server on the Internet, such as www.example.com, the connection first goes through the VPN, then gets routed back out to the Internet from the 198.51.100.1 interface.



The following procedure explains how to configure this service. You must configure both endpoints of the VPN tunnel.

Before you begin

This procedure assumes you are using the default setting for permitting VPN traffic, which subjects the VPN traffic to the access control policy. In the running configuration, this is represented by the **no sysopt connection permit-vpn** command. If you instead enabled **sysopt connection permit-vpn** through FlexConfig, or by selecting the **Bypass Access Control policy for decrypted traffic** option in RA VPN connection profiles, the steps that configure access control rules are not needed.

Procedure

Step 1 (Site A, main site.) Configure the site-to-site VPN connection to remote Site B.

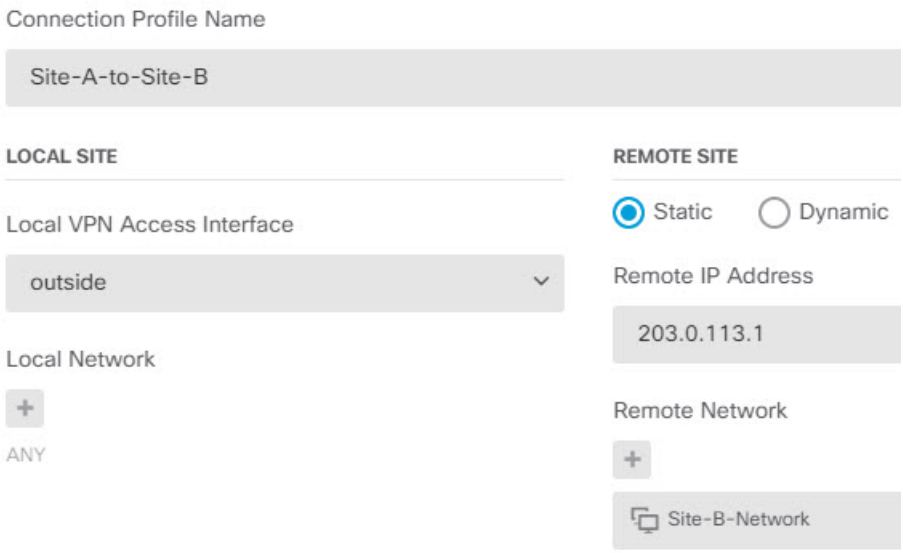
a) Click **Device**, then click **View Configuration** in the Site-to-Site VPN group.

b) Click + to add a new connection.

c) Define the endpoints as follows, and then click **Next**:

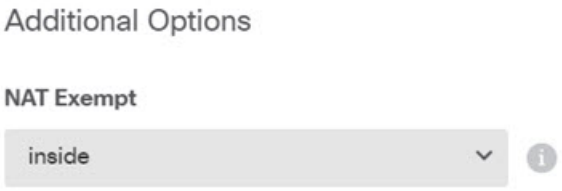
- **Connection Profile Name**—Give the connection a meaningful name, for example, Site-A-to-Site-B.
- **Local VPN Access Interface**—Select the outside interface.
- **Local Network**—Keep the default, Any.
- **Remote IP Address**—Enter the IP address of the remote peer's outside interface. In this example, 203.0.113.1.
- **Remote Network**—Click +, then select the network object that defines the remote peer's protected network. In this example, 192.168.2.0/24. You can click **Create New Network** to create the object now.

The following graphic shows how the first step should look.



d) Define the privacy configuration, then click **Next**.

- **IKE Policy**—The IKE settings have no impact on hair pinning. Simply select the IKE versions, policies, and proposals that fit your security needs. Make note of the local and remote pre-shared keys you enter: you will need these when configuring the remote peer.
- **NAT Exempt**—Select the inside interface.



- **Diffie Helman Group for Perfect Forward Secrecy**—This setting has no impact on hair pinning. Configure it as you see fit.

e) Click **Finish**.

The connection summary is copied to the clipboard. You can paste it into a text file or other document to help you configure the remote peer.

Step 2 (Site A, main site.) Configure the NAT rule to translate all connections going out the outside interface to ports on the outside IP address (interface PAT).

When you complete the initial device configuration, the system creates a NAT rule named InsideOutsideNatRule. This rule applies interface PAT to IPv4 traffic from any interface that exits the device through the outside interface. Because the outside interface is included in "Any" source interface, the rule you need already exists, unless you edited it or deleted it.

The following procedure explains how to create the rule you need.

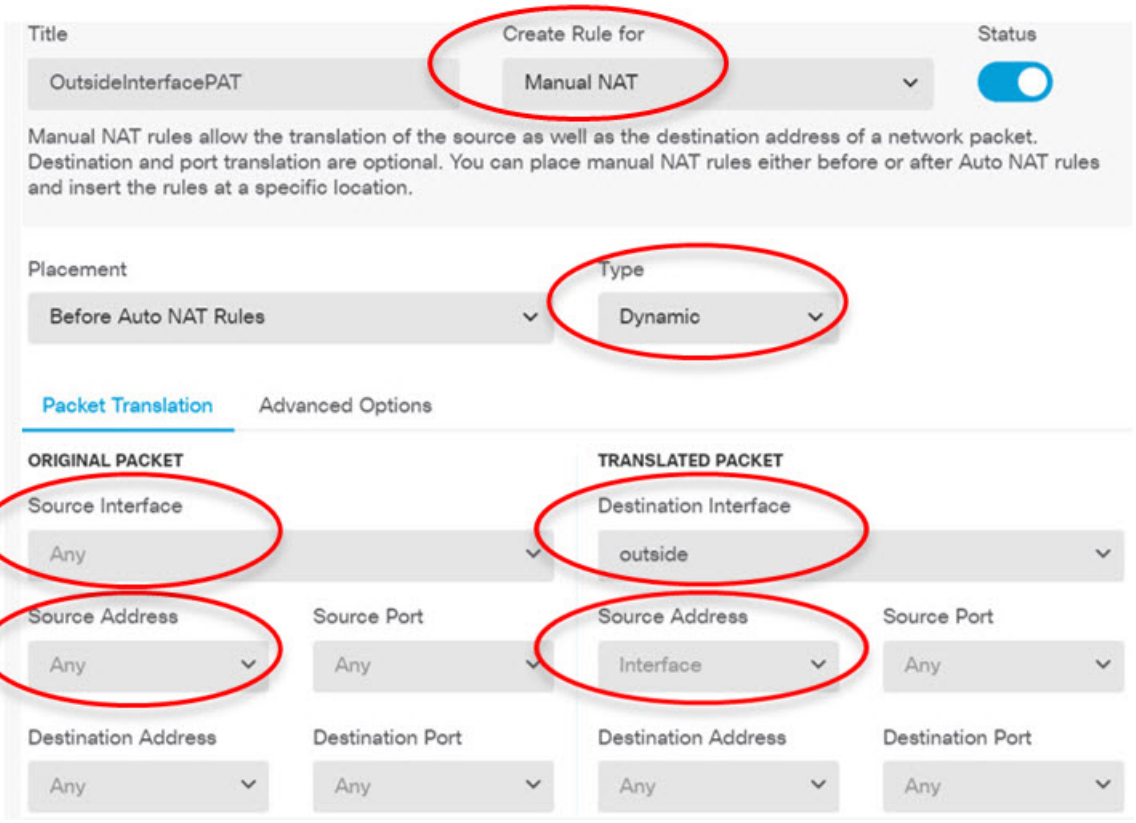
a) Click **Policies** > **NAT**.
b) Do one of the following:

- To edit the InsideOutsideNatRule, mouse over the **Action** column and click the edit icon ( ).
- To create a new rule, click +.

c) Configure a rule with the following properties:

- **Title**—For a new rule, enter a meaningful name without spaces. For example, OutsideInterfacePAT.
- **Create Rule For**—**Manual NAT**.
- **Placement**—**Before Auto NAT Rules** (the default).
- **Type**—**Dynamic**.
- **Original Packet**—For **Source Address**, select either Any or any-ipv4. For **Source Interface**, ensure that you select Any (which is the default). For all other Original Packet options, keep the default, Any.
- **Translated Packet**—For **Destination Interface**, select outside. For **Translated Address**, select **Interface**. For all other Translated Packet options, keep the default, Any.

The following graphic shows the simple case where you select Any for the source address.



d) Click **OK**.

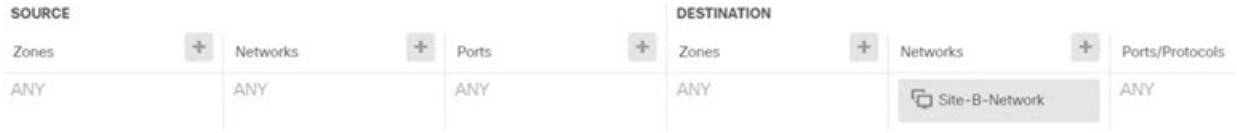
Step 3 (Site A, main site.) Configure an access control rule to allow access to the protected network on Site B.

Simply creating a VPN connection does not automatically allow traffic on the VPN. You need to ensure that your access control policy allows traffic to the remote network.

The following procedure shows how to add a rule specifically for the remote network. Whether you need an additional rule depends on your existing rules.

a) Click **Policies** > **Access Control**.
b) Click **+** to create a new rule.
c) Configure a rule with the following properties:

- **Order**—Select a position in the policy before any other rule that might match these connections and block them. The default is to add the rule to the end of the policy. If you need to reposition the rule later, you can edit this option or simply drag and drop the rule to the right slot in the table.
- **Title**—Enter a meaningful name without spaces. For example, Site-B-Network.
- **Action**—**Allow**. You can select Trust if you do not want this traffic to be inspected for protocol violations or intrusions.
- **Source/Destination** tab—For **Destination** > **Network**, select the same object you used in the VPN connection profile for the remote network. Leave the default, Any, for all other Source and Destination options.



- **Application**, **URL**, and **Users** tabs—Leave the default settings on these tabs, that is, nothing selected.
- **Intrusion**, **File** tabs—You can optionally select intrusion or file policies to inspect for threats or malware.
- **Logging** tab—You can optionally enable connection logging.

d) Click **OK**.

Step 4 (Site A, main site.) Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.



b) Click the **Deploy Now** button.

You can wait until deployment completes, or click **OK** and check the task list or deployment history later. If you leave the window up, it will indicate that there are no pending changes after a successful deployment.

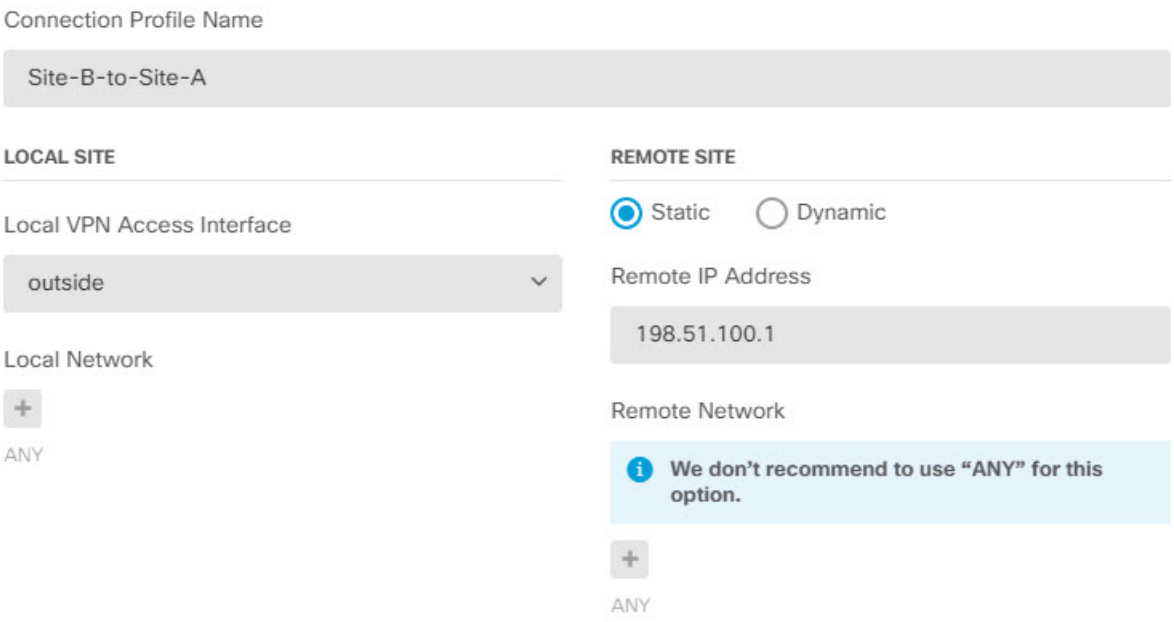
Step 5 (Site B, remote site.) Log into the remote site's device, and configure the site-to-site VPN connection to Site A.

Use the connection summary obtained from the Site A device configuration to help you configure the Site B side of the connection.

a) Click **Device**, then click **View Configuration** in the Site-to-Site VPN group.
b) Click **+** to add a new connection.
c) Define the endpoints as follows, and then click **Next**:

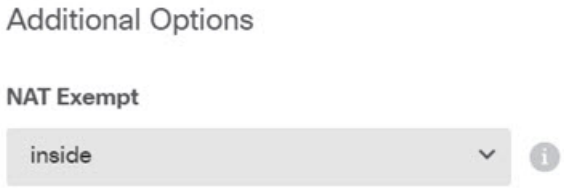
- **Connection Profile Name**—Give the connection a meaningful name, for example, Site-B-to-Site-A.
- **Local VPN Access Interface**—Select the outside interface.
- **Local Network**—Click **+**, then select the network object that defines the local protected network. In this example, 192.168.2.0/24. You can click **Create New Network** to create the object now.
- **Remote IP Address**—Enter the IP address of the main site's outside interface. In this example, 198.51.100.1.
- **Remote Network**—Keep the default, Any. Ignore the warning; it is not relevant for this use case.

The following graphic shows how the first step should look.



d) Define the privacy configuration, then click **Next**.

- **IKE Policy**—The IKE settings have no impact on hair pinning. Configure the same or compatible options as those on Site A's end of the VPN connection. You must configure the pre-shared keys correctly: switch the local and remote keys (for IKEv2) as configured on the Site A device. For IKEv1, there is just one key, which must be the same on both peers.
- **NAT Exempt**—Select the inside interface.



- **Diffie Helman Group for Perfect Forward Secrecy**—This setting has no impact on hair pinning. Match the setting used on Site A's end of the VPN connection.

e) Click **Finish**.

Step 6 (Site B, remote site.) Delete all NAT rules for the protected network so that all traffic leaving the site must go through the VPN tunnel.

Performing NAT on this device is unnecessary because the Site A device will do the address translation. But please examine your specific situation. If you have multiple internal networks and not all of them are participating in this VPN connection, do not delete NAT rules that you need for those networks.

a) Click **Policies** > **NAT**.
b) Do one of the following:

- To delete rules, mouse over the Action column and click the delete icon ( ).

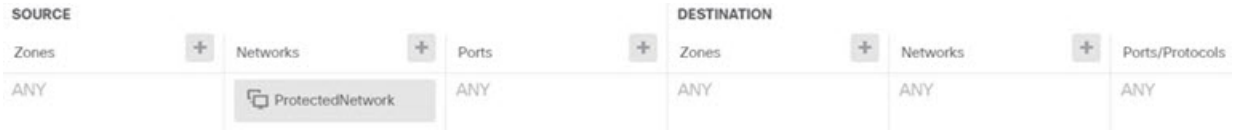
- To edit rules so they no longer apply to the protected network, mouse over the Action column and click the edit icon ( ).

Step 7 (Site B, remote site.) Configure an access control rule to allow access from the protected network to the Internet.

The following example allows traffic from the protected network to any destination. You can adjust this to meet your specific requirements. You can also precede the rule with block rules to filter out undesirable traffic. Another option is to configure the block rules on the Site A device.

a) Click **Policies** > **Access Control**.
b) Click **+** to create a new rule.
c) Configure a rule with the following properties:

- **Order**—Select a position in the policy before any other rule that might match these connections and block them. The default is to add the rule to the end of the policy. If you need to reposition the rule later, you can edit this option or simply drag and drop the rule to the right slot in the table.
- **Title**—Enter a meaningful name without spaces. For example, Protected-Network-to-Any.
- **Action**—**Allow**. You can select Trust if you do not want this traffic to be inspected for protocol violations or intrusions.
- **Source/Destination** tab—For **Source** > **Network**, select the same object you used in the VPN connection profile for the local network. Leave the default, Any, for all other Source and Destination options.



- **Application**, **URL**, and **Users** tabs—Leave the default settings on these tabs, that is, nothing selected.
- **Intrusion**, **File** tabs—You can optionally select intrusion or file policies to inspect for threats or malware.
- **Logging** tab—You can optionally enable connection logging.

d) Click **OK**.

Step 8 (Site B, remote site.) Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.



b) Click the **Deploy Now** button and wait for deployment to finish.

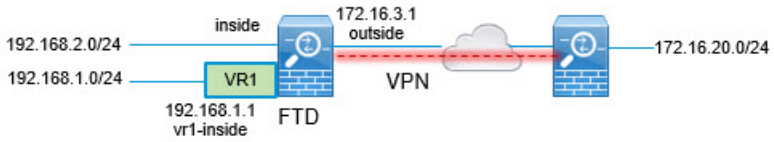
You can wait until deployment completes, or click **OK** and check the task list or deployment history later. If you leave the window up, it will indicate that there are no pending changes after a successful deployment.

How to Secure Traffic from Networks in Multiple Virtual Routers over a Site-to-Site VPN

If you configure multiple virtual routers on a device, you must configure the site-to-site VPN in the global virtual router. You cannot configure site-to-site VPN on an interface that is assigned to a custom virtual router.

Because the routing tables for virtual routers are separate, you must create static routes if you need to secure the connections from or to networks hosted within custom virtual routers over the site-to-site VPN. You also need to update the site-to-site VPN connection to include these additional networks.

Consider the following example. In this case, the site-to-site VPN is defined on the outside interface at 172.16.3.1. This VPN can include the inside network 192.168.2.0/24 without extra configuration, because the inside interface is also part of the global virtual router. But, if you need to provide site-to-site VPN services to the 192.168.1.0/24 network, which is part of the VR1 virtual router, you must configure static routes going both ways, and add the network to the site-to-site VPN configuration.



Before you begin

This example assumes that you have already configured the site-to-site VPN between the 192.168.2.0/24 local network and the 172.16.20.0/24 external network, defined the virtual routers, and configured and assigned the interfaces to the appropriate virtual routers.

Procedure

Step 1 Configure the route leak from the Global virtual router to VR1.

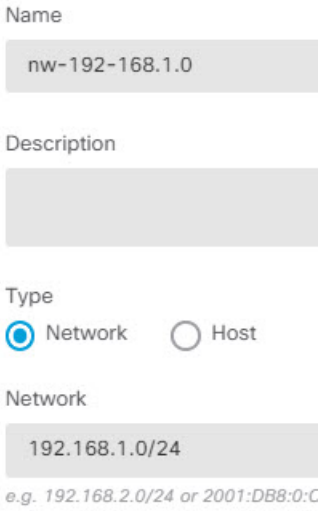
This route allows endpoints protected by the external (remote) end of the site-to-site VPN to access the 192.168.1.0/24 network in the VR1 virtual router.

a) Choose **Device** > **Routing** > **View Configuration**.

b) Click the view icon ( ) for the Global virtual router.

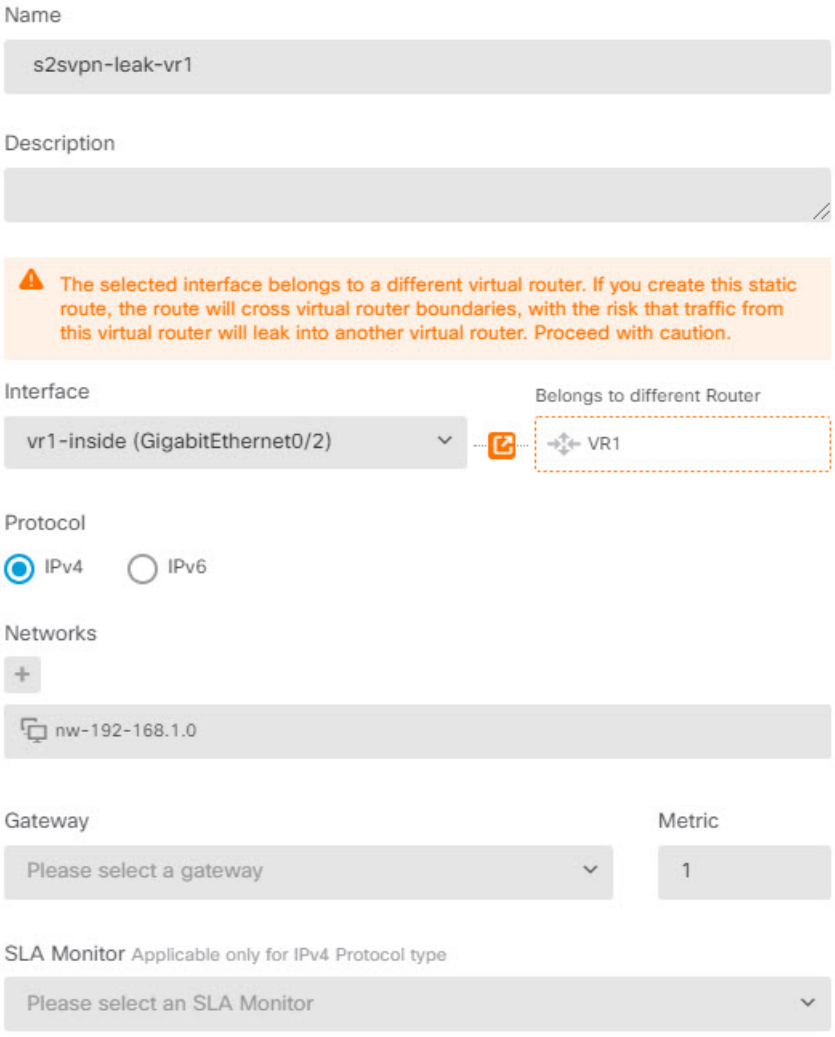
c) On the **Static Routing** tab for the Global router, click + and configure the route:

- **Name**—Any name will do, such as **s2svpn-leak-vr1**.
- **Interface**—Select **vr1-inside**.
- **Protocol**—Select **IPv4**.
- **Networks**—Select an object that defines the 192.168.1.0/24 network. Click **Create New Network** to create the object now if necessary.



- **Gateway**—Leave this item blank. When leaking a route into another virtual router, you do not select the gateway address.

The dialog box should look similar to the following:



d) Click **OK**.

Step 2 Configure the route leak from VR1 to the Global virtual router.

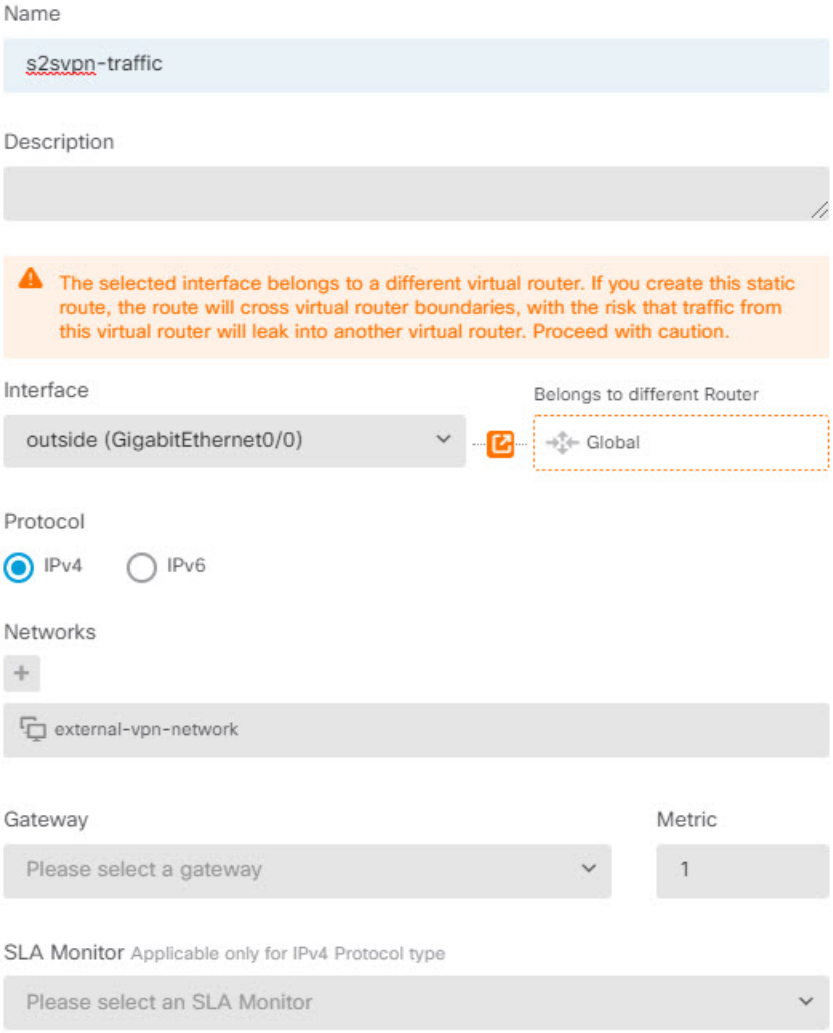
This route allows endpoints on the 192.168.1.0/24 network to initiate connections that will traverse the site-to-site VPN tunnel. For this example, the remote endpoint is protecting the 172.16.20.0/24 network.

a) Choose **VR1** from the virtual routers drop-down list to switch to the VR1 configuration.

b) On the **Static Routing** tab for the VR1 virtual router, click **+** and configure the route:

- **Name**—Any name will do, such as **s2svpn-traffic**.
- **Interface**—Select **outside**.
- **Protocol**—Select **IPv4**.
- **Networks**—Select the object you created for the protected networks of the remote endpoint, for example, **external-vpn-network**.
- **Gateway**—Leave this item blank. When leaking a route into another virtual router, you do not select the gateway address.

The dialog box should look similar to the following:



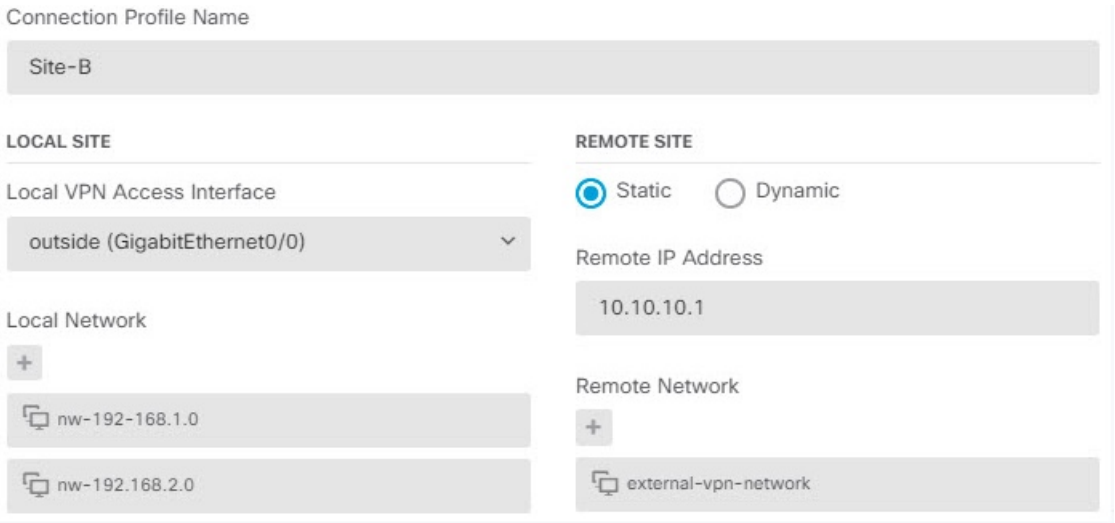
c) Click **OK**.

Step 3 Add the 192.168.1.0/24 network to the site-to-site VPN connection profile.

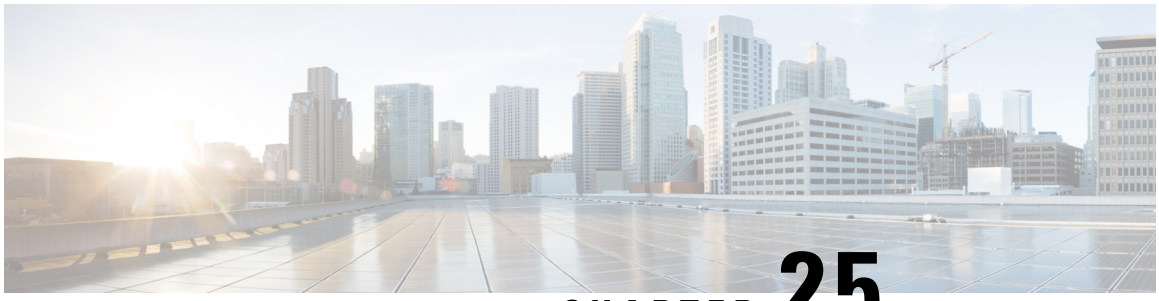
a) Choose **Device** > **Site-to-Site VPN** > **View Configuration**.

b) Click the edit icon ( ) for the connection profile.

c) On the first page of the wizard, click + under **Local Network** and add the object for the 192.168.1.0/24 network.



d) Complete the wizard.



CHAPTER 25

Remote Access VPN

Remote Access virtual private network (VPN) allows individual users to connect to your network from a remote location using a computer or other supported iOS or Android device connected to the Internet. This allows mobile workers to connect from their home networks or a public Wi-Fi network, for example.

The following topics explain how to configure remote access VPN for your network.

- Remote Access VPN Overview, on page 671
- Licensing Requirements for Remote Access VPN, on page 677
- Guidelines and Limitations for Remote Access VPN, on page 678
- Configuring Remote Access VPN, on page 678
- Managing the Remote Access VPN Configuration, on page 684
- Monitoring Remote Access VPN, on page 698
- Troubleshooting Remote Access VPNs, on page 698
- Examples for Remote Access VPN, on page 701

Remote Access VPN Overview

You can use the Firewall Device Manager to configure remote access VPN over SSL using the AnyConnect Client sofware.

When the AnyConnect Client negotiates an SSL VPN connection with the Firewall Threat Defense device, it connects using Transport Layer Security (TLS) or Datagram Transport Layer Security (DTLS). DTLS avoids latency and bandwidth problems associated with some SSL connections and improves the performance of real-time applications that are sensitive to packet delays. The client and the Firewall Threat Defense device negotiate the TLS/DTLS version to use. DTLS is used if the client supports it.

Maximum Concurrent VPN Sessions By Device Model

There is a maximum limit to the number of concurrent VPN sessions allowed on a device based on the device model. This limit is designed so that system performance does not degrade to unacceptable levels. The limit is a combination of remote access VPN users plus site-to-site VPN peers. Use these limits for capacity planning.

Device Model	Maximum Concurrent VPN Sessions
Firepower 1010	75
Firepower 1120	150

Device Model	Maximum Concurrent VPN Sessions
Firepower 1140	400
Firepower 2110	1500
Firepower 2120	3500
Firepower 2130	7500
Firepower 2140	10,000
Secure Firewall 3110	3000
Secure Firewall 3120	6000
Secure Firewall 3130	15,000
Secure Firewall 3140	20,000
Firepower 4100 series, all models	10,000
Firepower 9300 appliance, all models	20,000
Firewall Threat Defense Virtual: FTDv5	50
Firewall Threat Defense Virtual: FTDv10, FTDv20, FTDv30	250
Firewall Threat Defense Virtual: FTDv50	750
Firewall Threat Defense Virtual: FTDv100	10,000
ISA 3000	25

Downloading the AnyConnect Client Software

Before you can configure a remote access VPN, you must download the AnyConnect Client software to your workstation. You will need to upload these packages when defining the VPN.

You should download the latest AnyConnect Client version, to ensure that you have the latest features, bug fixes, and security patches. Regularly update the packages on the Firewall Threat Defense device.

Note You can upload one AnyConnect Client package per operating system: Windows, Mac, and Linux. You cannot upload multiple versions for a given OS type.

Obtain the AnyConnect Client software packages from software.cisco.com. You need to download the “Full Installation Package” versions of the clients.

How Users Can Install the AnyConnect Client Software

To complete a VPN connection, your users must install the AnyConnect Client software. You can use your existing software distribution methods to install the software directly. Or, you can have users install the AnyConnect Client directly from the Firewall Threat Defense device.

Users must have Administrator rights on their workstations to install the software.

Once the AnyConnect Client is installed, if you upload new AnyConnect Client versions to the system, the AnyConnect Client will detect the new version on the next VPN connection the user makes. The system will automatically prompt the user to download and install the updated client software. This automation simplifies software distribution for you and your clients.

If you decide to have users initially install the software from the Firewall Threat Defense device, tell users to perform the following steps.



Note Android and iOS users should download the AnyConnect Client from the appropriate App Store.

Procedure

Step 1 Using a web browser, open **https://***ravpn-address*, where *ravpn-address* is the IP address or hostname of the outside interface on which you are allowing VPN connections.

You identify this interface when you configure the remote access VPN. The system prompts the user to log in.

If you changed the port for remote access VPN connections, users must include the custom port in the URL. For example, if you changed the port to 4443: https://ravpn.example.com:4443

Step 2 Log into the site.

Users are authenticated using the directory server configured for the remote access VPN. Log in must be successful to continue.

If log in is successful, the system determines if the user already has the required version of the AnyConnect Client. If the AnyConnect Client is absent from the user's computer, or is down-level, the system automatically starts installing the AnyConnect Client software.

When installation is finished, AnyConnect Client completes the remote access VPN connection.

Controlling User Permissions and Attributes Using RADIUS and Group Policies

You can apply user authorization attributes (also called user entitlements or permissions) to RA VPN connections from an external RADIUS server or from a group policy defined on the Firewall Threat Defense device. If the Firewall Threat Defense device receives attributes from the external AAA server that conflict with those configured on the group policy, then attributes from the AAA server always take precedence.

The Firewall Threat Defense device applies attributes in the following order:

1. User attributes defined on the external AAA server—The server returns these attributes after successful user authentication or authorization.
2. Group policy configured on the Firewall Threat Defense device—If a RADIUS server returns the value of the RADIUS CLASS attribute IETF-Class-25 (OU= group-policy) for the user, the Firewall Threat Defense device places the user in the group policy of the same name and enforces any attributes in the group policy that are not returned by the server.
3. Group policy assigned by the connection profile—The connection profile has the preliminary settings for the connection, and includes a default group policy applied to the user before authentication. All users connecting to the Firewall Threat Defense device initially belong to this group, which provides any attributes that are missing from the user attributes returned by the AAA server, or the group policy assigned to the user.

Firewall Threat Defense devices support RADIUS attributes with vendor ID 3076. If the RADIUS server you use does not have these attributes defined, you must manually define them. To define an attribute, use the attribute name or number, type, value, and vendor code (3076).

The following topics explain the supported attributes based on whether the values are defined in the RADIUS server, or whether they are values the system sends to the RADIUS server.

Attributes Sent to the RADIUS Server

RADIUS attributes 146 and 150 are sent from the Firewall Threat Defense device to the RADIUS server for authentication and authorization requests. All of the following attributes are sent from the Firewall Threat Defense device to the RADIUS server for accounting start, interim-update, and stop requests.

Table 15: Attributes Firewall Threat Defense Sends to RADIUS

Attribute	Attribute Number	Syntax, Type	Single or Multi-valued	Description or Value
Client Type	150	Integer	Single	The type of client that is connecting to the VPN: 2 = AnyConnect Client SSL VPN
Session Type	151	Integer	Single	The type of connection: 1 = AnyConnect Client SSL VPN
Tunnel Group Name	146	String	Single	The name of the connection profile that was used to establish the session, as defined on the Firewall Threat Defense device. The name can be 1 - 253 characters.

Attributes Received from the RADIUS Server

The following user authorization attributes are sent to the Firewall Threat Defense device from the RADIUS server.

Table 16: RADIUS Attributes Sent to Firewall Threat Defense

Attribute	Attribute Number	Syntax, Type	Single or Multi-valued	Description or Value
Access-List-Inbound	86	String	Single	Both of the Acess-List attributes take the name of an ACL that is configured on the Firewall Threat Defense device. Create these ACLs using the Smart CLI Extended Access List object type (select **Device** > **Advanced Configuration** > **Smart CLI** > **Objects**). These ACLs control traffic flow in the inbound (traffic entering the Firewall Threat Defense device) or outbound (traffic leaving the Firewall Threat Defense device) direction.
Access-List-Outbound	87	String	Single	
Address-Pools	217	String	Single	The name of a network object defined on the Firewall Threat Defense device that identifies a subnet, which will be used as the address pool for clients connecting to the RA VPN. Define the network object on the **Objects** page.
Banner1	15	String	Single	The banner to display when the user logs in.
Banner2	36	String	Single	The second part of the banner to display when the user logs in. Banner2 is appended to Banner1.
Group-Policy	25	String	Single	The group policy to use in the connection. You must create the group policy on the RA VPN **Group Policy** page. You can use one of the following formats: • *group policy name* • OU=*group policy name* • OU=*group policy name*;
Simultaneous-Logins	2	Integer	Single	The number of separate simultaneous connections the user is allowed to establish, 0 - 2147483647.
VLAN	140	Integer	Single	The VLAN on which to confine the user's connection, 0 - 4094. You must also configure this VLAN on a subinterface on the Firewall Threat Defense device.

Two-Factor Authentication

You can configure two-factor authentication for the RA VPN. With two-factor authentication, the user must supply a username and static password, plus an additional item such as an RSA token or a Duo passcode. Two-factor authentication differs from using a second authentication source in that two-factor is configured on a single authentication source, with the relationship to the RSA/Duo server tied to the primary authentication source. The exception is Duo LDAP, where you configure the Duo LDAP server as the secondary authentication source.

The system has been tested with RSA tokens and Duo passcode pushed to mobile for the second factor in conjunction with any RADIUS or AD Server as the first factor in the two-factor authentication process.

RSA Two-Factor Authentication

You can configure RSA using one of the following approaches. See the RSA documentation for information about the RSA-side configuration.

- Define the RSA Server directly in the Firewall Device Manager as a RADIUS server, and use the server as the primary authentication source in the RA VPN.

  When using this approach, the user must authenticate using a username that is configured in the RSA RADIUS server, and concatenate the password with the one-time temporary RSA token, separating the password and token with a comma: *password*,*token*.

  In this configuration, it is typical to use a separate RADIUS server (such as one supplied in Cisco ISE) to provide authorization services. You would configure the second RADIUS server as the authorization and, optionally, accounting server.

- Integrate the RSA server with a RADIUS or AD server that supports direct integration, and configure the RA VPN to use the non-RSA RADIUS or AD server as the primary authentication source. In this case, the RADIUS/AD server uses RSA-SDI to delegate and orchestrate the two-factor authentication between the client and RSA Server.

  When using this approach, the user must authenticate using a username that is configured in the non-RSA RADIUS or AD server, and concatenate the password with the one-time temporary RSA token, separating the password and token with a comma: *password*,*token*.

  In this configuration, you would also use the non-RSA RADIUS server as the authorization and, optionally, accounting server.

Duo Two-Factor Authentication Using RADIUS

You can configure the Duo RADIUS server as the primary authentication source. This approach uses the Duo RADIUS Authentication Proxy.

For the detailed steps to configure Duo, please see https://duo.com/docs/cisco-firepower.

You would then configure Duo to forward authentication requests directed to the proxy server to use another RADIUS server, or an AD server, as the first authentication factor, and the Duo Cloud Service as the second factor.

When using this approach, the user must authenticate using a username that is configured on both the Duo Authentication Proxy and the associated RADIUS/AD server, and the password for the username configured in the RADIUS/AD server, followed by one of the following Duo codes:

- *Duo-passcode*. For example, *my-password*,**12345**.
- **push**. For example, *my-password*,**push**. Use **push** to tell Duo to send a push authentication to the Duo Mobile app, which the user must have already installed and registered.
- **sms**. For example, *my-password*,**sms**. Use **sms** to tell Duo to send an SMS message with a new batch of passcodes to the user's mobile device. The user's authentication attempt will fail when using **sms**. The user must then re-authenticate and enter the new passcode as the secondary factor.
- **phone**. For example, *my-password*,**phone**. Use **phone** to tell Duo to perform phone callback authentication.

If the username/password is authenticated, the Duo Authentication Proxy contacts the Duo Cloud Service, which validates that the request is from a valid configured proxy device and then pushes a temporary passcode to the mobile device of the user as directed. When the user accepts this passcode, the session is marked authenticated by Duo and the RA VPN is established.

Duo Two-Factor Authentication using LDAP

You can use the Duo LDAP server as the secondary authentication source in conjunction with a Microsoft Active Directory (AD) or RADIUS server as the primary source. With Duo LDAP, the secondary authentication validates the primary authentication with a Duo passcode, push notification, or phone call.

The Firewall Threat Defense device communicates with Duo LDAP using LDAPS over port TCP/636.

Note that the Duo LDAP server provides authentication services only, it does not provide identity services. Thus, if you use Duo LDAP as a primary authentication source, you will not see usernames associated with RA VPN connections in any dashboards, and you will not be able to write access control rules for these users.

When using this approach, the user must authenticate using a username that is configured on both the RADIUS/AD server and the Duo LDAP server. When prompted to log in by the AnyConnect Client, the user provides the RADIUS/AD password in the primary **Password** field, and for the **Secondary Password**, provides one of the following to authenticate with Duo. For more details, see https://guide.duo.com/anyconnect.

- *Duo passcode*—Authenticate using a passcode, either generated with Duo Mobile, sent via SMS, generated by your hardware token, or provided by an administrator. For example, 1234567.
- **push**—Push a login request to your phone, if you have installed and activated the Duo Mobile app. Review the request and tap **Approve** to log in.
- **phone**—Authenticate using a phone callback.
- **sms**—Request a Duo passcode in a text message. The login attempt will fail. Log in again using the new passcode.

For a detailed explanation and example of using Duo LDAP, see How to Configure Two-Factor Authentication using Duo LDAP, on page 709.

Licensing Requirements for Remote Access VPN

Your base device license must meet export requirements before you can configure remote access VPN. When you register the device, you must do so with a Smart Software Manager account that is enabled for export-controlled features. You also cannot configure the feature using the evaluation license.

In addition, you need to purchase and enable a remote access VPN license, any of the following: AnyConnect Plus, AnyConnect Apex, or AnyConnect VPN Only. These licenses are treated the same for Firewall Threat Defense devices, even though they are designed to allow different feature sets when used with ASA Software-based headends.

To enable the license, select **Device** > **Smart License** > **View Configuration**, then select the appropriate license in the RA VPN License group. You need to have the license available in your Smart Software Manager account. For more information about enabling licenses, see Enabling or Disabling Optional Licenses, on page 89.

For more information, see the *Cisco AnyConnect Ordering Guide*, http://www.cisco.com/c/dam/en/us/products/collateral/security/anyconnect-og.pdf. There are also other data sheets available on http://www.cisco.com/c/en/us/products/security/anyconnect-secure-mobility-client/datasheet-listing.html.

Guidelines and Limitations for Remote Access VPN

Please keep the following guidelines and limitations in mind when configuring RA VPN.

- You cannot configure both the Firewall Device Manager access (HTTPS access in the management access list) and remote access SSL VPN on the same interface for the same TCP port. For example, if you configure remote access SSL VPN on the outside interface, you cannot also open the outside interface for HTTPS connections on port 443. If you configure both features on the same interface, ensure that you change the HTTPS port for at least one of these services to avoid a conflict.
- The RA VPN outside interface is a global setting. You cannot configure separate connection profiles on different interfaces.
- You cannot use overlapping addresses in the source address of a NAT rule and a remote access VPN address pool.
- If you configure two-factor authentication using RADIUS and RSA tokens, the default authentication timeout of 12 seconds is too quick to allow successful authentication in most cases. You can increase the authentication timeout value by creating a custom AnyConnect Client profile and applying it to the RA VPN connection profile, as described in Configure and Upload Client Profiles, on page 680. We recommend an authentication timeout of at least 60 seconds, so that users have enough time to authenticate and then paste the RSA token, and for the round-trip verification of the token.
- Issuing commands such as **curl** against the RA VPN headend is not directly supported, and might not have desirable results. For example, the headend does not respond to HTTP HEAD requests.
- For RA VPN, you cannot use multiple SAML Server objects with the same Identity Provider (IDP) Entity ID URL, even if other attributes are different. The entity ID must be unique across all SAML server objects used in RA VPN connection profiles.

Configuring Remote Access VPN

To enable remote access VPN for your clients, you need to configure a number of separate items. The following procedure provides the end to end process.

Procedure

Step 1 Configure licenses.

You need to enable two licenses:

- When you register the device, you must do so with a Smart Software Manager account that is enabled for export-controlled features. The base license must meet export control requirements before you can configure remote access VPN. You also cannot configure the feature using the evaluation license. For the procedure to register the device, see Registering the Device, on page 87.
- A remote access VPN license. For details, see Licensing Requirements for Remote Access VPN, on page 677. To enable the license, see Enabling or Disabling Optional Licenses, on page 89.

Step 2 Configure Certificates.

Certificates are required to authenticate SSL connections between the clients and the device. You can use the pre-defined DefaultInternalCertificate for the VPN, or create your own.

If you use an encrypted connection for the directory realm used for authentication, you must upload a trusted CA certificate.

For more information on certificates and how to upload them, see Configuring Certificates, on page 150.

Step 3 (Optional.) Configure TLS/SSL settings.

By default, the system will allow remote users to connect to the remote access VPN using any TLS version and encryption cipher supported by the system. However, you can limit the TLS/DTLS versions, ciphers, and Diffie-Hellman groups allowed to enforce a more secure connection. See Configuring TLS/SSL Cipher Settings, on page 784.

Step 4 (Optional.) Configure and Upload Client Profiles, on page 680.

Step 5 Configure the identity source used for authenticating remote users.

You can use the following sources for user accounts that are allowed to log into the remote access VPN. Alternatively, you can use client certificates for authentication, either alone or in conjunction with an identity source.

- Active Directory identity realm—As a primary authentication source. The user accounts are defined in your Active Directory (AD) server. See Configuring AD Identity Realms, on page 161.
- RADIUS server group—As a primary or secondary authentication source, and for authorization and accounting. See Configure RADIUS Server Groups, on page 167.
- LocalIdentitySource (the local user database)—As a primary or fallback source. You can define users directly on the device and not use an external server. If you use the local database as a fallback source, ensure that you define the same usernames/passwords as the ones defined in the external server. See Configure Local Users, on page 175.
- Duo LDAP server—As a primary or secondary authentication source. Although you can use a Duo LDAP server as the primary source, this is not the normal configuration. You would normally use it as the secondary source to provide two-factor authentication in conjunction with a primary Active Directory or RADIUS server. For details, see How to Configure Two-Factor Authentication using Duo LDAP, on page 709.
- **SAML**—Use a SAML server as the primary authentication. When using SAML, you cannot configure a fallback or a secondary authentication source. For details, see Configure SAML Servers, on page 172.

Step 6 (Optional.) Configure Group Policies for RA VPN, on page 692

The group policy defines user-related attributes. You can configure group policies to provide differential access to resources based on group membership. Alternatively, you can use the default policy for all connections.

Step 7 Configure an RA VPN Connection Profile, on page 685.

Step 8 Allow Traffic Through the Remote Access VPN, on page 682.

Step 9 Verify the Remote Access VPN Configuration, on page 683.

If you encounter problems completing a connection, see Troubleshooting Remote Access VPNs, on page 698.

Step 10 (Optional.) Enable the identity policy and configure a rule for passive authentication.

If you enable passive user authentication, users who logged in through the remote access VPN will be shown in the dashboards, and they will also be available as traffic-matching criteria in policies. If you do not enable passive authentication, RA VPN users will be available only if they match an active authentication policy. You must enable the identity policy to get any username information in the dashboards or for traffic matching.

Configure and Upload Client Profiles

AnyConnect Client profiles are downloaded to clients along with the AnyConnect Client software. These profiles define many client-related options, such as auto connect on startup and auto reconnect, and whether the end user is allowed to change the option from the AnyConnect Client preferences and advanced settings.

If you configure a fully-qualified hostname (FQDN) for the outside interface when configuring the remote access VPN connection, the system creates a client profile for you. This profile enables the default settings. You need to create and upload client profiles only if you want non-default behavior. Note that client profiles are optional: if you do not upload one, AnyConnect Client will use default settings for all profile-controlled options.



Note You must include the Firewall Threat Defense device's outside interface in the VPN profile's server list in order for the AnyConnect Client to display all user controllable settings on the first connection. If you do not add the address or FQDN as a host entry in the profile, then filters do not apply for the session. For example, if you create a certificate match and the certificate properly matches the criteria, but you do not add the device as a host entry in that profile, the certificate match is ignored.

You can create profiles for the AnyConnect Client as well as for a variety of modules that you can optionally use with AnyConnect Client, such as the AMP Enabler. Although you can upload profiles for any of these modules, the Firewall Device Manager supports the creation of the AnyConnect Client Profile only. However, you can upload any kind of profile through the Firewall Device Manager, then use the Firewall Threat Defense API (from API Explorer) to change the profile type for the object. The profiles page shows all profiles of any type, although the list does not indicate the profile type. The procedure explains how to do this.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create AnyConnect Client profile objects while editing a profile property by clicking the **Create New AnyConnect Client Profile** link shown in the object list.

Before you begin

Before you can upload client profiles, you must do the following.

- Download and install the stand-alone AnyConnect Client "Profile Editor - Windows / Standalone installer (MSI)." The installation file is for Windows only, and has the file name anyconnect-profileeditor-win-<version>-k9.msi, where <version> is the AnyConnect Client version (the file name is subject to change). For example, anyconnect-profileeditor-win-4.3.04027-k9.msi. You must also install Java JRE 1.6 (or higher) before installing the profile editor. Obtain the AnyConnect Client profile editor from software.cisco.com. Note that this package contains all of the profile editors, not just the one for the VPN client.
- Use the profile editor to create the profiles you need. You should specify the hostname or IP address of the outside interface in the profile. For detailed information, see the editor's online help.

Procedure

Step 1 Select **Objects**, then select **AnyConnect Client Profiles** from the table of contents.

Step 2 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon ( ) for the object.
- To download the profile associated with an object, click the download icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 3 Enter a name and optionally, a description, for the object.

If you are uploading a module profile, use the object name to indicated the module type to make it easier for you to distinguish it from the AnyConnect Client Profiles.

Step 4 Click **Upload** and select the file you created using the Profile Editor.

Step 5 Click **Open** to upload the profile.

Step 6 Click **OK** to add the object.

Step 7 If the profile you created is actually a different type than AnyConnect Client Profile, complete the following steps to alter the profile type of the object.

a) Click the more options button ( ) and choose **API Explorer**.

The system opens the API Explorer in a separate tab or window, depending on your browser settings.

b) Open the AnyConnectClientProfile resource.

c) Select the GET /object/anyconnectclientprofiles method and click the **Try It Out!** button.

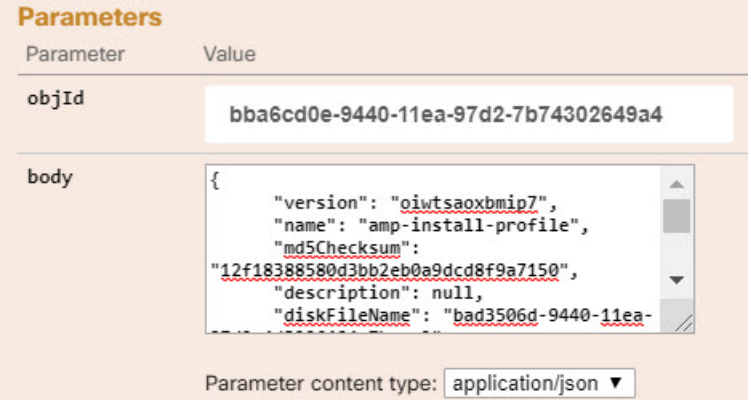
Each profile object will be represented like the following. The highlighted attribute is the one you need to change.

```
{
      "version": "oiwtsaoxbmip7",
      "name": "amp-install-profile",
      "md5Checksum": "12f18388580d3bb2eb0a9dcd8f9a7150",
      "description": null,
      "diskFileName": "bad3506d-9440-11ea-97d2-4d3296494e7b.xml",
      "anyConnectModuleType": "ANY_CONNECT_CLIENT_PROFILE",
      "id": "bba6cd0e-9440-11ea-97d2-7b74302649a4",
      "type": "anyconnectclientprofile",
      "links": {
        "self": "https://10.89.5.38/api/fdm/v6/object/
anyconnectclientprofiles/bba6cd0e-9440-11ea-97d2-7b74302649a4"
      }
```

d) Find your object in the output, select the code and use Ctrl+click to copy it to the clipboard.

e) Select the PUT /object/anyconnectclientprofiles/{objId} method, and paste the contents into the **body** field.

f) Copy the **id** value and paste it into the **objId** edit box above the body. You can also find the object ID at the end of the “self” URL.



g) In the object body, find the **anyConnectModuleType** field and replace the value with the one for your profile type. Choose from DART, FEEDBACK, WEB_SECURITY, ANY_CONNECT_CLIENT_PROFILE, AMP_ENABLER, NETWORK_ACCESS_MANAGER, NETWORK_VISIBILITY, START_BEFORE_LOGIN, ISE_POSTURE, UMBRELLA.

h) Again in the **body**, delete the **links** attribute, including the comma after the **type** value.

The object body should look similar to the following:

```
{
      "version": "oiwtsaoxbmip7",
      "name": "amp-install-profile",
      "md5Checksum": "12f18388580d3bb2eb0a9dcd8f9a7150",
      "description": null,
      "diskFileName": "bad3506d-9440-11ea-97d2-4d3296494e7b.xml",
      "anyConnectModuleType": "AMP_ENABLER",
      "id": "bba6cd0e-9440-11ea-97d2-7b74302649a4",
      "type": "anyconnectclientprofile"
      }
```

i) Click **Try It Out!** Examine the response to verify that the object was modified correctly. You should get a response code of 200, and a response body that echos your changes. You can use the GET method to further verify the results.

Allow Traffic Through the Remote Access VPN

You can use one of the following techniques to enable traffic flow in the remote access VPN tunnel.

- Configure the **sysopt connection permit-vpn** command, which exempts traffic that matches the VPN connection from the access control policy. The default for this command is **no sysopt connection permit-vpn**, which means VPN traffic must also be allowed by the access control policy.

  This is the more secure method to allow traffic in the VPN, because external users cannot spoof IP addresses in the remote access VPN address pool. The downside is that the VPN traffic will not be inspected, which means that intrusion and file protection, URL filtering, or other advanced features will not be applied to the traffic. This also means that no connection events will be generated for the traffic, and thus statistical dashboards will not reflect VPN connections.

  To configure this command, select the **Bypass Access Control policy for decrypted traffic** option in your RA VPN connection profiles.

- Create access control rules to allow connections from the remote access VPN address pool. This method ensures that VPN traffic is inspected and advanced services can be applied to the connections. The downside is that it opens the possibility for external users to spoof IP addresses and thus gain access to your internal network.

Verify the Remote Access VPN Configuration

After you configure the remote access VPN, and deploy the configuration to the device, verify that you can make remote connections.

If you encounter problems, read through the troubleshooting topics to help isolate and correct the problems. See Troubleshooting Remote Access VPNs, on page 698.

Procedure

Step 1 From an external network, establish a VPN connection using the AnyConnect Client.

Using a web browser, open **https://***ravpn-address*, where *ravpn-address* is the IP address or hostname of the outside interface on which you are allowing VPN connections. If necessary, install the client software and complete the connection. See How Users Can Install the AnyConnect Client Software, on page 673.

If you changed the port for remote access VPN connections, you must include the custom port in the URL. For example, if you changed the port to 4443: https://ravpn.example.com:4443

If you configured group URLs, also try those URLs.

Step 2 Log into the device CLI as explained in Logging Into the Command Line Interface (CLI), on page 15. Alternatively, open the CLI Console.

Step 3 Use the **show vpn-sessiondb** command to view summary information about current VPN sessions.

The statistics should show your active AnyConnect Client session, and information on cumulative sessions, the peak concurrent number of sessions, and inactive sessions. Following is sample output from the command.

```
> show vpn-sessiondb
----------------------------------------------------------------------------
VPN Session Summary
----------------------------------------------------------------------------
                            Active : Cumulative : Peak Concur : Inactive
                           ------------------------------------------------
AnyConnect Client          :      1 :         49 :           3 :        0
  SSL/TLS/DTLS             :      1 :         49 :           3 :        0
Clientless VPN             :      0 :          1 :           1
  Browser                  :      0 :          1 :           1
----------------------------------------------------------------------------
Total Active and Inactive  :      1             Total Cumulative :     50
Device Total VPN Capacity  :  10000
Device Load                :     0%
----------------------------------------------------------------------------

----------------------------------------------------------------------------
Tunnels Summary
----------------------------------------------------------------------------
                            Active : Cumulative : Peak Concurrent
                           ------------------------------------------------
Clientless                 :      0 :          1 :               1
```

```
AnyConnect-Parent            :      1 :         49 :               3
SSL-Tunnel                   :      1 :         46 :               3
DTLS-Tunnel                  :      1 :         46 :               3
------------------------------------------------------------------------------
Totals                       :      3 :        142
------------------------------------------------------------------------------

------------------------------------------------------------------------------
IPv6 Usage Summary
------------------------------------------------------------------------------
                               Active : Cumulative : Peak Concurrent
                              ---------------------------------------------
AnyConnect SSL/TLS/DTLS      :        :            :
  Tunneled IPv6              :      1 :         20 :           2
------------------------------------------------------------------------------
```

Step 4 Use the **show vpn-sessiondb anyconnect** command to view detailed information about current VPN sessions.

Detailed information includes encryption used, bytes transmitted and received, and other statistics. If you use your VPN connection, you should see the bytes transmitted/received numbers change as you re-issue this command.

```
> show vpn-sessiondb anyconnect

Session Type: AnyConnect

Username     : priya                  Index        : 4820
Assigned IP  : 172.18.0.1             Public IP    : 192.168.2.20
Assigned IPv6: 2009::1
Protocol     : AnyConnect-Parent SSL-Tunnel DTLS-Tunnel
License      : AnyConnect Premium
Encryption   : AnyConnect-Parent: (1)none  SSL-Tunnel: (1)AES-GCM-256  DTLS-Tunnel: (1)AES256
Hashing      : AnyConnect-Parent: (1)none  SSL-Tunnel: (1)SHA384  DTLS-Tunnel: (1)SHA1
Bytes Tx     : 27731                  Bytes Rx     : 14427
Group Policy : MyRaVpn|Policy         Tunnel Group : MyRaVpn
Login Time   : 21:58:10 UTC Mon Apr 10 2017
Duration     : 0h:51m:13s
Inactivity   : 0h:00m:00s
VLAN Mapping : N/A                    VLAN         : none
Audt Sess ID : c0a800fd012d400058ebfff2
Security Grp : none                   Tunnel Zone  : 0
```

Managing the Remote Access VPN Configuration

Remote access VPN connection profiles define the characteristics that allow external users to make a VPN connection to the system using the AnyConnect Client. Each profile defines the AAA servers and certificates used to authenticate users, the address pool for assigning users IP addresses, and the group policies that define a variety of user-oriented attributes.

You would create multiple profiles if you need to provide variable services to different user groups, or if you have different authentication sources. For example, if your organization merges with a different organization that uses different authentication servers, you can create a profile for the new group that uses those authentication servers.

Procedure

Step 1 Click **View Configuration** in the **Device** > **Remote Access VPN** group.

The group shows summary information on how many connection profiles and group policies are currently configured.

Step 2 Click **Connection Profiles** in the table of contents if it is not already selected.

Step 3 Do any of the following:

- Click the + button to create a new connection profile. For detailed instructions, see Configure an RA VPN Connection Profile, on page 685.
- Click the view button ( ) to open a summary of the connection profile and connection instructions. Within the summary, you can click **Edit** to make changes.
- Click the delete button ( ) to delete a connection profile that you no longer need.
- Select **Group Policies** in the table of contents to define the user-oriented attributes for the connection profiles. See Configure Group Policies for RA VPN, on page 692.

Configure an RA VPN Connection Profile

You can create a remote access VPN connection profile to allow your users to connect to your inside networks when they are on external networks, such as their home network. Create separate profiles to accommodate different authentication methods.

Before you begin

Before configuring the remote access (RA) VPN connection:

- Download the required AnyConnect Client software packages from software.cisco.com to your workstation.
- The outside interface, the one that terminates remote access VPN connections, cannot also have a management access list that allows HTTPS connections on the same port. Either configure a different port for management access (see Configuring the HTTPS Port for Management Access on Data Interfaces, on page 749), or configure a different port for the connection profile. Both services use port 443 by default, so one must change.

Procedure

Step 1 Click **View Configuration** in the **Device** > **Remote Access VPN** group.

The group shows summary information on how many connection profiles and group policies are currently configured.

Step 2 Click **Connection Profiles** in the table of contents if it is not already selected.

Step 3 Do one of the following:

- Click the + button to create a new connection profile.
- Click the view button ( ) to open a summary of the connection profile and connection instructions. Within the summary, you can click **Edit** to make changes.

Step 4 Configure the basic connection attributes.

- **Connection Profile Name**—The name for this connection, up to 50 characters without spaces. For example, MainOffice. You cannot use an IP address as the name.

  Note
  The name you enter here is what users will see in the connection list in the AnyConnect Client client. Choose a name that will make sense to your users.

- **Group Alias**, **Group URL**—Aliases contain alternate names or URLs for a specific connection profile. VPN users can choose an alias name in the AnyConnect Client client in the list of connections when they connect to the Firewall Threat Defense device. The connection profile name is automatically added as a group alias. Aliases can be up to 31 characters.

  You can also configure the list of group URLs, which your endpoints can select while initiating the Remote Access VPN connection. If users connect using the group URL, the system will automatically use the connection profile that matches the URL. This URL would be used by clients who do not yet have the AnyConnect Client client installed.

  Add as many group aliases and URLs as required. These aliases and URLs must be unique across all connection profiles defined on the device. Group URLs must start with **https://**.

  For example, you might have the alias Contractor and the group URL https://ravpn.example.com/contractor. Once the AnyConnect Client client is installed, the user would simply select the group alias in the AnyConnect Client VPN drop-down list of connections.

Step 5 Configure the primary and optionally, secondary identity sources.

These options determine how remote users authenticate to the device to enable the remote access VPN connection. The simplest approach is to use AAA only and then select an AD realm or use the LocalIdentitySource. You can use the following approaches for **Authentication Type**:

- **AAA Only**—Authenticate and authorize users based on username and password. For details, see Configure AAA for a Connection Profile, on page 688.
- **Client Certificate Only**—Authenticate users based on client device identity certificate. For details, see Configure Certificate Authentication for a Connection Profile, on page 691.
- **AAA and ClientCertificate**—Use both username/password and client device identity certificate.
- **SAML**—Use a SAML server as the primary authentication. When using SAML, you cannot configure a fallback or a secondary authentication source. For details, see Configure AAA for a Connection Profile, on page 688.

Step 6 Configure the address pool for clients.

The address pool defines the IP addresses that the system can assign to remote clients when they establish a VPN connection. For more information, see Configure Client Addressing for RA VPN, on page 692.

Step 7 Click **Next**.

Step 8 Select the **Group Policy** to use for this profile.

The group policy sets terms for user connections after the tunnel is established. The system includes a default group policy named DfltGrpPolicy. You can create additional group policies to provide the services you require.

When you select a group policy, you are shown a summary of the group characteristics. Click **Edit** in the summary to make changes.

If the group policy you need does not yet exist, click **Create New Group Policy** in the drop-down list.

For detailed information about group policies, see Configure Group Policies for RA VPN, on page 692.

Step 9 Click **Next**.

Step 10 Configure the global settings.

These options apply to every connection profile. After you create the first connection profile, these options are pre-configured for each subsequent profile. If you make changes, you are changing every configured connection profile.

- **Certificate of Device Identity**—Select the internal certificate used to establish the identity of the device. Clients must accept this certificate to complete a secure VPN connection. If you do not already have a certificate, click **Create New Internal Certificate** in the drop-down list. You must configure a certificate.
- **Outside Interface**—The interface to which users connect when making the remote access VPN connection. Although this is normally the outside (Internet-facing) interface, choose whichever interface is between the device and the end users you are supporting.
- **Fully-qualified Domain Name for the Outside Interface**—The name of the interface, for example, ravpn.example.com. If you specify a name, the system can create a client profile for you.

  Note
  You are responsible for ensuring that the DNS servers used in the VPN and by clients can resolve this name to the outside interface's IP address. Add the FQDN to the relevant DNS servers.

- **Port**—The TCP port to use for RA VPN connections. The default is 443. If you need to connect to the Firewall Device Manager on the same interface used for RA VPN, you must change the port number either for the connection profile or for the Firewall Device Manager. Both services use 443 by default. Note that users will have to include the port number in the URL if you change the port for remote access VPN connections.
- **Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)**—Whether to subject VPN traffic to the access control policy. Decrypted VPN traffic is subjected to access control policy inspection by default. Enabling the **Bypass Access Control policy for decrypted traffic** option bypasses the access control policy, but for remote access VPN, the VPN Filter ACL and the authorization ACL downloaded from the AAA server are still applied to VPN traffic.

  Note that if you select this option, the system configures the **sysopt connection permit-vpn** command, which is a global setting. This will also impact the behavior of site-to-site VPN connections. Also, you cannot make different selections for this option across your connection profiles: the feature is either on or off for all profiles.

  If you do not select this option, it might be possible for external users to spoof IP addresses in your remote access VPN address pool, and thus gain access to your network. This can happen because you will need to create access control rules that allow your address pool to have access to internal resources. If you use

access control rules, consider using user specifications to control access, rather than source IP address alone.

The downside of selecting this option is that the VPN traffic will not be inspected, which means that intrusion and file protection, URL filtering, or other advanced features will not be applied to the traffic. This also means that no connection events will be generated for the traffic, and thus statistical dashboards will not reflect VPN connections.

- **NAT Exempt**—Enable NAT Exempt to exempt traffic to and from the remote access VPN endpoints from NAT translation. If you do not exempt VPN traffic from NAT, ensure that the existing NAT rules for the outside and inside interfaces do not apply to the RA VPN pool of addresses. NAT exempt rules are manual static identity NAT rules for a given source/destination interface and network combination, but they are not reflected in the NAT policy, they are hidden. If you enable NAT Exempt, you must also configure the following.

  Note that this is a global option; it applies to all connection profiles. Thus, simply add interfaces and inside networks, do not replace them, or you will be changing the NAT exempt settings for all the other connection profiles that you have already defined.

  - **Inside Interfaces**—Select the interfaces for the internal networks remote users will be accessing. NAT rules are created for these interfaces.
  - **Inside Networks**—Select the network objects that represent internal networks remote users will be accessing. The networks list must contain the same IP types as the address pools you are supporting.

- **AnyConnect Packages**—The AnyConnect Client full installation software images that you will support on RA VPN connections. For each package, the filename, including extensions, can be no more than 60 characters. You can upload separate packages for Windows, Mac, and Linux endpoints. However, you cannot configure different packages for different connection profiles. If you already configured a package for another profile, the package is pre-selected. Changing it will change it for all profiles.

  Download the packages from software.cisco.com. If the endpoint does not already have the right package installed, the system prompts the user to download and install the package after the user authenticates.

Step 11 Click **Next**.

Step 12 Review the summary.

First, verify that the summary is correct.

Then, click **Instructions** to see what end users need to do to initially install the AnyConnect Client software and test that they can complete a VPN connection. Click **Copy** to copy these instructions to the clipboard, and then distribute them to your users.

Step 13 Click **Finish**.

What to do next

Ensure that traffic is allowed in the VPN tunnel, as explained in Allow Traffic Through the Remote Access VPN, on page 682.

Configure AAA for a Connection Profile

Authentication, Authorization, and Accounting (AAA) servers use username and password to determine if a user is allowed access to the remote access VPN. If you use RADIUS servers, you can distinguish authorization

levels among authenticated users, to provide differential access to protected resources. You can also use RADIUS accounting services to keep track of usage.

When configuring AAA, you must configure a primary identity source. Secondary and fallback sources are optional. Use a secondary source if you want to implement dual authentication, for example, using RSA tokens or DUO.

Primary Identity Source Options

- **Primary Identity Source for User Authentication**—The primary identity source used for authenticating remote users. End users must be defined in this source or the optional fallback source to complete a VPN connection. Select one of the following:
  - An Active Directory (AD) identity realm. If the realm you need does not yet exist, click **Create New Identity Realm**.
  - A RADIUS server group.
  - LocalIdentitySource (the local user database)—You can define users directly on the device and not use an external server.
  - A Duo LDAP server. However, this is best used as a secondary authentication source to provide two-factor authentication, as described in How to Configure Two-Factor Authentication using Duo LDAP, on page 709. If you use it as a primary source, you will not get user identity information, and you will not see user information in the dashboards, nor will you be able to write user-based access control rules.
  - A SAML server. If you use a SAML server, you cannot configure a fallback or secondary authentication source. You can use RADIUS as an authorization server, but you must configure the RADIUS server so that authentication is not required. That is, so the RADIUS server will provide authorization information after the connection is authenticated by SAML.
- **SAML Login Experience**—If you select SAML as the primary authentication source, you need to select which client browser to use to complete the web authentication:
  - **VPN Client embedded browser**—The VPN client uses its embedded browser for web authentication, so the authentication applies to the VPN connection only. This is the default and requires no further configuration.
  - **Default OS Browser**—The VPN client uses the system's default browser for web authentication. This option enables single sign-on (SSO) between your VPN authentication and other corporate logins. Also choose this option if you want to support web authentication methods, such as biometric authentication, that cannot be performed in the embedded browser.

    You must upload a package that enables web authentication in the browser. Obtain packages from software.cisco.com. Note that the package you upload is used by all connection profiles that use SAML with the default OS browser; the packages are global, not connection-profile specific.
- **Fallback Local Identity Source**—If the primary source is an external server, you can select the LocalIdentitySource as a fallback in case the primary server is unavailable. If you use the local database as a fallback source, ensure that you define the same local usernames/passwords as the ones defined in the external server.

Advanced Options—Click the **Advanced** link and configure the following options:

- **Strip options**—A realm is an administrative domain. Enabling the following options allows the authentication to be based on the username alone. You can enable any combination of these options. However, you must select both check boxes if your server cannot parse delimiters.
  - **Strip Identity Source Server from Username**—Whether to remove the identity source name from the username before passing the username on to the AAA server. For example, if you select this option and the user enters domain\username as the username, the domain is stripped off from the username and sent to AAA server for authentication. By default this option is unchecked.
  - **Strip Group from Username**—Whether to remove the group name from the username before passing the username on to the AAA server. This option applies to names given in the username@domain format; the option strips the domain and @ sign. By default this option is unchecked.
- **Enable Password Management**—Whether to allow the user to change the password when it expires. If you do not select this option, when the user's password expires, the AnyConnect Client will refuse the connection and the user must go and change the password on the AAA server. If you select this option, AnyConnect Client prompts the user to change the password when it expires, which is much more convenient for the user. Select one of the following options. Also, ensure that you enable MSCHAPv2 on the AAA server.
  - **Notify user x days prior to password expiration** (LDAP only)—Starting the number of days you specify, warn the user of the upcoming password expiration. You can set the warning from 1-180 days, with 14 being the default.
  - **Notify user on the day of password expiration**—The user is not warned, but is still prompted to change the password when the password expires. Even if you set a warning period, RADIUS users always get this behavior.

Secondary Identity Source

- **Secondary Identity Source for User Authorization**—The optional second identity source. If the user successfully authenticates with the primary source, the user is prompted to authenticate with the secondary source. You can select an AD realm, RADIUS server group, Duo LDAP server, or the local identity source.
- **Advanced options**—Click the **Advanced** link and configure the following options:
  - **Fallback Local Identity Source for Secondary**—If the secondary source is an external server, you can select the LocalIdentitySource as a fallback in case the secondary server is unavailable. If you use the local database as a fallback source, ensure that you define the same local usernames/passwords as the ones defined in the secondary external server.
  - **Use Primary Username for Secondary Login**—By default, when using a secondary identity source, the system will prompt for both username and password for the secondary source. If you select this option, the system prompts for the secondary password only, and uses the same username for the secondary source that was authenticated against the primary identity source. Select this option if you configure the same usernames in both the primary and secondary identity sources.
  - **Username for Session Server**—After successful authentication, the username is shown in events and statistical dashboards, is used to determine matches for user- or group-based SSL decryption and access control rules, and is used for accounting. Because you are using two authentication sources, you need to tell the system whether to use the Primary or Secondary username as the user identity. By default, the primary name is used.

- **Password Type**—How to obtain the password for the secondary server. This field applies only if you select **AAA and Client Certificate** for the authentication type, and for the certificate options, you select both **Prefill username from certificate on user login window** and **Hide username in login window**. The default is **Prompt**, which means the user is asked to enter the password.

  Select **Primary Identity Source Password** to automatically use the password entered when the user authenticated to the primary server.

  Select **Common Password** to use the same password for every user, then enter that password in the **Common Password** field.

Additional Options

- **Authorization Server**—The RADIUS server group that has been configured to authorize remote access VPN users.

  After authentication is complete, authorization controls the services and commands available to each authenticated user. Authorization works by assembling a set of attributes that describe what the user is authorized to perform, their actual capabilities, and restrictions. Were you not to use authorization, authentication alone would provide the same access to all authenticated users. For information on configuring RADIUS for authorization, see Controlling User Permissions and Attributes Using RADIUS and Group Policies, on page 673.

  Note that if the system obtains authorization attributes from the RADIUS server that overlap those defined in the group policy, the RADIUS attributes override the group policy attributes.

- **Accounting Server**—(Optional.) The RADIUS server group to use to account for the remote access VPN session.

  Accounting tracks the services users are accessing as well as the amount of network resources they are consuming. The Firewall Threat Defense device reports user activity to the RADIUS server. Accounting information includes when sessions start and stop, usernames, the number of bytes that pass through the device for each session, the service used, and the duration of each session. You can then analyze the data for network management, client billing, or auditing. You can use accounting alone or together with authentication and authorization.

Configure Certificate Authentication for a Connection Profile

You can use certificates installed on the client device to authenticate remote access VPN connections. When using certificate authentication, ensure that the trusted CA certificate used to validate the remote access user connections includes the **SSL Client** option for **Validation Usage**.

When using client certificates, you can still configure a secondary identity source, fallback source, and authorization and accounting servers. These are AAA options; for details, see Configure AAA for a Connection Profile, on page 688.

Following are the certificate-specific attributes. You can configure these attributes separately for the primary and secondary identity sources. Configuring a secondary source is optional.

- **Username from Certificate**—Select one of the following:
  - **Map Specific Field**—Use the certificate elements in the order of **Primary Field** and **Secondary Field**. The defaults are CN (Common Name) and OU (Organizational Unit). Select the options that work for your organization. The fields are combined to provide the username, and this is the name used in events, dashboards, and for matching purposes in SSL decryption and access control rules.

- **Use entire DN (distinguished name) as username**—The system automatically derives the username from the DN fields.

- **Advanced options**—Click the **Advanced** link and configure the following options:
  - **Prefill username from certificate on user login window**—Whether to fill in the username field with the retrieved username when prompting the user to authenticate.
  - **Hide username in login window**—If you select the **Prefill** option, you can hide the username, which means the user cannot edit the username in the password prompt.

Configure Client Addressing for RA VPN

There must be a way for the system to provide an IP address to endpoints that connect to the remote access VPN. These addresses can be provided by the AAA server, a DHCP server, an IP address pool configured in the group policy, or an IP address pool configured in the connection profile. The system tries these resources in that order, and stops when it obtains an available address, which it then assigns to the client. Thus, you can configure multiple options to create a failsafe in case of an unusual number of concurrent connections.

Use one or more of the following methods to configure the address pool for a connection profile.

- **AAA Server**—First, configure a network object on the Firewall Threat Defense device that specifies a subnet for the address pool. Then, in the RADIUS server, configure the Address-Pools (217) attribute for the user with the object name. Also, specify the RADIUS server for authentication in the connection profile.
- **DHCP**—First, configure a DHCP server with one or more IPv4 address ranges for the RA VPN (you cannot configure IPv6 pools using DHCP). Then, create a host network object with the IP address of the DHCP server. You can then select this object in the **DHCP Servers** attribute of the connection profile. You can configure up to 10 DHCP servers.

  If the DHCP server has multiple address pools, you can use the **DHCP Scope** attribute in the group policy that you attach to the connection profile to select which pool to use. Create a host network object with the network address of the pool. For example, if the DHCP pool contains 192.168.15.0/24 and 192.168.16.0/24, setting the DHCP scope to 192.168.16.0 will ensure that an address from the 192.168.16.0/24 subnet will be selected.
- **Local IP address pools**—First, create up to six network objects that specify subnets. You can configure separate pools for IPv4 and IPv6. Then, select these objects in the **IPv4 Address Pool** and **IPv6 Address Pool** options, either in the group policy, or in the connection profile. You do not need to configure both IPv4 and IPv6, just configure the address scheme you want to support.

  You also do not need to configure the pool in both the group policy and the connection profile. The group policy overrides the connection profile settings, so if you configure the pools in the group policy, leave the options empty in the connection profile.

  Note that the pools are used in the order in which you list them.

Configure Group Policies for RA VPN

A group policy is a set of user-oriented attribute/value pairs for remote access VPN connections The connection profile uses a group policy that sets terms for user connections after the tunnel is established. Group policies

let you apply whole sets of attributes to a user or a group of users, rather than having to specify each attribute individually for each user.

The system includes a default group policy named DfltGrpPolicy. You can create additional group policies to provide the services you require.

Procedure

Step 1 Click **View Configuration** in the **Device** > **Remote Access VPN** group.

The group shows summary information on how many connection profiles and group policies are currently configured.

Step 2 Click **Group Policies** in the table of contents.

Step 3 Do any of the following:

- Click the + button to create a new group. See the following topics for explanations of the attributes on the pages of the group policy:
  - General Attributes, on page 693
  - Session Settings Attributes, on page 694
  - Address Assignment Attributes, on page 694
  - Split Tunneling Attributes, on page 695
  - AnyConnect Client Attributes, on page 696
  - Traffic Filters Attributes, on page 697
  - Windows Browser Proxy Attributes, on page 698
- Click the edit button ( ) to edit an existing group policy.
- Click the delete button ( ) to delete a group that you no longer need. The group cannot be currently used in a connection profile.

General Attributes

The general attributes of a group policy define the name of the group and some other basic settings. The Name attribute is the only required attribute.

- **Name**—The name of the group policy. The name can be up to 64 characters, spaces are allowed.
- **Description**—A description of the group policy. The description can be up to 1,024 characters.
- **DNS Server**—Select the DNS server group that defines the DNS servers clients should use for domain name resolution when connected to the VPN. If the group you need is not yet defined, click **Create DNS Group** and create it now.

- **Banner**—The banner text, or welcome message, to present to users at login. The default is no banner. The length can be up to 496 characters. The AnyConnect Client supports partial HTML. To ensure that the banner displays properly to remote users, use the
 tag to indicate line breaks.
- **Default Domain**—The default domain name for users in the RA VPN. For example, example.com. This domain is added to hostnames that are not fully-qualified, for example, serverA instead of serverA.example.com.
- **AnyConnect Client Profiles**—Click **+** and select the AnyConnect Client Profiles to use for this group. If you configure a fully-qualified domain name for the outside interface (in the connection profile), a default profile will be created for you. Alternatively, you can upload your own client profile. Create these profiles using the standalone AnyConnect Client Profile Editor, which you can download and install from software.cisco.com. If you do not select a client profile, the AnyConnect Client uses default values for all options. The items in this list are AnyConnect Client Profile objects rather than the profiles themselves. You can create (and upload) new profiles by clicking **Create New AnyConnect Client Profile** in the drop-down list.

  You can select AnyConnect Client module profiles, such as AMP Enabler, in addition to the AnyConnect Client Profile. You can select one profile per module type.

Session Settings Attributes

The session settings of a group policy control how long users can connect through the VPN and how many separate connections they can establish.

- **Maximum Connection Time**—The maximum length of time, in minutes, that users are allowed to stay connected to the VPN without logging out and reconnecting, from 1- 4473924 or blank. The default is unlimited (blank), but the idle timeout still applies.
- **Connection Time Alert Interval**—If you specify a maximum connection time, the alert interval defines the amount of time before the maximum time is reached to display a warning to the user about the upcoming automatic disconnect. The user can choose to end the connection and reconnect to restart the timer. The default is 1 minute. You can specify 1 to 30 minutes.
- **Idle Time**—The length of time, in minutes, that the VPN connection can be idle before it is automatically closed, from 1-35791394. If there is no communication activity on the connection for this consecutive number of minutes, the system stops the connection. The default is 30 minutes.
- **Idle Time Alert Interval**—The amount of time before the idle time is reached to display a warning to the user about the upcoming automatic disconnect due to an idle session. Any activity resets the timer. The default is 1 minute. You can specify 1 to 30 minutes.
- **Simultaneous Login Per User**—The maximum number of simultaneous connections allowed for a user. The default is 3. You can specify 1 to 2147483647 connections. Allowing a large number of simultaneous connections might compromise security and affect performance.

Address Assignment Attributes

The address assignment attributes of a group policy define the IP address pool for the group. The pool defined here overrides the pool defined in any connection profile that uses this group. Leave these settings blank if you want to use the pool defined in the connection profile.

- **IPv4 Address Pool**, **IPv6 Address Pool**—These options define the address pools for the remote endpoints. Clients are assigned an address from these pools based on the IP version they use to make the VPN

connection. Select a network object that defines a subnet for each IP type you want to support. Leave the list empty if you do not want to support that IP version. For example, you could define an IPv4 pool as 10.100.10.0/24. The address pool cannot be on the same subnet as the IP address for the outside interface.

You can specify a list of up to six address pools to use for local address allocation. The order in which you specify the pools is significant. The system allocates addresses from these pools in the order in which the pools appear.

- **DHCP Scope**—If you configure DHCP servers for the address pool in the connection profile, the DHCP scope identifies the subnets to use for the pool for this group. The DHCP server must also have addresses in the same subnet identified by the scope. The scope allows you to select a subset of the address pools defined in the DHCP server to use for this specific group.

  If you do not define a network scope, the DHCP server assigns IP addresses in the order of the address pools configured. It goes through the pools until it identifies an unassigned address.

  To specify a scope, select a network object that contains a routeable address on the same subnet as the desired pool, but not within the pool. The DHCP server determines which subnet this IP address belongs to and assigns an IP address from that pool.

  We recommend using the IP address of an interface whenever possible for routing purposes. For example, if the pool is 10.100.10.2-10.100.10.254, and the interface address is 10.100.10.1/24, use 10.100.10.1 as the DHCP scope. Do not use the network number. Click **Create New Network** if the object does not yet exist. You can use DHCP for IPv4 addressing only. If the address you choose is not an interface address, you might need to create a static route for the scope address.

Split Tunneling Attributes

The split tunneling attributes of a group policy define how the system should handle traffic meant for the internal network vs. externally-directed traffic. Split tunneling directs some network traffic through the VPN tunnel (encrypted) and the remaining network traffic outside the VPN tunnel (unencrypted or in clear text).

- **IPv4 Split Tunneling**, **IPv6 Split Tunneling**—You can specify different options based on whether the traffic uses IPv4 or IPv6 addressing, but the options for each are the same. If you want to enable split tunneling, specify one of the options that requires you to select network objects.
  - **Allow all traffic over tunnel**—Do no split tunneling. Once the user makes an RA VPN connection, all the user's traffic goes through the protected tunnel. This is the default. It is also considered the most secure option.
  - **Allow specified traffic over the tunnel**—Select the network objects that define destination network and host addresses. Any traffic to these destinations goes through the protected tunnel. Traffic to any other destination is routed by the client to connections outside the tunnel (such as a local Wi-Fi or network connection).
  - **Exclude networks specified below**—Select the network objects that define destination network or host addresses. Any traffic to these destinations is routed by the client to connections outside the tunnel. Traffic to any other destination goes through the tunnel.
- **Split DNS**—You can configure the system to send some DNS requests through the secure connection, while allowing the client to send other DNS requests to the DNS servers configured on the client. You can configure the following DNS behavior:

- **Send DNS Request as per split tunnel policy**—With this option, DNS requests are handled the same way as the split tunnel options are defined. If you enable split tunneling, DNS requests are sent based on the destination addresses. If you do not enable split tunneling, all DNS requests go over the protected connection.
- **Always send DNS requests over tunnel**—Select this option if you enable split tunneling, but you want all DNS requests sent through the protected connection to the DNS servers defined for the group.
- **Send only specified domains over tunnel**—Select this option if you want your protected DNS servers to resolve addresses for certain domains only. Then, specify those domains, separating domain names with commas. For example, example.com, example1.com. Use this option if you want your internal DNS servers to resolve names for internal domains, while external DNS servers handle all other Internet traffic.

AnyConnect Client Attributes

The AnyConnect Client attributes of a group policy define some SSL and connection settings used by the AnyConnect Client for a remote access VPN connection.

SSL Settings

- **Enable Datagram Transport Layer Security (DTLS)**—Whether to allow the AnyConnect Client to use two simultaneous tunnels: an SSL tunnel and a DTLS tunnel. Using DTLS avoids latency and bandwidth problems associated with some SSL connections and improves the performance of real-time applications that are sensitive to packet delays. If you do not enable DTLS, AnyConnect Client users establishing SSL VPN connections connect with an SSL tunnel only.
- **DTLS Compression**—Whether to compress Datagram Transport Layer Security (DTLS) connections for this group using LZS. DTLS Compression is disabled by default.
- **SSL Compression**—Whether to enable data compression, and if so, the method of data compression to use, **Deflate**, or **LZS**. SSL Compression is **Disabled** by default. Data compression speeds up transmission rates, but also increases the memory requirement and CPU usage for each user session. Therefore, SSL compression decreases the overall throughput of the device.
- **SSL Rekey Method**, **SSL Rekey Interval**—The client can rekey the VPN connection, renegotiating the crypto keys and initialization vectors, to increase the security of the connection. Disable rekeying by selecting **None**. To enable rekey, select **New Tunnel** to create a new tunnel each time. (The **Existing Tunnel** option results in the same action as **New Tunnel**.) If you enable rekeying, also set the rekey interval, which is 4 minutes by default. You can set the interval to 4-10080 minutes (1 week).

Connection Settings

- **Ignore the DF (Don't Fragment) bit**—Whether to ignore the Don't Fragment (DF) bit in packets that need fragmentation. Select this option to allow the forced fragmentation of packets that have the DF bit set, so that these packets can pass through the tunnel.
- **Client Bypass Protocol**—Allows you to configure how the secure gateway manages IPv4 traffic (when it is expecting only IPv6 traffic), or how it manages IPv6 traffic (when it is expecting only IPv4 traffic).

  When the AnyConnect Client makes a VPN connection to the headend, the headend assigns it an IPv4, IPv6, or both an IPv4 and IPv6 address. If the headend assigns the AnyConnect Client connection only

an IPv4 address or only an IPv6 address, you can configure the Client Bypass Protocol to drop network traffic for which the headend did not assign an IP address (default, disabled, not checked), or allow that traffic to bypass the headend and be sent from the client unencrypted or "in the clear" (enabled, checked).

For example, assume that the secure gateway assigns only an IPv4 address to the AnyConnect Client connection and the endpoint is dual-stacked. When the endpoint attempts to reach an IPv6 address, if Client Bypass Protocol is disabled, the IPv6 traffic is dropped; however, if Client Bypass Protocol is enabled, the IPv6 traffic is sent from the client in the clear.

- **MTU**—The maximum transmission unit (MTU) size for SSL VPN connections established by the AnyConnect Client. The default is 1406 bytes. The range is 576 to 1462 bytes.
- **Keepalive Messages Between AnyConnect and VPN Gateway**—Whether to exchange keepalive messages between peers to demonstrate that they are available to send and receive data in the tunnel. Keepalive messages transmit at set intervals. The default interval is 20 seconds, the valid range is 15 to 600 seconds.
- **DPD on Gateway Side Interval**, **DPD on Client Side Interval**—Enable Dead Peer Detection (DPD) to ensure that the VPN gateway or VPN client quickly detects when the peer is no longer responding. You can separately enable gateway or client DPD. The default interval is 30 seconds for sending DPD messages. The interval can be 5-3600 seconds.

Traffic Filters Attributes

The traffic filter attributes of a group policy define restrictions you want to place on users assigned to the group. You can use these attributes instead of creating access control policy rules to restrict RA VPN users to specific resources, based on host or subnet address and protocol, or on VLAN.

By default, RA VPN users are not restricted by the group policy from accessing any destination on your protected network.

- **Access List Filter**—Restrict access using an extended access control list (ACL). Select the Smart CLI Extended ACL object, or click **Create Extended Access List** and create it now.

  The extended ACL lets you filter based on source address, destination address, and protocol (such as IP or TCP). ACLs are evaluated on a top-down, first-match basis, so ensure that you place specific rules before more general rules. There is an implicit "deny any" at the end of the ACL, so if your intention is to simply deny access to a few subnets while allowing all other access, ensure that you include a "permit any" rule at the end of the ACL. The VPN filter applies to initial connections only. It does not apply to secondary connections, such as a SIP media connection, that are opened due to the action of application inspection.

  Because you cannot create network objects while editing an extended ACL Smart CLI object, you should create the ACL before editing the group policy. Otherwise, you might need to simply create the object, then go back later to create the network objects and then all the access control entries that you need. To create the ACL, go to **Device** > **Advanced Configuration** > **Smart CLI** > **Objects**, create an object, and select **Extended Access List** as the object type. For an example, see How to Control RA VPN Access By Group, on page 733.
- **Restrict VPN to VLAN**—Also called "VLAN mapping," this attribute specifies the egress VLAN interface for sessions to which this group policy applies. The system forwards all traffic from this group to the selected VLAN.

Use this attribute to assign a VLAN to the group policy to simplify access control. Assigning a value to this attribute is an alternative to using an ACL to filter traffic on a session. Ensure that you specify a VLAN number that is defined on a subinterface on the device. Values range from 1 to 4094.

Windows Browser Proxy Attributes

The Windows browser proxy attributes of a group policy determine how, and whether, a proxy defined on the user's browser operates.

You can select one of the following values for **Browser Proxy During VPN Session**:

- **No change in endpoint settings**—Allow the user to configure (or not configure) a browser proxy for HTTP, and use the proxy if it is configured.
- **Disable browser proxy**—Do not use the proxy defined for the browser, if any. No browser connections will go through the proxy.
- **Auto detect settings**—Enable the use of automatic proxy server detection in the browser for the client device.
- **Use custom settings**—Define a proxy that should be used by all client devices for HTTP traffic. Configure the following settings:
  - **Proxy Server IP or Hostname**, **Port**—The IP address, or hostname, of the proxy server, and the port used for proxy connections by the proxy server. The host and port combined cannot exceed 100 characters.
  - **Browser Exemption List**—Connections to the hosts/ports in the exemption list do not go through the proxy. Add all of the host/port values for destinations that should not use the proxy. For example, www.example.com port 80. Click the **Add** link to add items to the list. Click the trash can icon to delete items. The entire proxy exception list, combining all addresses and ports, cannot be longer than 255 characters.

Monitoring Remote Access VPN

To monitor and troubleshoot remote access VPN connections, open the CLI console or log into the device CLI and use the following commands.

- **show vpn-sessiondb** displays information about VPN sessions. You can reset these statistics using the **clear vpn-sessiondb** command.
- **show webvpn** *keyword* displays information about the remote access VPN configuration, including statistics and the AnyConnect images installed. Enter **show webvpn ?** to see the available keywords.
- **show aaa-server** displays statistics about the directory server used with remote access VPN.

Troubleshooting Remote Access VPNs

Remote access VPN connection issues can originate in the client or in the Firewall Threat Defense device configuration. The following topics cover the main troubleshooting problems you might encounter.

Troubleshooting SSL Connection Problems

If the user cannot make the initial, non-AnyConnect Client, SSL connection to the outside IP address to download the AnyConnect Client, do the following:

1. If you configured a non-default port for the remote access VPN connection profile, ensure the user is including the port number in the URL. For example: https://ravpn.example.com:4443
2. From the client workstation, verify that you can ping the IP address of the outside interface. If you cannot, determine why there is no route from the user's workstation to the address.
3. From the client workstation, verify that you can ping the fully-qualified domain name (FQDN) of the outside interface, the one defined in the remote access (RA) VPN connection profile. If you can ping the IP address but not the FQDN, then you need to update the DNS servers used by the client and RA VPN connection profile to add the FQDN-to-IP-address mapping.
4. Verify that the user is accepting the certificate presented by the outside interface. The user should accept it permanently.
5. Examine the RA VPN connection configuration and verify that you selected the correct outside interface. A common mistake is to select an inside interface, the one facing the internal networks, rather than the outside interface, which faces the RA VPN users.
6. If SSL encryption is properly configured, use an external sniffer to verify whether the TCP three-way handshake is successful.

Troubleshooting AnyConnect Client Download and Installation Problems

If the user can make an SSL connection to the outside interface, but cannot download and install the AnyConnect Client package, consider the following:

- Ensure that you uploaded an AnyConnect Client package for the client's operating system. For example, if the user's workstation runs Linux, but you did not upload a Linux AnyConnect Client image, there is no package that can be installed.
- For Windows clients, the user must have Administrator rights to install software.
- For Windows clients, the workstation must enable ActiveX or install Java JRE 1.5 or higher, with JRE 7 recommended.
- For Safari browsers, Java must be enabled.
- Try different browsers, one might fail where another succeeds.

Troubleshooting AnyConnect Client Connection Problems

If the user was able to connect to the outside interface, download, and install the AnyConnect Client, but could not then complete a connection using AnyConnect Client, consider the following:

- If you are using DHCP to provide IP addresses to the client, and the client cannot obtain an address, check the NAT rules. Any NAT rule that applies to the RA VPN network should include the route lookup option. Route lookup can help ensure the DHCP requests are sent to the DHCP server through an appropriate interface.

- If authentication fails, verify that the user is entering the correct username and password, and that the username is defined correctly in the authentication server. The authentication server must also be available through one of the data interfaces.

Note If the authentication server is on an external network, you need to configure a site-to-site VPN connection to the external network, and include the remote access VPN interface address within the VPN. For details, see How to Use a Directory Server on an Outside Network with Remote Access VPN, on page 720.

- If you configured a fully-qualified domain name (FQDN) for the outside interface in the remote access (RA) VPN connection profile, verify that you can ping the FQDN from the client device. If you can ping the IP address but not the FQDN, then you need to update the DNS servers used by the client and RA VPN connection profile to add the FQDN-to-IP-address mapping. If you are using the default AnyConnect Client profile that is generated when you specify an FQDN for the outside interface, the user will need to edit the server address to use the IP address until DNS is updated.
- Verify that the user is accepting the certificate presented by the outside interface. The user should accept it permanently.
- If the user's AnyConnect Client includes multiple connection profiles, that they are selecting the right one.
- If everything seems right on the client end, make an SSH connection to the Firewall Threat Defense device, and enter the **debug webvpn** command. Examine the messages issued during a connection attempt.

Troubleshooting RA VPN Traffic Flow Problems

If the user can make a secure remote access (RA) VPN connection, but cannot send and receive traffic, do the following:

1. Have the client disconnect, then reconnect. Sometimes this eliminates the problem.
2. In the AnyConnect Client, check the traffic statistics to determine whether both the sent and received counters are increasing. If the received packet count stays at zero, the Firewall Threat Defense device is not returning any traffic. There is likely a problem in the Firewall Threat Defense configuration. Common problems include the following:
   - Access rules are blocking traffic. Check the access control policy for rules that prevent traffic between the inside networks and the RA VPN address pool. You might need to create an explicit Allow rule if your default action is to block traffic.
   - The VPN filter is blocking traffic. Check the ACL traffic filter or VLAN filter configured in the group policy for the connection profile. You might need to make adjustments in the ACL or change the VLAN, depending on how (or if) you are filtering traffic based on group policy.
   - NAT rules are not being bypassed for the RA VPN traffic. Ensure that NAT exempt is configured for the RA VPN connection for every inside interface. Alternatively, ensure that the NAT rules do not prevent communication between the inside networks and interfaces and the RA VPN address pool and outside interface.

- Routes are misconfigured. Ensure that all defined routes are valid and functioning correctly. For example, if you have a static IP address defined for the outside interface, ensure that the routing table includes a default route (for 0.0.0.0/0 and ::/0).
- Ensure that the DNS server and domain name configured for the RA VPN are correct, and that the client system is using the correct ones. Verify that the DNS servers are reachable.
- If you enable split tunneling in the RA VPN, check whether traffic to the specified inside networks is going through the tunnel, while all other traffic is bypassing the tunnel (so that the Firewall Threat Defense device does not see it).

3. Make an SSH connection to the Firewall Threat Defense device and verify that traffic is being sent and received for the remote access VPN. Use the following commands.
   - **show webvpn anyconnect**
   - **show vpn-sessiondb**

Examples for Remote Access VPN

The following are examples of configuring remote access VPN.

How to Implement RADIUS Change of Authorization

RADIUS Change of Authorization (CoA), also known as dynamic authorization, provides end-point security for the Firewall Threat Defense remote access VPN. A key challenge for RA VPNs is to secure the internal network against compromised end points and to secure the end point itself when it is affected by viruses or malware, by remediating the attack on the endpoint. There is a need to secure the endpoint and the internal network in all phases, that is, before, during, and after the RA VPN session. The RADIUS CoA feature helps in achieving this goal.

If you use Cisco Identity Services Engine (ISE) RADIUS servers, you can configure Change of Authorization policy enforcement.

The ISE Change of Authorization feature provides a mechanism to change the attributes of an authentication, authorization, and accounting (AAA) session after it is established. When a policy changes for a user or user group in AAA, ISE sends CoA messages to the Firewall Threat Defense device to reinitialize authentication and apply the new policy. An Inline Posture Enforcement Point (IPEP) is not required to apply access control lists (ACLs) for each VPN session established with the Firewall Threat Defense device.

The attributes you can change during CoA are redirect URL, redirect ACL, and security group tag.

The following topics explain how CoA works, and how to configure it.

System Flow for Change of Authorization

Cisco ISE has a client posture agent that assesses an endpoint's compliance for criteria such as processes, files, registry entries, antivirus protection, antispyware protection, and firewall software installed on the host. Administrators can then restrict network access until the endpoint is in compliance or can elevate local user privileges so they can establish remediation practices. ISE Posture performs a client-side evaluation. The client receives the posture requirement policy from ISE, performs the posture data collection, compares the results against the policy, and sends the assessment results back to ISE.

Following is the system flow between the Firewall Threat Defense device, ISE, and the RA VPN client for Change of Authorization (CoA) processing.

1. The remote user starts an RA VPN session, using the AnyConnect Client, with the Firepower Threat Defense device.
2. The Firewall Threat Defense device sends a RADIUS Access-Request message for that user to the ISE server.
3. Because the client posture is unknown at this point, ISE matches the user to the authorization policy that is configured for unknown posture. This policy defines the following cisco-av-pair options, which ISE sends to the Firewall Threat Defense in a RADIUS Access-Accept response.
   - **url-redirect-acl=***acl_name*, where *acl_name* is the name of an extended ACL that is configured on the Firewall Threat Defense device. This ACL defines which user traffic should be redirected to the ISE server, which is HTTP traffic. For example:

     ```
     url-redirect-acl=redirect
     ```

   - **url-redirect=***url*, where the URL is the one to which traffic should be redirected. For example:

     ```
     url-redirect=https://ise2.example.com:8443/guestportal/gateway?sessionId=xx&action=cpp
     ```

   You must configure DNS for data interfaces so that the hostname can be resolved. If you also configure traffic filtering in the group policy for the connection profile, ensure that the client pool can reach the ISE server through the port (TCP/8443 in the example).
4. The Firewall Threat Defense device sends a RADIUS Accounting-Request start packet and receives a response from ISE. The accounting request includes all of the details of the session, including the session ID, the external IP address of the VPN client, and the IP address of the Firewall Threat Defense device. ISE uses the session ID to identify that session. The Firewall Threat Defense device also sends periodic interim account information, where the most important attribute is the Framed-IP-Address with the IP address that is assigned to the client by the Firewall Threat Defense device.
5. While in an unknown posture state, the Firewall Threat Defense device redirects traffic from the client that matches the redirect ACL to the redirect URL. ISE determines if the client has the required posture compliance module, and prompts the user to install it if necessary.
6. After the agent is installed on the client device, it automatically performs the checks that are configured in the ISE posture policy. The client communicates directly with ISE. It sends a posture report to ISE, which can include multiple exchanges using the SWISS protocol and ports TCP/UDP 8905.
7. When ISE receives the posture report from the agent, it processes the authorization rules once again. This time, the posture result is known and a different rule now matches the client. ISE sends a RADIUS CoA packet, which includes the downloadable ACL (DACL) for either compliant or non-compliant endpoints. For example, the compliant DACL might permit all access, while the non-compliant DACL denies all access. The contents of the DACL are up to the ISE administrator.
8. The Firewall Threat Defense device removes the redirection. If it does not have the DACLs cached, it must send an Access-Request in order to download them from ISE. The specific DACL is attached to the VPN session; it does not become part of the device configuration.
9. The next time that the RA VPN user tries to access the web page, the user can access the resources that are permitted by the DACL that is installed on the Firewall Threat Defense device for the session.

Note If the endpoint fails to satisfy any mandatory requirement and if a manual remediation is required, then a remediation window opens in the AnyConnect Client, displaying the items that require action. The remediation window runs in the background so that the updates on network activity do not pop up and interfere or cause disruption. A user can click **Details** in the ISE Posture tile portion of the AnyConnect Client to see what has been detected and what updates are needed before the user can join the network.

Configure Change of Authorization on the Firewall Threat Defense Device

Most of the Change of Authorization policy is configured in the ISE server. However, you must configure the Firewall Threat Defense device to connect to ISE correctly. The following procedure explains how to configure the Firewall Threat Defense side of the configuration.

Before you begin

If you use hostnames in any object, ensure that you configure DNS servers for use with the data interfaces, as explained in Configuring DNS for Data and Management Traffic, on page 763. You typically need to configure DNS anyway to have a fully-functional system.

Procedure

Step 1 Configure the extended access control list (ACL) for redirecting initial connections to ISE.

The purpose of the redirect ACL is to send initial traffic to ISE so that ISE can assess the client posture. The ACL should send HTTPS traffic to ISE, but not traffic that is already destined for ISE, or traffic that is directed to a DNS server for name resolution. A sample redirect ACL might look like the following:

```
access-list redirect extended deny ip any host <ISE server IP>
access-list redirect extended deny ip any host <DNS server IP>
access-list redirect extended deny icmp any any
access-list redirect extended permit tcp any any eq www
```

However, note that ACLs have an implicit “deny any any” as the last access control entry (ACE). In this example, the last ACE, which matches TCP port www (that is, port 80), will not match any traffic that matches the first 3 ACEs, so those are redundant. You could simply create an ACL with the last ACE and get the same results.

Note that in a redirect ACL, the permit and deny actions simply determine which traffic matches the ACL, with permit matching and deny not matching. No traffic is actually dropped, denied traffic is simply not redirected to ISE.

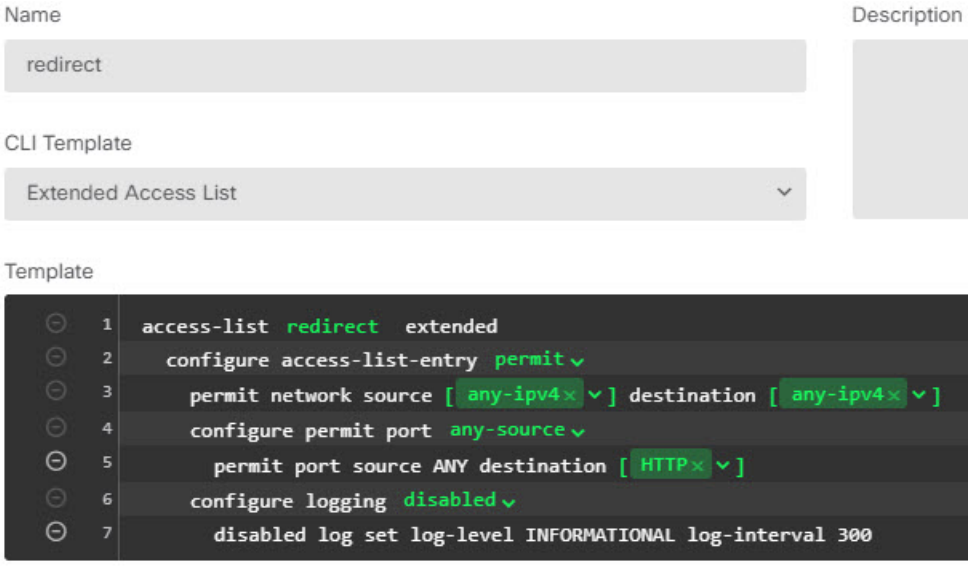
To create the redirect ACL, you need to configure a Smart CLI object.

a) Choose **Device** > **Advanced Configuration** > **Smart CLI** > **Objects**.
b) Click **+** to create a new object.
c) Enter a name for the ACL. For example, **redirect**.
d) For **CLI Template**, select **Extended Access List**.
e) Configure the following in the **Template** body:

- configure access-list-entry action = permit

- source-network = any-ipv4
- destination-network = any-ipv4
- configure permit port = any-source
- destination-port = HTTP
- configure logging = disabled

The ACE should look like the following:



f) Click **OK**.

This ACL will be configured the next time you deploy changes. You do not need to use the object in any other policy to force deployment.

Note
This ACL applies to IPv4 only. If you also want to support IPv6, simply add a second ACE with all the same attributes, except select any-ipv6 for the source and destination networks. You can also add the other ACEs to ensure traffic to the ISE or DNS server is not redirected. You will first need to create host network objects to hold the IP addresses of those servers.

Step 2 Configure a RADIUS server group for dynamic authorization.

There are several critical options that you must select correctly in the RADIUS server and server group objects to enable Change of Authorization, also known as dynamic authorization. The following procedure focuses on these attributes. For more details on these objects, see RADIUS Servers and Groups, on page 165.

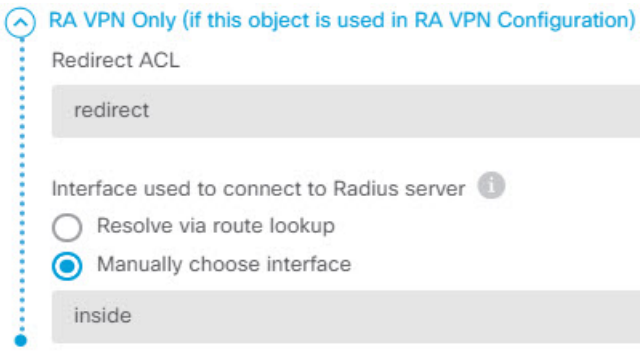
a) Choose **Objects** > **Identity Sources**.
b) Click **+** > **RADIUS Server**.
c) Enter a name for the server, and the hostname/IP address of the ISE RADIUS server, authentication port, and secret key configured on the server. Adjust the timeout if desired. These options are not directly related to dynamic authorization.
d) Click the RA VPN Only link and configure the following options:

- **Redirect ACL**—Select the extended ACL you created for redirection. In this example, the ACL named redirect.
- **Interface used to connect to Radius server**—Select **Manually Choose Interface**, and select the interface through which the server can be reached. You must select a specific interface so that the system can correctly enable the CoA listener on the interface.

  If the server is on the same network as the management address, which means you will select the diagnostic interface, you must also configure an IP address on the diagnostic interface. Having a management IP address is not sufficient. Go to **Device** > **Interfaces**, and configure an IP address on the diagnostic interface that is on the same subnet as the management IP address.

  If you also use this server for the Firewall Device Manager administrative access, this interface is ignored. Administrative access attempts are always authenticated through the management IP address.

The following example shows the options configured for the inside interface.



e) Click **OK** to save the server object.

If you have a redundant setup, with multiple duplicate ISE RADIUS servers, create server objects for each of these servers.

f) Click + > **RADIUS Server Group**.
g) Enter a name for the server group, and adjust the dead time and maximum attempts if desired.
h) Select the **Dynamic Authorization** option, and change the port number if your ISE server is configured to use a different port. Port 1700 is the default port used for listening for CoA packets.
i) If the RADIUS server is configured to use an AD server for authenticating users, select the **Realm that Supports the RADIUS Server** that specifies the AD server used in conjunction with this RADIUS server. If the realm does not already exist, click **Create New Identity Realm** at the bottom of the list and configure it now.
j) Under **RADIUS Server**, click + and select the server object you created for RA VPN.
k) Click **OK** to save the server group object.

Step 3 Choose **Device** > **RA VPN** > **Connection Profiles**, and create a connection profile that uses this RADIUS server group.

Use **AAA Authentication** (either only or with certificates), and select the server group in the **Primary Identity Source for User Authentication**, **Authorization**, and **Accounting** options.

Configure all other options as needed for your organization.

Note

If the DNS servers are reached through the VPN network, edit the group policy used in the connection profile to configure the **Split DNS** option on the Split Tunneling Attributes page.

Configure Change of Authorization in ISE

Most of the Change of Authorization configuration is done in the ISE server. ISE has a posture assessment agent that runs on the endpoint device, and ISE communicates directly with the device to determine posture stance. The Firewall Threat Defense device essentially waits for instructions from ISE on how to handle a given end user.

A full discussion of configuring posture assessment policies is outside the scope of this document. However, the following procedure explains some of the basics. Use this as a starting point for configuring ISE. Note that the exact command paths, page names, and attribute names can change from release to release. The version of ISE you are using might use different terminology or organization.

The minimum supported ISE release is 2.2 patch 1.

Before you begin

This procedure assumes you have already configured users in the ISE RADIUS server.

Procedure

Step 1 Choose **Administration** > **Network Resources** > **Network Devices** > **Network Devices**, add the Firewall Threat Defense device to the ISE Network Device inventory, and configure the RADIUS settings.

Select the **RADIUS Authentication Settings**, and configure the same **Shared Secret** that is configured in the Firewall Threat Defense RADIUS server object. If desired, change the **CoA Port** number and ensure you configure the same port in the Firewall Threat Defense RADIUS server group object.

Step 2 Choose **Policy** > **Policy Elements** > **Results** > **Authorization** > **Downloadable ACLs**.

Create 2 downloadable ACLs (DACL), one for use by compliant endpoints, one for non-compliant endpoints.

For example, you might allow all access for compliant endpoints (permit ip any any), while denying all access to non-compliant endpoints (deny ip any any). You can make these DACLs as complex as you require, to provide the exact access users should have based on their compliance state. You will use these DACLs in authorization profiles.

Step 3 Choose **Policy** > **Policy Elements** > **Results** > **Authorization** > **Authorization Profile** and configure the required profiles.

You need profiles for the following states. Minimum attributes for each are listed.

- **Unknown**—The unknown posture profile is the default posture profile. Every endpoint is matched to this policy when they initially establish the RA VPN connection. The point of this rule is to apply the redirect ACL and URL, and to download the posture agent if it is not already on the endpoint. Endpoints can remain attached to this profile if the agent is not installed, or if installation fails. Otherwise, after assessing the posture, endpoints move to the compliant or non-compliant profiles.

  Minimum attributes include the following:

- **Name**—For example, PRE_POSTURE.
- **Access Type**—Select **ACCESS_ACCEPT**.
- **Common Tasks**—Select **Web Redirection (CWA, MDM, NSP, CPP)**, then select **Client Provisioning (Posture)**, and enter the name of the redirect ACL you configured on the Firewall Threat Defense device. In **Value**, select **Client Provisioning Portal** if it is not already selected.
- The **Attribute Details** should show two cisco-av-pair values, for url-redirect-acl and url-redirect. ISE will send this data to the Firewall Threat Defense device, which will apply the criteria to the RA VPN user session.

- **Compliant**—After the posture assessment completes, if the endpoint meets all requirements configured for the endpoint, the client is considered compliant and gets this profile. You would typically give this client full access.

  Minimum attributes include the following:

  - **Name**—For example, FULL_ACCESS.
  - **Access Type**—Select **ACCESS_ACCEPT**.
  - **Common Tasks**—Select **DACL Name**, and select the downloadable ACL for compliant users, for example, PERMIT_ALL_TRAFFIC. ISE will send the ACL to the Firewall Threat Defense device, which will apply it to the user session. This DACL will replace the initial redirect ACL for the user session.

- **Non-compliant**—If the posture assessment determines that the endpoint does not meet all requirements, there is a countdown during which the client can bring the endpoint into compliance, for example, by installing required updates. The AnyConnect Client informs the user of the compliance issues. During the countdown, the endpoint remains in the unknown compliance state. If the endpoint remains non-compliant after the countdown expires, the session is marked non-compliant and it gets the non-compliant profile. You would typically prevent all access for this endpoint, or at least restrict access in some way.

  Minimum attributes include the following:

  - **Name**—For example, Non_Compliant.
  - **Access Type**—Select **ACCESS_ACCEPT**.
  - **Common Tasks**—Select **DACL Name**, and select the downloadable ACL for non-compliant users, for example, DENY_ALL_TRAFFIC. ISE will send the ACL to the Firewall Threat Defense device, which will apply it to the user session. This DACL will replace the initial redirect ACL for the user session.

Step 4 Choose **Policy** > **Policy Elements** > **Results** > **Client Provisioning** > **Resources** and configure the following resources:

- **AnyConnect package**—The head end package file, which you download from software.cisco.com. You need separate packages for the client platforms you support, so you might need to configure multiple types, such as AnyConnectDesktopWindows.
- **ISE Posture Configuration File (Type: AnyConnectProfile)**—This configuration file defines the settings that the compliance module uses to evaluate the end user's device. This file also defines the length of time the user has to bring a non-compliant device into compliance.

- **Compliance Module Package (Type: ComplianceModule)**—The AnyConnect Client Compliance Module file is the file which will be pushed down to the installed AnyConnect package to check endpoint compliance. Download this file using the **Add Resource from Cisco Site** command. Ensure that you download the correct module based on the AnyConnect Client packages you have configured, or users will get download failures. You can also find these files on software.cisco.com in the AnyConnect Client listings in the ISEComplianceModule folder.
- **AnyConnect Configuration File (Type: AnyConnectConfig)**—These AnyConnect Client release-specific settings define the **AnyConnect Package**, **Compliance Module**, and **ISE Posture** to apply. Because the packages are OS-specific, create separate configuration files for each client OS you will support (for example, Windows, MAC, Linux).

Step 5 Choose **Policy** > **Client Provisioning** and configure the client provisioning policy.

Create new rules, for example, with names like CoA_ClientProvisionWin, for each operating system that should implement CoA. Select the appropriate operating system for the rule, and in **Results**, select the AnyConnect Client configuration file you created for the OS as the **Agent**.

Disable the default OS-specific rules that you are replacing.

Step 6 Configure the posture policy.

In this step, you develop the posture requirements that make sense for your organization.

- Choose **Policy** > **Policy Elements** > **Conditions** > **Posture**, and define the simple posture conditions that need to be met. For example, you might require that the user have certain applications installed.
- Choose **Policy** > **Policy Elements** > **Results** > **Posture** > **Requirements**, and define the compliance module requirement for the endpoint.
- Choose **Policy** > **Posture** > **Posture Policy** and configure the policies for the supported operating systems.

Step 7 Choose **Policy** > **Policy Sets** > **Default** > **Authorization Policy** and create the policy.

Add rules for each of the compliant conditions. These sample values are based on the examples in previous steps.

- Unknown, for pre-posture and posture download.
  - Name—For example, PRE_POSTURE
  - Conditions—”Session-PostureStatus EQUALS Unknown” AND “Radius-NAS-Port-Type EQUALS Virtual”
  - Profiles—For example, PRE_POSTURE
- Compliant, for clients that satisfy posture requirements.
  - Name—For example, FULL_ACCESS
  - Conditions—”Session-PostureStatus EQUALS Compliant” AND “Radius-NAS-Port-Type EQUALS Virtual”
  - Profiles—For example, FULL_ACCESS
- Non-compliance, for clients that fail posture requirements.
  - Name—For example, NON-COMPLIANT

- Conditions—"Session-PostureStatus EQUALS NonCompliant" AND "Radius-NAS-Port-Type EQUALS Virtual"
- Profiles—For example, Non_Compliant

Step 8 (Optional.) Choose **Administration** > **Settings** > **Posture** > **Reassessments** and enable posture reassessment.

By default, posture is assessed at connection time only. You can enable posture reassessment to periodically check the posture of connected endpoints. You can set the reassessment interval to determine how often this occurs.

If the system fails reassessment, you can define how the system should respond. You can allow the user to continue (remain connected), log the user off, or ask the user to remediate the system.

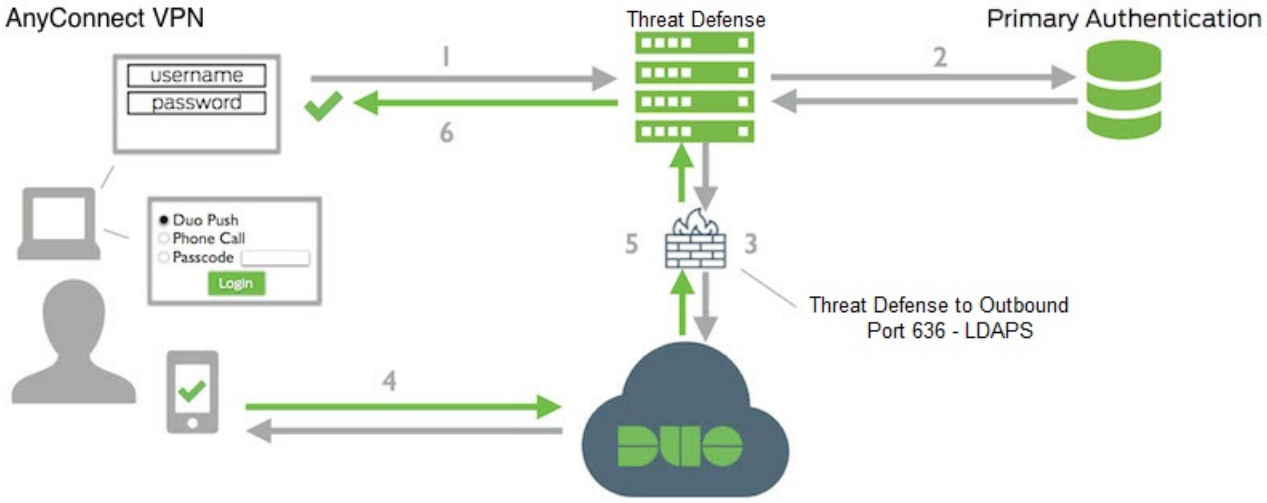
How to Configure Two-Factor Authentication using Duo LDAP

You can use the Duo LDAP server as the secondary authentication source in conjunction with a Microsoft Active Directory (AD) or RADIUS server as the primary source. With Duo LDAP, the secondary authentication validates the primary authentication with a Duo passcode, push notification, or phone call.

The following topics explain the configuration in more detail.

System Flow for Duo LDAP Secondary Authentication

The following graphic shows how the Firewall Threat Defense and Duo work together to provide two-factor authentication using LDAP.



Following is an explanation of the system flow:

1. The user makes a remote access VPN connection to the Firewall Threat Defense device and provides username and password.
2. Firewall Threat Defense authenticates this primary authentication attempt with the primary authentication server, which might be Active Directory or RADIUS.

3. If the primary authentication works, the Firewall Threat Defense sends a request for secondary authentication to the Duo LDAP server.
4. Duo then authenticates the user separately, through push notification, text message with a passcode, or a telephone call. The user must complete this authentication successfully.
5. Duo responds to the Firewall Threat Defense device to indicate whether the user authenticated successfully.
6. If the secondary authentication was successful, the Firewall Threat Defense device establishes a remote access VPN connection with the user's AnyConnect Client.

Configure Duo LDAP Secondary Authentication

The following procedure explains the end-to-end process of configuring two-factor authentication, using Duo LDAP as the secondary authentication source, for remote access VPN. Note that you must have an account with Duo, and obtain some information from Duo, to complete this configuration.

Procedure

Step 1 Create a Duo account and obtain the integration key, secret key, and API hostname.

Following is an overview of the process. For details, please see the Duo web site, https://duo.com.

a) Sign up for a Duo account.
b) Log in to the Duo Admin Panel and navigate to **Applications**.
c) Click **Protect an Application** and locate Cisco SSL VPN in the applications list. Click **Protect this Application** to get your integration key, secret key, and API hostname. For help, see the Duo *Getting Started* guide, https://duo.com/docs/getting-started.

Step 2 Create a Duo LDAP identity source for the Duo LDAP server.

You must use the Firewall Threat Defense API to create the Duo LDAP object; you cannot create it using the Firewall Device Manager. You can either use the API Explorer, or write your own client application, to create the object. The following procedure explains how to create the object using API Explorer.

a) Log into the Firewall Device Manager, click the more options button (⋮), and choose **API Explorer**.

The system opens the API Explorer in a separate tab or window, depending on your browser settings.

b) (Optional.) Obtain the values needed to identify the interface the system should use to connect to the Duo LDAP server.

If you do not specify an interface, the system uses the routing table. If necessary, you can create a static route for the Duo LDAP server. Alternatively, you can specify the interface to use in the Duo LDAP object. If you want to specify the interface, use the various GET methods in the Interfaces group to obtain the needed values. You can use physical, subinterface, EtherChannel, or VLAN interfaces. For example, to get the values for a physical interface, use the GET /devices/default/interfaces method and find the object for the interface you need to use. You need the following values from the interface object:

- id
- type
- version

- name

c) Click on the **DuoLDAPIdentitySource** heading to open the group.
d) Click on the **POST /object/duoldapidentitysources** method.
e) Under the **Parameters** heading, for the **body** element, click in the **Example Value** display box in the **Data Type** column on the right. This action loads the example into the body value edit box.
f) In the **body value** edit box, do the following:

- Delete the following attribute lines: **version**, **id**. (These attributes are needed for PUT calls but not for POST.)
- For **name**, enter a name for the object, such as Duo-LDAP-server.
- For **description**, either enter a meaningful description of the object for your reference purposes, or delete the attribute line.
- For **apiHostname**, enter the API Hostname that you obtained from your Duo account. The hostname should look like the following, with the X's replaced with your unique value: API-XXXXXXXX.DUOSECURITY.COM. Uppercase is not required.
- For **port**, enter the TCP port to use for LDAPS. This should be 636 unless you have been told by Duo to use a different port. Note that you must ensure that your access control list allows traffic to the Duo LDAP server through this port.
- For **timeout**, enter the timeout, in seconds, to connect to the Duo server. The value can be 1-300 seconds. The default is 120. To use the default, either enter 120 or delete the attribute line.
- For **integrationKey**, enter the integration key that you obtained from your Duo account.
- For **secretKey**, enter the secret key that you obtained from your Duo account. This key will subsequently be masked.
- For **interface**, either enter the id, type, version, and name values of the interface to use to connect to the Duo LDAP server, or delete the 6 lines used to define the interface attribute, including the trailing closing brace.
- For **type**, leave the value as duoldapidentitysource.

For example, the object body might look like the following, where the apiHostname and integrationKey are obfuscated, but the intentionally faked secret key is shown:

```
{
  "name": "Duo-LDAP-server",
  "description": "Duo LDAP server for RA VPN",
  "apiHostname": "API-XXXXXXXX.DUOSECURITY.COM",
  "port": 636,
  "timeout": 120,
  "integrationKey": "XXXXXXXXXXXXXXXXXXXX",
  "secretKey": "123456789",
  "type": "duoldapidentitysource"
}
```

g) Click the **Try It Out!** button.

The system will issue the **curl** command to post the object to the device configuration. You are shown the curl command, the response body, and the response code. If you created a valid body, you should see **200** in the **Response Code** field.

If you made an error, look at the response body for an error message. You can correct the body value and try again.

h) Return to the Firewall Device Manager by clicking **Device** in the top menu.

i) Click **Objects**, then click **Identity Sources** in the table of contents.

Your Duo LDAP object should appear in the list. If it does not, go back to the API Explorer and try to create the object again. You can use the GET method to check whether it was actually created.

Note that you can delete the object using the Firewall Device Manager, but you cannot edit it or see its contents. You must use the API for those actions. The relevant methods are shown in the **DuoLDAPIdentitySource** group.

Step 3 Upload the trusted CA certificate for the Duo web site to the Firewall Device Manager.

The Firewall Threat Defense system must have the certificate needed to validate the connection to the Duo LDAP server. You can obtain and upload the certificate using this procedure, which was done with the Google Chrome browser. The exact steps for your browser might differ. Alternatively, you can go directly to https://www.digicert.com/digicert-root-certificates.htm and download the certificate, but the following procedure is generic and you can use it to obtain root trusted CA certificates for any site.

a) Open https://duo.com in your browser.

b) Click the site information link in the browser's URL field, then click the **Certificate** link. This action opens the certificate information dialog box.

c) Click the **Certificate Path** tab, and select the root (top) level of the path. In this case, DigiCert.

d) With DigiCert selected, click **View Certificate**. This action will open a new Certificate dialog box, and the General tab should indicate that it was issued to DigiCert High Assurance EV Root CA. This is the root CA certificate that you need to upload to the Firewall Device Manager.

e) Click the **Details** tab, then click the **Copy to File** button to start the certificate download wizard.

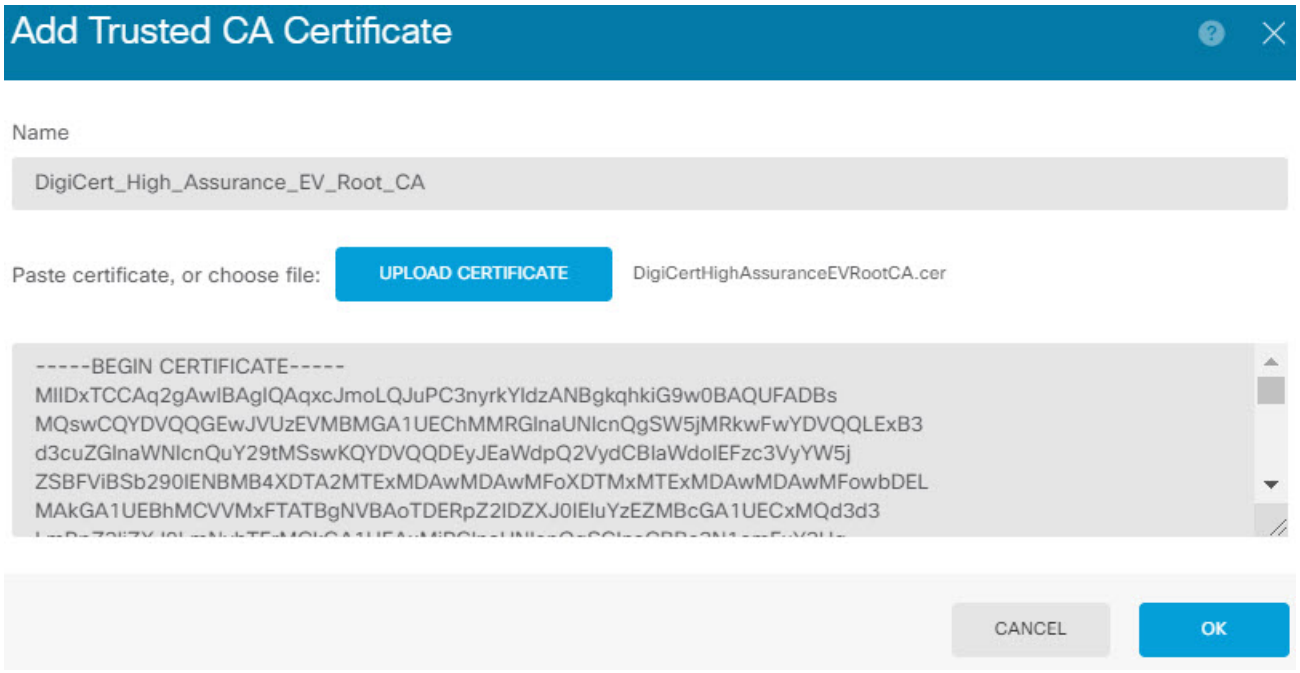
f) Use the wizard to download the certificate to your workstation. Download using the default DER format.

g) In the Firewall Device Manager, choose to **Objects** > **Certificates**.

h) Click **+** > **Add Trusted CA Certificate**.

i) Enter a name for the certificate, for example, DigiCert_High_Assurance_EV_Root_CA. (Spaces are not allowed.)

j) Click **Upload Certificate** and select the file you downloaded.



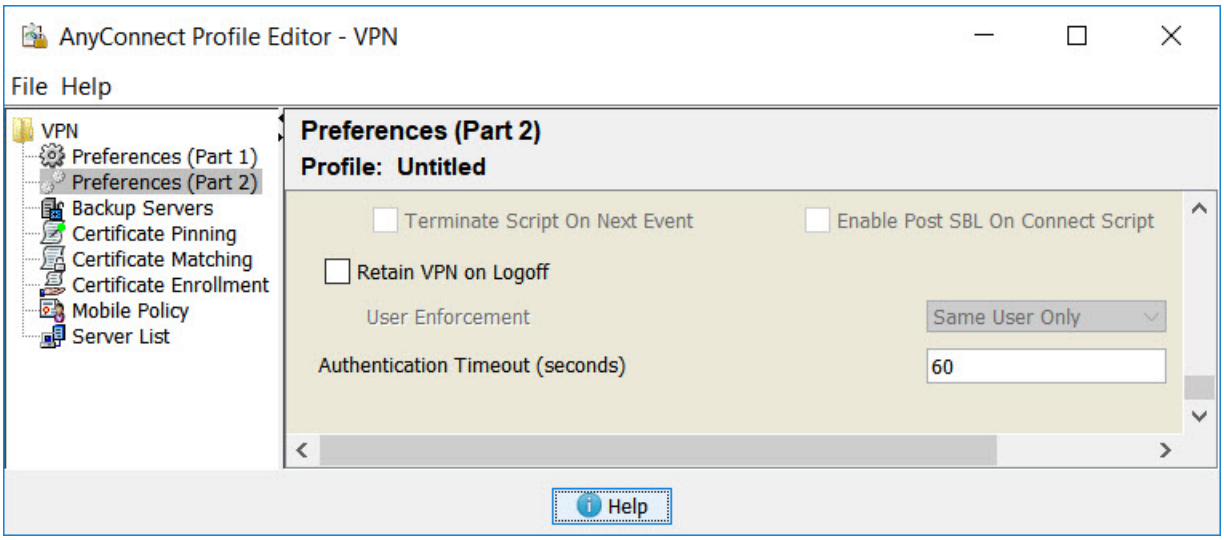
k) Click **OK**.

Step 4 Use the AnyConnect Client Profile Editor to create a profile that specifies 60 seconds or more for authentication timeout.

You need to give users extra time to obtain the Duo passcode and complete the secondary authentication. We recommend at least 60 seconds.

For details on creating AnyConnect Client profiles and uploading them, see Configure and Upload Client Profiles, on page 680. The following procedure explains how to configure the authentication timeout only, and then upload the profile to the Firewall Threat Defense. If you want to change other settings, you can do so now.

a) If you have not already done so, download and install the AnyConnect Client profile editor package. You can find this in the Cisco Software center (software.cisco.com) in the folder for your AnyConnect Client version.
b) Open the AnyConnect Client **VPN Profile Editor**.
c) Select **Preferences (Part 2)** in the table of contents, scroll to the end of the page, and change **Authentication Timeout** to 60 (or more). The following image is from the AnyConnect 4.7 VPN Profile Editor; previous or subsequent versions might be different.



d) Choose **File** > **Save**, and save the profile XML file to your workstation with an appropriate name, for example, duo-ldap-profile.xml.

You can now close the VPN Profile Editor application.

e) In the Firewall Device Manager, choose **Objects** > **AnyConnect Client Profiles**.
f) Click + to create a new profile object.
g) Enter a **Name** for the object. For example, Duo-LDAP-profile.
h) Click **Upload**, and select the XML file you created.
i) Click **OK**.

Step 5 Create a group policy and select the AnyConnect Client profile in the policy.

The group policy that you assign to a user controls many aspects of the connection. The following procedure explains how to assign the profile XML file to the group. For more details about what you can do with group policies, see Configure Group Policies for RA VPN, on page 692.

a) Click **View Configuration** in **Device** > **Remote Access VPN**.
b) Choose **Group Policies** in the table of contents.
c) Either edit DfltGrpPolicy, or click + and create a new group policy. For example, if you need a single remote access VPN connection profile for all users, editing the default group policy is appropriate.
d) On the General page, configure the following properties:

- **Name**—For a new profile, enter a name. For example, Duo-LDAP-group.
- **AnyConnect Client Profiles**—Click + and select the AnyConnect Client profile object you created.

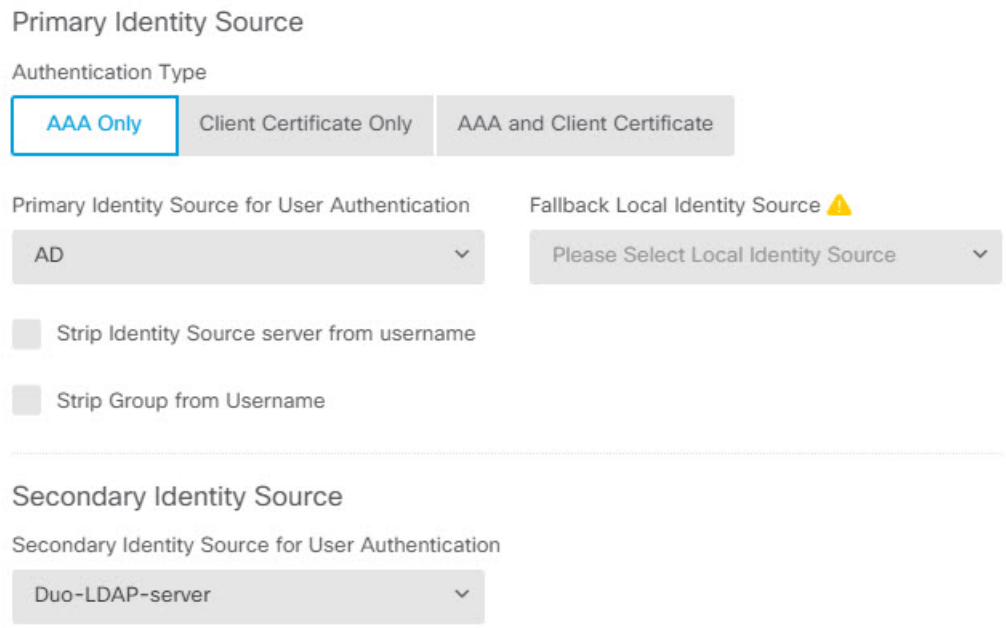
e) Click **OK** to save the group profile.

Step 6 Create or edit the remote access VPN connection profile to use for Duo-LDAP secondary authentication.

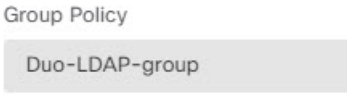
There are many steps to configuring a connection profile, which are explained in Configure an RA VPN Connection Profile, on page 685. The following procedure just mentions the key changes to make to enable Duo-LDAP as the secondary authentication source, and to apply the AnyConnect Client profile. For new connection profiles, you must configure the rest of the required fields. For this procedure, we assume you are editing an existing connection profile, and you simply need to change these two settings.

a) On the RA VPN page, choose **Connection Profiles** in the table of contents.
b) Either edit an existing connection profile, or create a new one.
c) Under Primary Identity Source, configure the following:

- **Authentication Type**—Choose either **AAA Only** or **AAA and Client Certificate**. You cannot configure two-factor authentication unless you use AAA.
- **Primary Identity Source for User Authentication**—Select your primary Active Directory or RADIUS server. Note that you can select a Duo-LDAP identity source as the primary source. However, Duo-LDAP provides authentication services only, not identity services, so if you use it as a primary authentication source, you will not see usernames associated with RA VPN connections in any dashboards, and you will not be able to write access control rules for these users. (You can configure fallback to the local identity source if you want to.)
- **Secondary Identity Source**—Select the Duo-LDAP identity source.



d) Click **Next**.
e) On the Remote User Experience page, select the **Group Policy** you created or edited.



f) Click **Next** on this page and the next page, Global Settings.
g) Click **Finish** to save your changes to the connection profile.

Step 7 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.



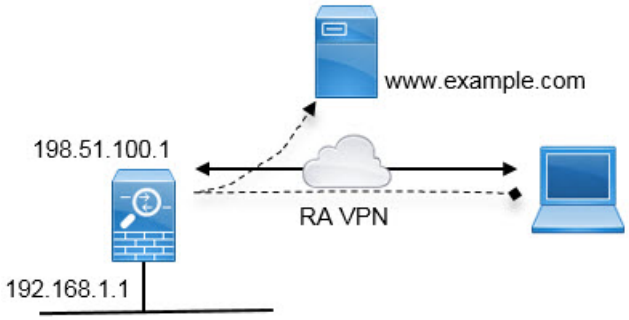
b) Click the **Deploy Now** button.

You can wait until deployment completes, or click **OK** and check the task list or deployment history later.

How to Provide Internet Access on the Outside Interface for Remote Access VPN Users (Hair Pinning)

In remote access VPN, you might want users on the remote networks to access the Internet through your device. However, because the remote users are entering your device on the same interface that faces the Internet (the outside interface), you need to bounce Internet traffic right back out of the outside interface. This technique is sometimes called hair pinning.

The following graphic shows an example. There is a remote access VPN configured on the outside interface, 198.51.100.1. You want to split the remote user's VPN tunnel, so that Internet-bound traffic goes back out the outside interface, while traffic to your internal networks continue through the device. Thus, when a remote user wants to go to a server on the Internet, such as www.example.com, the connection first goes through the VPN, then gets routed back out to the Internet from the 198.51.100.1 interface.



The following procedure explains how to configure this service.

Before you begin

This example assumes that you have already registered the device, applied a remote access VPN license, and uploaded the AnyConnect Client image. It also assumes that you have configured the identity realm, which is also used in Identity policies.

Procedure

Step 1 Configure the remote access VPN connection.

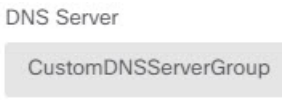
The configuration requires a customized group policy in addition to the connection profile. Because hair pinning is a common configuration, and the required settings in the group policy are generally applicable, in this example we will edit the default group policy instead of creating a new group policy. You can take either approach.

a) Click **View Configuration** in the **Device** > **Remote Access VPN** group.

b) Click **Group Policies** in the table of contents, then click the edit icon ( ) for the DfltGrpPolicy object.

c) Make the following changes to the default group policy:

- On the **General** page, in **DNS Server**, select the DNS server group that defines the servers VPN endpoints should use to resolve domain names.



- On the **Split Tunneling** page, for both **IPv4** and **IPv6 Split Tunneling**, select the **Allow all traffic over tunnel option**. This is the default setting, so it might already be configured correctly.



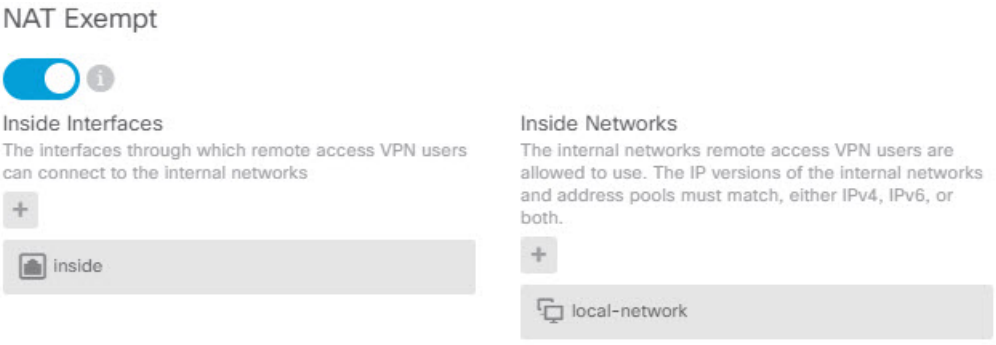
Note
This is a critical setting to enable hair-pinning. You want all traffic to go to the VPN gateway, whereas split tunneling is a way to allow remote clients to directly access local or Internet sites outside of the VPN.

d) Click **OK** to save the changes to the default group policy.
e) Click **Connection Profiles** and either edit an existing profile or create a new one.
f) In the connection profile, page through the wizard and configure all options as you would for any other RA VPN configuration. However, you must configure the following options correctly to enable hair-pinning:

- **Group Policy**, in step 2. Select the group policy you customized for hair-pinning.



- **NAT Exempt**, in step 3. Enable this feature. Select the inside interface, then select a network object that defines the internal networks. In this example, the object should specify 192.168.1.0/24. RA VPN traffic going to the internal network will not get address translation. However, because hair-pinned traffic is going out the outside interface, it will still be NAT'ed because the NAT exemption applies to the inside interface only. Note that if you have other connection profiles defined, you need to add to the existing settings, as the configuration applies to all connection profiles.



Note
The **NAT Exempt** option is the other critical setting for the hair pin configuration.

g) (Optional.) In the **Global Settings** step, select the **Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)** option.

By selecting this option, you remove the need to configure access control rules to allow traffic from RA VPN pool addresses. This option provides improved security (external users cannot spoof addresses in the pool), but it means that RA VPN traffic is exempt from inspection, including URL filtering and intrusion protection. Consider the pros and cons before deciding on this option.

h) Review the RA VPN configuration, then click **Finish**.

Step 2 Configure the NAT rule to translate all connections going out the outside interface to ports on the outside IP address (interface PAT).

When you complete the initial device configuration, the system creates a NAT rule named InsideOutsideNatRule. This rule applies interface PAT to IPv4 traffic from any interface that exits the device through the outside interface. Because the outside interface is included in "Any" source interface, the rule you need already exists, unless you edited it or deleted it.

The following procedure explains how to create the rule you need.

a) Click **Policies** > **NAT**.

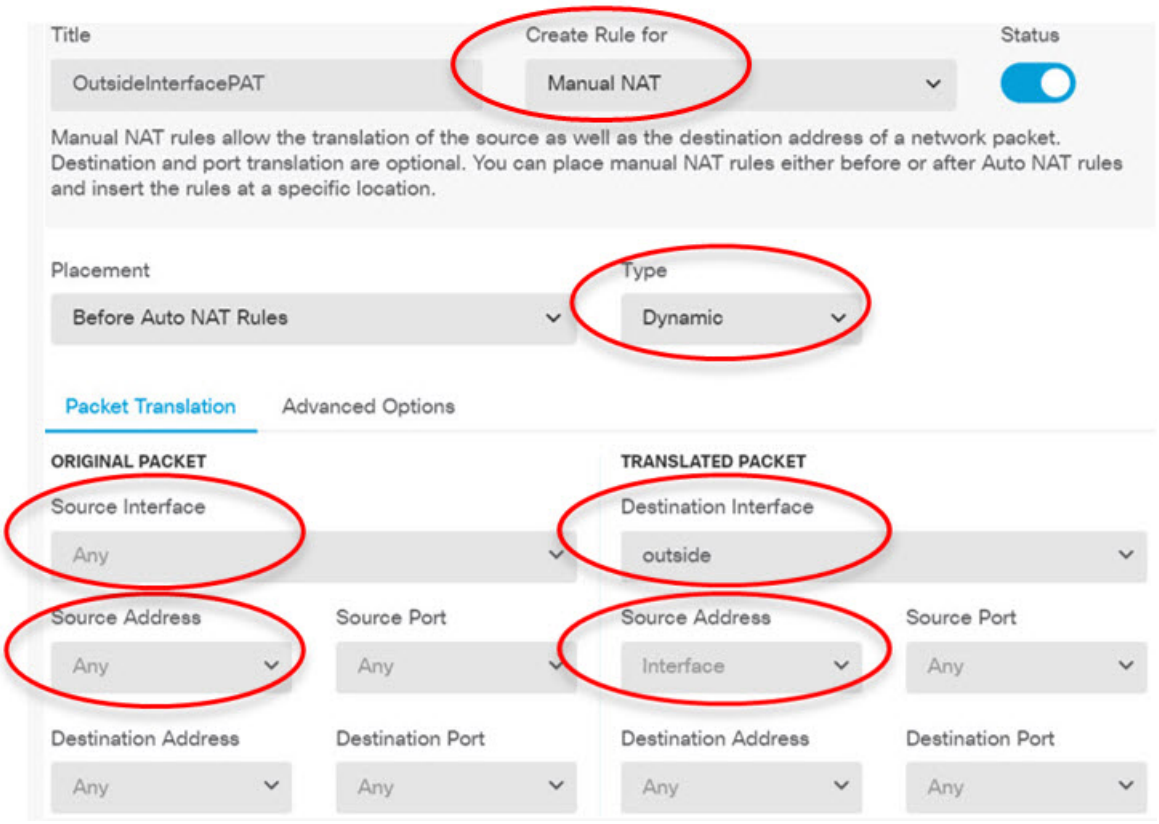
b) Do one of the following:

- To edit the InsideOutsideNatRule, mouse over the **Action** column and click the edit icon ( ).
- To create a new rule, click **+**.

c) Configure a rule with the following properties:

- **Title**—For a new rule, enter a meaningful name without spaces. For example, OutsideInterfacePAT.
- **Create Rule For**—**Manual NAT**.
- **Placement**—**Before Auto NAT Rules** (the default).
- **Type**—**Dynamic**.
- **Original Packet**—For **Source Address**, select either Any or any-ipv4. For **Source Interface**, ensure that you select Any (which is the default). For all other Original Packet options, keep the default, Any.
- **Translated Packet**—For **Destination Interface**, select outside. For **Translated Address**, select **Interface**. For all other Translated Packet options, keep the default, Any.

The following graphic shows the simple case where you select Any for the source address.



d) Click **OK**.

Step 3 (If you do not configure **Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)** in the connection profile.) Configure an access control rule to allow access from the remote access VPN address pool.

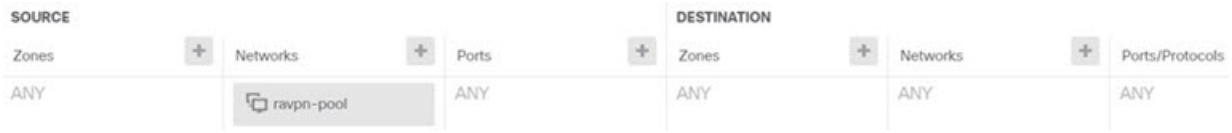
If you select the **Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)** in the connection profile, traffic from RA VPN pool addresses bypasses the access control policy. You cannot write access control rules that will apply to the traffic. You need to write rules only if you disable the option.

The following example allows traffic from the address pool to any destination. You can adjust this to meet your specific requirements. You can also precede the rule with block rules to filter out undesirable traffic.

a) Click **Policies** > **Access Control**.
b) Click **+** to create a new rule.
c) Configure a rule with the following properties:

- **Order**—Select a position in the policy before any other rule that might match these connections and block them. The default is to add the rule to the end of the policy. If you need to reposition the rule later, you can edit this option or simply drag and drop the rule to the right slot in the table.
- **Title**—Enter a meaningful name without spaces. For example, RAVPN-address-pool.
- **Action**—**Allow**. You can select Trust if you do not want this traffic to be inspected for protocol violations or intrusions.

- **Source/Destination** tab—For **Source** > **Network**, select the same object you used in the RA VPN connection profile for the address pool. Leave the default, Any, for all other Source and Destination options.



- **Application**, **URL**, and **Users** tabs—Leave the default settings on these tabs, that is, nothing selected.
- **Intrusion**, **File** tabs—You can optionally select intrusion or file policies to inspect for threats or malware.
- **Logging** tab—You can optionally enable connection logging.

d) Click **OK**.

Step 4 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.



b) Click the **Deploy Now** button.

You can wait until deployment completes, or click **OK** and check the task list or deployment history later.

How to Use a Directory Server on an Outside Network with Remote Access VPN

You can configure a remote access VPN to allow mobile workers and telecommuters to securely connect to your internal networks. Security of the connection depends on your directory server, which authenticates the user connection to ensure that only authorized users can gain entry.

If your directory server is on an outside network rather than an inside network, you need to configure a site-to-site VPN connection from the outside interface to the network that includes the directory server. **There is one trick to the site-to-site VPN configuration:** you must include the outside interface address of the remote access VPN device within the "inside" networks of the site-to-site VPN connection, and also in the remote networks for the device behind which the directory server resides. This will be explained further in the following procedure.



Note If you use the data interfaces as a gateway for the virtual management interface, this configuration also enables usage of the directory for identity policies. If you do not use data-interfaces as the management gateway, ensure that there is a route from the management network to the inside network that participates in the site-to-site VPN connection.

This use case implements the following network scenario.

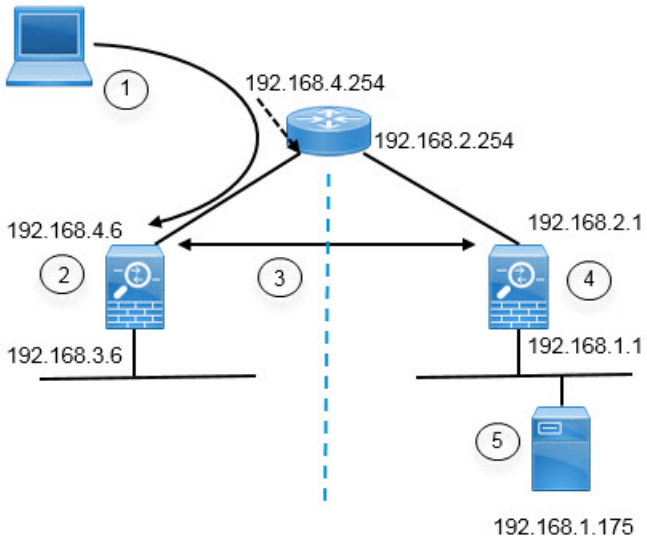


Figure Callout	Description
1	Remote access host that makes a VPN connection to 192.168.4.6. Clients will get an address in the 172.18.1.0/24 address pool.
2	Site A, which hosts the remote access VPN.
3	The site-to-site VPN tunnel between the outside interfaces of the Site A and Site B the Firewall Threat Defense devices.
4	Site B, which hosts the directory server.
5	The directory server, on the inside network of Site B.

Before you begin

This use case assumes that you followed the device setup wizard to establish a normal baseline configuration. Specifically:

- There is an Inside_Outside_Rule access control rule that allows (or trusts) traffic going from the inside_zone to the outside_zone.
- The inside_zone and outside_zone security zones contain the inside and outside interfaces (respectively).
- There is an InsideOutsideNATRule that performs interface PAT for all traffic coming from inside interfaces going to the outside interface. On devices that use an inside bridge group by default, there might be several rules for interface PAT.
- There is a static IPv4 route for 0.0.0.0/0 that points to the outside interface. This example assumes that you are using static IP addresses for the outside interfaces, but you could also use DHCP and obtain the static route dynamically. For this example, we are assuming the following static routes:
  - Site A: outside interface, gateway is 192.168.4.254.
  - Site B: outside interface, gateway is 192.168.2.254.

Procedure

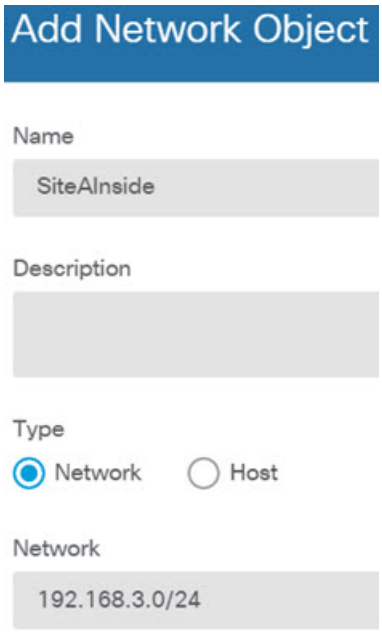
Step 1 Configure the site-to-site VPN connection on **Site B**, which hosts the directory server.

a) Click **Device**, then click **View Configuration** in the Site-to-Site VPN group.
b) Click the + button.
c) Configure the following options for **Endpoint Settings**.

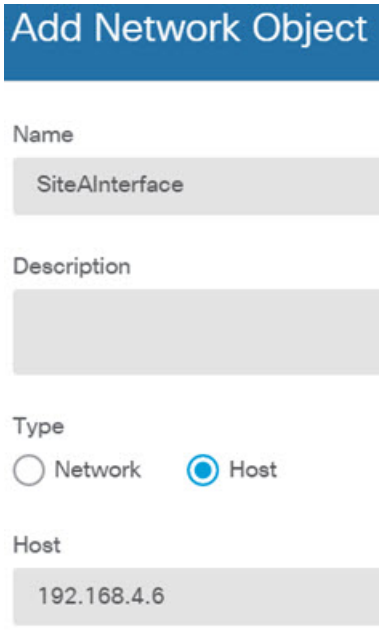
- **Connection Profile Name**—Enter a name, for example, SiteA (to indicate that the connection is to Site A).
- **Local Site**—These options define the local endpoint.
  - **Local VPN Access Interface**—Select the **outside** interface (the one with the 192.168.2.1 address in the diagram).
  - **Local Network**—Click + and select the network object that identifies the local network that should participate in the VPN connection. Because the directory server is on this network, it can participate in the site-to-site VPN. Assuming that the object does not already exist, click **Create New Network** and configure an object for the 192.168.1.0/24 network. After saving the object, select it in the drop-down list and click **OK**.



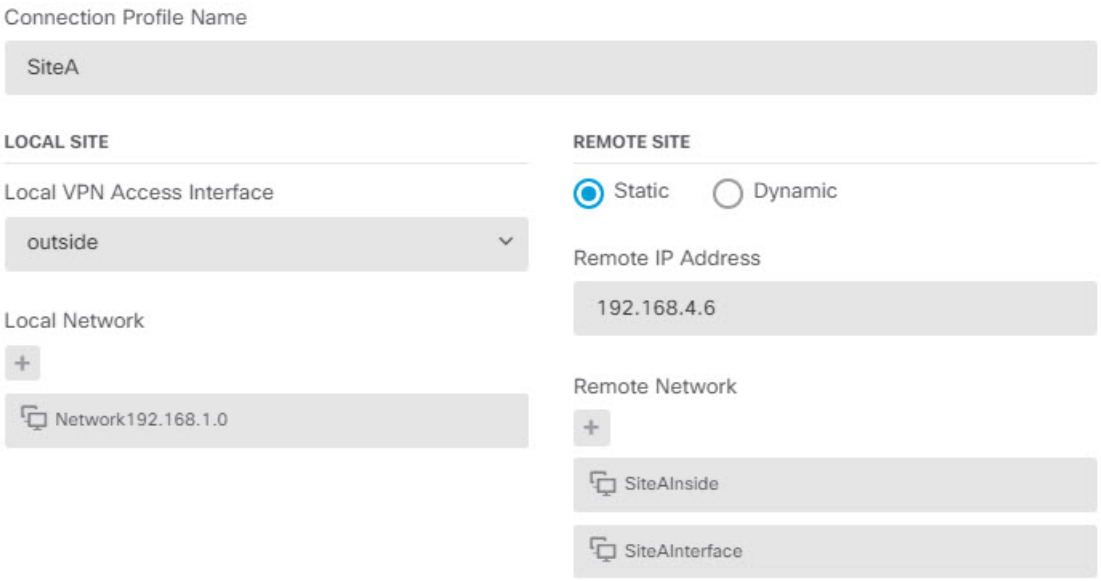
- **Remote Site**—These options define the remote endpoint.
  - **Remote IP Address**—Enter 192.168.4.6, which is the IP address of the remote VPN peer's interface that will host the VPN connection.
  - **Remote Network**—Click + and select the network objects that identify the remote networks that should participate in the VPN connection. Click **Create New Network**, configure the following objects, then select them in the list.
    1. SiteAInside, Network, 192.168.3.0/24.



2. SiteAInterface, Host, 192.168.4.6. **This is key: you must include the remote access VPN connection point address as part of the remote network for the site-to-site VPN connection so that the RA VPN hosted on that interface can use the directory server.**



When you are finished, the endpoint settings should look like the following:



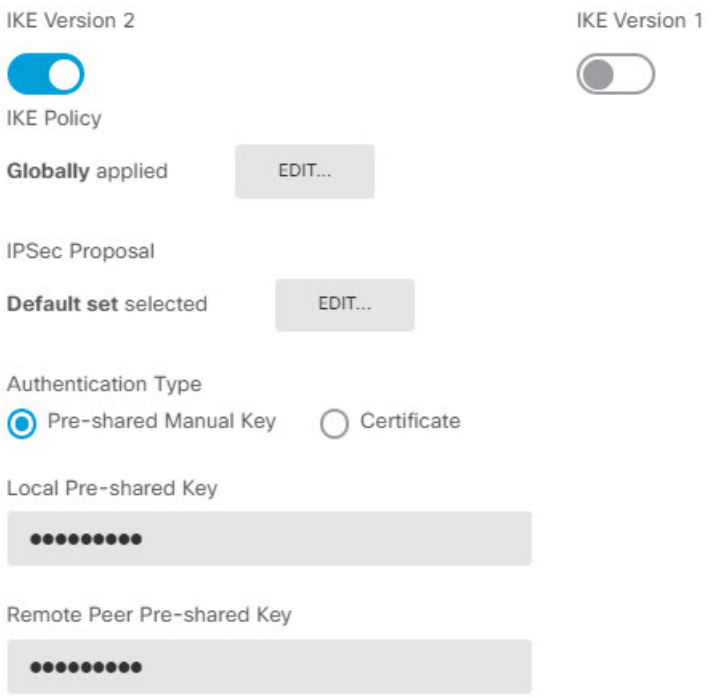
d) Click **Next**.

e) Define the privacy configuration for the VPN.

For this use case, we assume you qualify for export controlled features, which allows the use of strong encryption. Adjust these example settings to meet your needs and your license compliance.

- **IKE Version 2**, **IKE Version 1**—Keep the defaults, **IKE Version 2** enabled, **IKE Version 1** disabled.
- **IKE Policy**—Click **Edit** and enable **AES-GCM-NULL-SHA** and **AES-SHA-SHA**, and disable **DES-SHA-SHA**.
- **IPsec Proposal**—Click **Edit**. In the Select IPSec Proposals dialog box, click +, then click **Set Default** to choose the default AES-GCM proposals.
- **Local Preshared Key**, **Remote Peer Preshared Key**—Enter the keys defined on this device and on the remote device for the VPN connection. These keys can be different in IKEv2. The key can be 1-127 alphanumeric characters. **Remember these keys, because you must configure the same strings when creating the site-to-site VPN connection on the Site A device.**

The IKE policy should look like the following:



f) Configure the **Additional Options**.

- **NAT Exempt**—Select the interface that hosts the inside network, in this example, the **inside** interface. Typically, you do not want traffic within a site-to-site VPN tunnel to have their IP addresses translated. This option works only if the local network resides behind a single routed interface (not a bridge group member). If the local network is behind more than one routed interface, or one or more bridge group members, you must manually create the NAT exempt rules. For information on manually creating the required rules, see Exempting Site-to-Site VPN Traffic from NAT, on page 652.
- **Diffie-Helman Group for Perfect Forward Secrecy**—Select **Group 19**. This option determines 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. For an explanation of the options, see Decide which Diffie-Hellman modulus group to use, on page 632.

The options should look like the following.



g) Click **Next**.

h) Review the summary and click **Finish**.

The summary information is copied to the clipboard. You can paste the information in a document and use it to help you configure the remote peer, or to send it to the party responsible for configuring the peer.

i) Click the **Deploy Changes** icon in the upper right of the web page.



j) Click the **Deploy Now** button and wait for deployment to complete successfully.

Now the Site B device is ready to host one end of the site-to-site VPN connection.

Step 2 Log out of the **Site B** device and log into the **Site A** device.

Step 3 Configure the site-to-site VPN connection on **Site A**, which will host the remote access VPN.

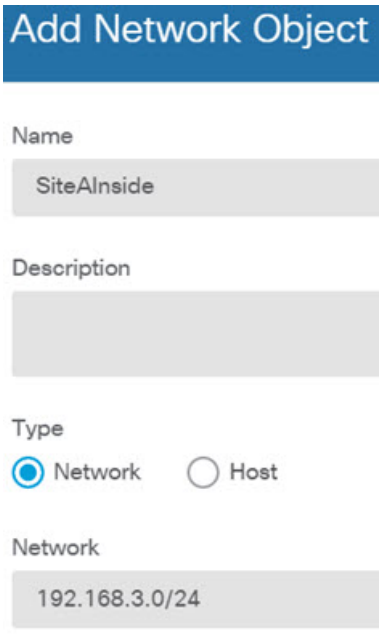
a) Click **Device**, then click **View Configuration** in the Site-to-Site VPN group.

b) Click the + button.

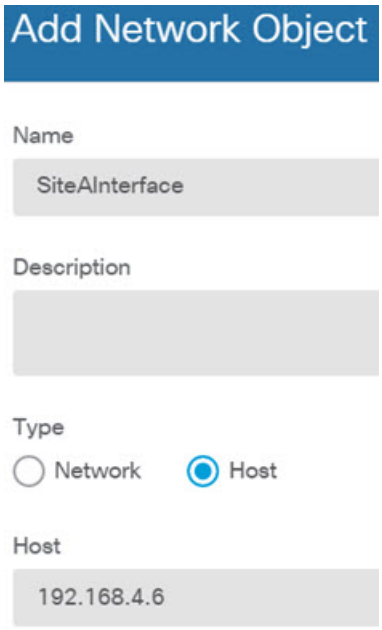
c) Configure the following options for **Endpoint Settings**.

- **Connection Profile Name**—Enter a name, for example, SiteB (to indicate that the connection is to Site B).
- **Local Site**—These options define the local endpoint.
  - **Local VPN Access Interface**—Select the **outside** interface (the one with the 192.168.4.6 address in the diagram).
  - **Local Network**—Click + and select the network objects that identify the local networks that should participate in the VPN connection. Click **Create New Network**, configure the following objects, then select them in the list. **Note that you created the same objects in the Site B device, but you have to create them again in the Site A device.**

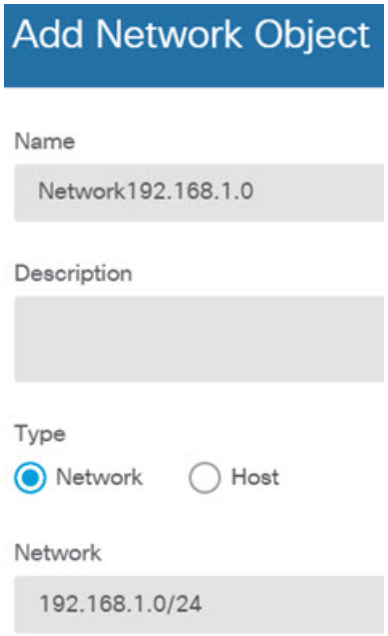
1. SiteAInside, Network, 192.168.3.0/24.



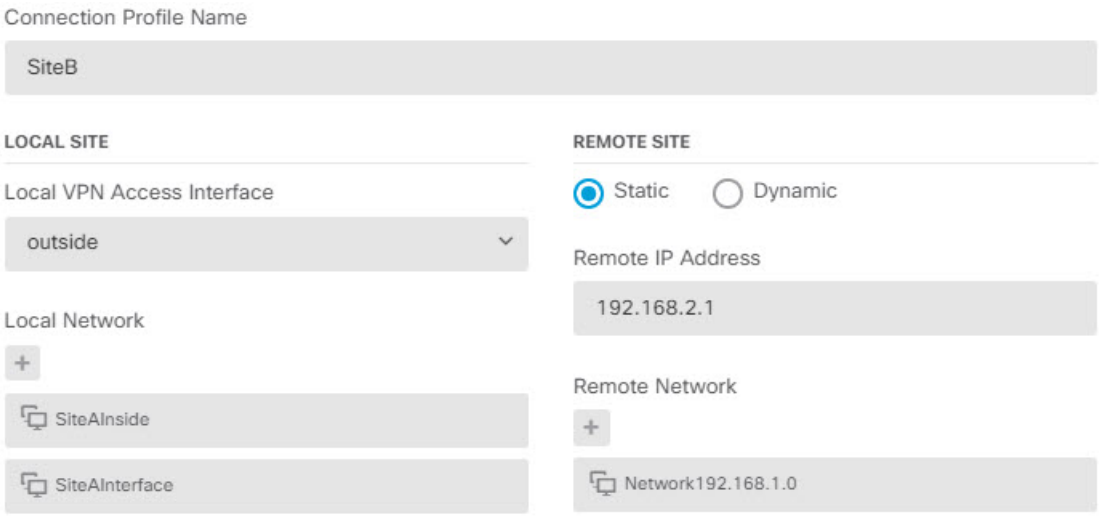
2. SiteAInterface, Host, 192.168.4.6. **This is key: you must include the remote access VPN connection point address as part of the inside network for the site-to-site VPN connection so that the RA VPN hosted on that interface can use the directory server on the remote network.**



- **Remote Site**—These options define the remote endpoint.
  - **Remote IP Address**—Enter 192.168.2.1, which is the IP address of the remote VPN peer's interface that will host the VPN connection.
  - **Remote Network**—Click + and select the network object that identifies the remote network that should participate in the VPN connection, the one that includes the directory server. Click **Create New Network** and configure an object for the 192.168.1.0/24 network. After saving the object, select it in the drop-down list and click **OK**. **Note that you created the same object in the Site B device, but you have to create it again in the Site A device.**



When you are finished, the endpoint settings should look like the following. Notice that the local/remote networks are flipped compared to the Site B settings. This is how the two ends of a point-to-point connection should always look.

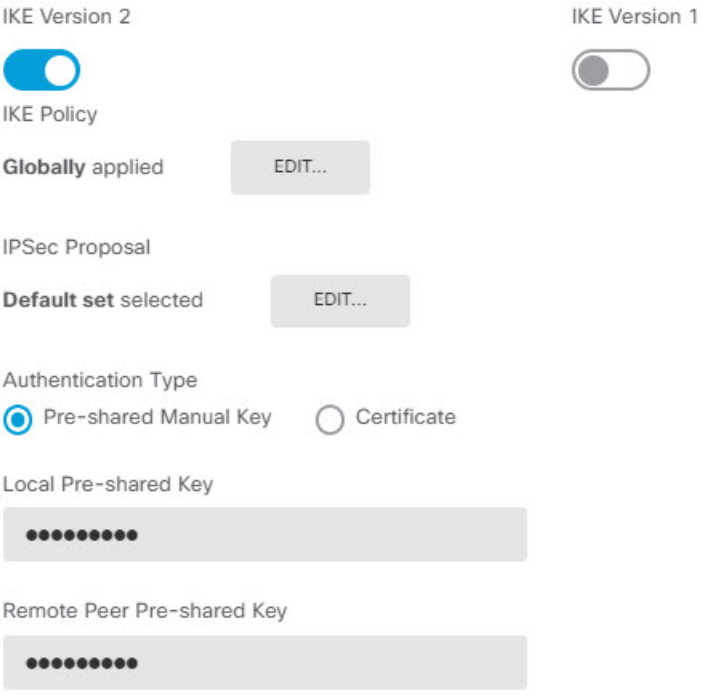


d) Click **Next**.

e) Define the privacy configuration for the VPN.

Configure the same IKE version, policy, and IPsec proposal, and the same preshared keys, as you did for the Site B connection, **but make sure that you reverse the Local and Remote preshared keys**.

The IKE policy should look like the following:



f) Configure the **Additional Options**.

- **NAT Exempt**—Select the interface that hosts the inside network, in this example, the **inside** interface. Typically, you do not want traffic within a site-to-site VPN tunnel to have their IP addresses translated. This option works only if the local network resides behind a single routed interface (not a bridge group member). If the local network is behind more than one routed interface, or one or more bridge group members, you must manually create the NAT exempt rules. For information on manually creating the required rules, see Exempting Site-to-Site VPN Traffic from NAT, on page 652.
- **Diffie-Helman Group for Perfect Forward Secrecy**—Select **Group 19**.

The options should look like the following.



g) Click **Next**.
h) Review the summary and click **Finish**.
i) Click the **Deploy Changes** icon in the upper right of the web page.



j) Click the **Deploy Now** button and wait for deployment to complete successfully.

Now the Site A device is ready to host the other end of the site-to-site VPN connection. Because Site B is already configured with compatible settings, the two devices should negotiate a VPN connection.

You can confirm the connection by logging into the device CLI and pinging the directory server. You can also use the **show ipsec sa** command to view the session information.

Step 4 Configure the directory server on **Site A**. Click **Test** to verify that there is a connection.

a) Select **Objects**, then select **Identity Sources** from the table of contents.
b) Click + > **AD**.
c) Configure the basic realm properties.

- **Name**—A name for the directory realm. For example, AD.
- **Type**—The type of directory server. Active Directory is the only supported type, and you cannot change this field.
- **Directory Username**, **Directory Password**—The distinguished username and password for a user with appropriate rights to the user information you want to retrieve. For Active Directory, the user does not need elevated privileges. You can specify any user in the domain. The username must be fully qualified; for example, Administrator@example.com (not simply Administrator).

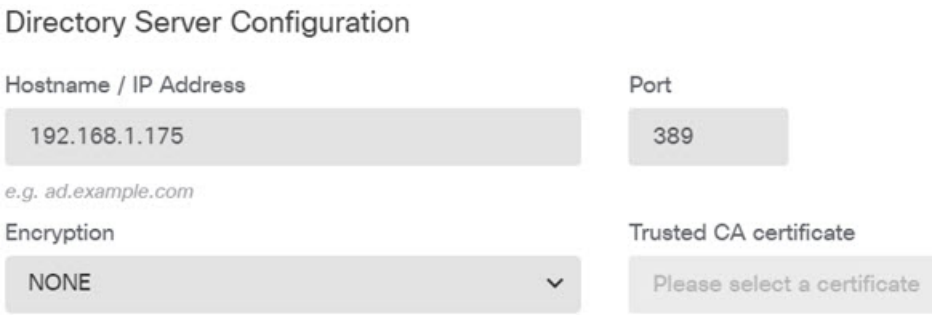
Note
The system generates ldap-login-dn and ldap-login-password from this information. For example, Administrator@example.com is translated as cn=adminisntrator,cn=users,dc=example,dc=com. Note that cn=users is always part of this translation, so you must configure the user you specify here under the common name “users” folder.

- **Base DN**—The directory tree for searching or querying user and group information, that is, the common parent for users and groups. For example, cn=users,dc=example,dc=com. For information on finding the base DN, see Determining the Directory Base DN, on page 160.
- **AD Primary Domain**— The fully qualified Active Directory domain name that the device should join. For example, example.com.



d) Configure the directory server properties.

- **Hostname/IP Address**—The hostname or IP address of the directory server. If you use an encrypted connection to the server, you must enter the fully-qualified domain name, not the IP address. For this example, enter 192.168.1.175.
- **Port**—The port number used for communications with the server. The default is 389. Use port 636 if you select LDAPS as the encryption method. For this example, keep 389.
- **Encryption**—To use an encrypted connection for downloading user and group information. The default is **None**, which means that user and group information is downloaded in clear text. For RA VPN, you can use **LDAPS**, which is LDAP over SSL. Use port 636 if you select this option. RA VPN does not support STARTTLS. For this example, select **None**.
- **Trusted CA Certificate**—If you select an encryption method, upload a Certificate Authority (CA) certificate to enable a trusted connection between the system and the directory server. 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 192.168.1.175 as the IP address but ad.example.com in the certificate, the connection fails.



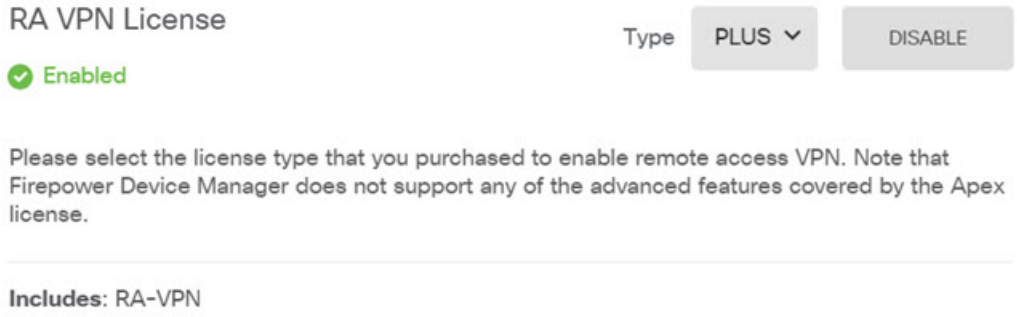
e) Click the **Test** button to verify the system can contact the server.

The system uses separate processes to access the server, so you might get errors indicating that the connection works for one type of use but not another, for example, available for Identity policies but not for remote access VPN. If the server cannot be reached, verify that you have the right IP address and host name, that the DNS server has an entry for the hostname, and so forth. Also, verify that the site-to-site VPN connection is working and that you included Site A's outside interface address in the VPN, and that NAT is not translating traffic for the directory server. You might also need to configure a static route for the server.

f) Click **OK**.

Step 5 Click **Device** > **Smart License** > **View Configuration**, and enable the RA VPN license.

When enabling the RA VPN license, select the type of license you purchased: Plus, Apex (or both), or VPN Only. For more information, see Licensing Requirements for Remote Access VPN, on page 677.

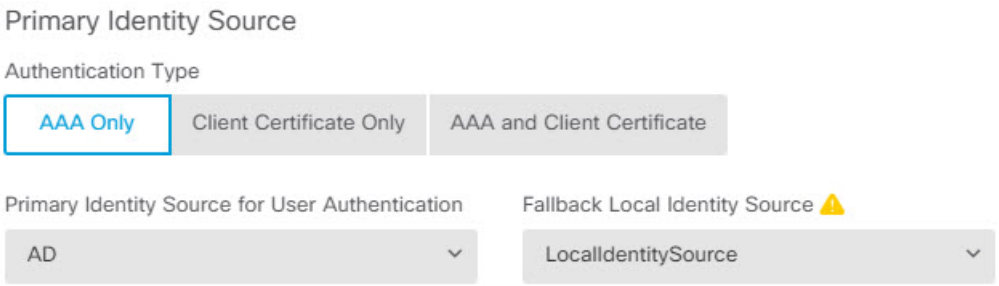


Step 6 Configure the remote access VPN on Site A.

a) Click **View Configuration** in the **Device** > **Remote Access VPN** group. Ensure that you are on the **Connection Profiles** page.

b) Create or edit a connection profile.

c) In the first step of the wizard, configure the profile name and then select the AD realm as the primary authentication source. You can optionally select the local database as fallback identity source.



d) Configure the address pool.

For this example, click **+**, then select **Create New Network** in the IPv4 address pool and create an object for the 172.18.1.0/24 network, then select the object. Clients are assigned an address from this pool. Leave the IPv6 pool blank. The address pool cannot be on the same subnet as the IP address for the outside interface.

The object should look like the following:



The pool specification should look like the following:



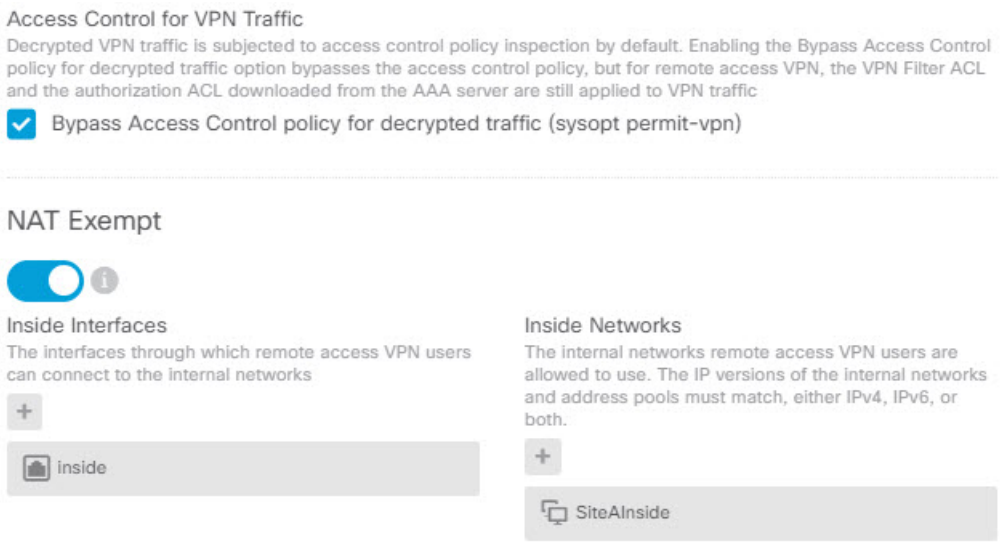
e) Click **Next**, then select an appropriate group policy.

Check the summary information about the policy you select. Ensure that the DNS servers are configured. If they are not, edit the policy now and configure DNS.

f) Click **Next**, and in global settings, select the **Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)** option, and configure the **NAT Exempt** options.

For **NAT Exempt**, you need to configure the following options. Note that if you have other connection profiles defined, you need to add to the existing settings, as the configuration applies to all connection profiles.

- **Inside Interfaces**—Select the **inside** interface. These are the interfaces for the internal networks remote users will be accessing. NAT rules are created for these interfaces.
- **Inside Networks**—Select the SiteAInside network object. These are the network objects that represent internal networks remote users will be accessing.



g) Upload the AnyConnect Client packages for the platforms you support.

h) Click **Next** and verify the settings.

First, verify that the summary is correct.

Then, click **Instructions** to see what end users need to do to initially install the AnyConnect Client software and test that they can complete a VPN connection. Click **Copy** to copy these instructions to the clipboard, and paste them in a text file or email.

i) Click **Finish**.

Step 7 Click the **Deploy Changes** icon in the upper right of the web page.



Step 8 Click the **Deploy Now** button and wait for deployment to complete successfully.

Now the Site A device is ready to accept RA VPN connections. Have an external user install the AnyConnect Client client and complete a VPN connection.

You can confirm the connection by logging into the device CLI and using the **show vpn-sessiondb anyconnect** command to view the session information.

How to Control RA VPN Access By Group

You can configure remote access VPN connection profiles to provide differential access to internal resources based on group policy. For example, if you want to provide unrestricted access to employees, but for contractors provide access to a single internal network and nothing else, you can use group policies to define different ACLs to restrict access appropriately.

The following example shows how to set up an RA VPN connection for contractors who should get access to the 192.168.2.0/24 internal subnet only. For regular employees, you can use the default group policy, which

does not have a traffic filter defined for the VPN. You can edit the default group policy if you want to apply restrictions to these users, and apply an ACL constructed as described below.

Before you begin

This procedure assumes that you have already created the identity source to use for the contractors. This might be a different source from the one you use for regular employees. Because the identity source is not strictly relevant to restricting access, we omit it from this example.

This example also assumes that the "inside2" interface is configured to host the 192.168.2.0/24 subnet, with the IP address 192.168.2.1 (any other address on the subnet is also acceptable).

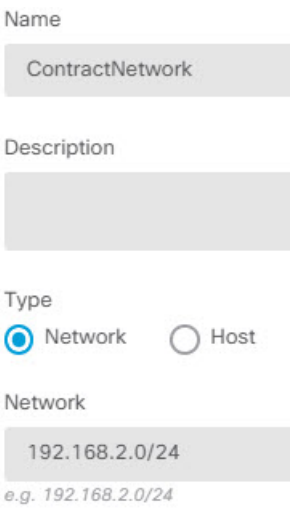
Procedure

Step 1 Configure the extended access control list (ACL) for restricting RA VPN traffic.

You need to first configure the network object that defines the target 192.168.2.0/24, then create the Smart CLI object that defines the access list. Because the ACL has an implicit deny at the end, you need only permit access to the subnet, and traffic directed to any IP address outside the subnet will be denied. This example applies to IPv4 only; you can also configure objects for restricting IPv6 access to particular subnets. Simply create the network object and add an IPv6-based ACE to the same ACL.

a) Choose **Objects** > **Networks**, and create the required object.

For example, name the object ContractNetwork. The object should look similar to the following:

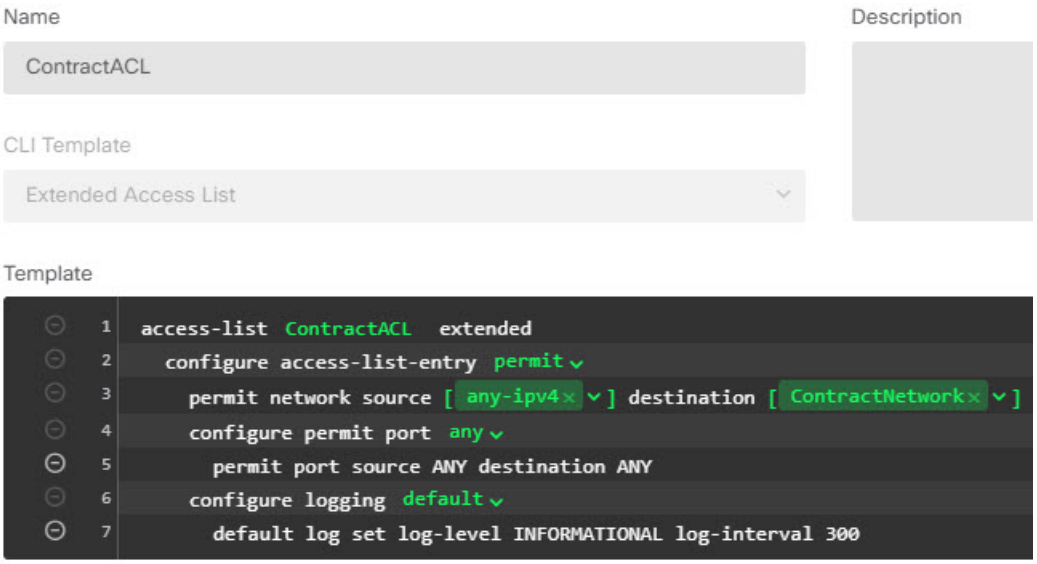


b) Choose **Device** > **Advanced Configuration** > **Smart CLI** > **Objects**.
c) Click **+** to create a new object.
d) Enter a name for the ACL. For example, **ContractACL**.
e) For **CLI Template**, select **Extended Access List**.
f) Configure the following in the **Template** body:

- configure access-list-entry action = permit
- source-network = any-ipv4

- destination-network = ContractNetwork object
- configure permit port = any
- configure logging = default

The ACE should look like the following:



g) Click **OK**.

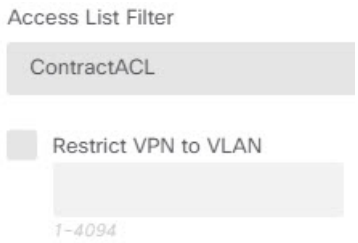
This ACL will be configured the next time you deploy changes. You do not need to use the object in any other policy to force deployment.

Step 2 Create a group policy that uses the ACL.

At minimum, you should also configure DNS servers for the group policy. You can configure other options as needed. The following procedure focuses on the one setting that is relevant for this use case.

a) Choose **Device** > **RA VPN** > **Group Policies**.
b) Click **+** to create a new group policy.
c) On the **General** page, enter a name for the policy, such as **ContractGroup**.
d) Click **Traffic Filters** in the table of contents.
e) For **Access List Filter**, select the ContractACL object.

For this example, leave the VLAN option empty. Note that you could alternatively set up a VLAN for filtering purposes, and configure a subinterface for the VLAN.



f) Click **OK** to save the group policy.

Step 3 Configure the connection profile for contractors.

a) On the RA VPN page, click **Connection Profiles** in the table of contents.

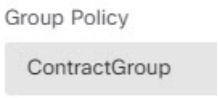
b) Click **+** to create a new connection profile.

c) Complete step 1 of the wizard and click **Next**.

Enter a name for the profile, for example, Contractors.

Configure the rest of the options as normal. This includes selecting the appropriate authentication source for the contractors, and defining an address pool.

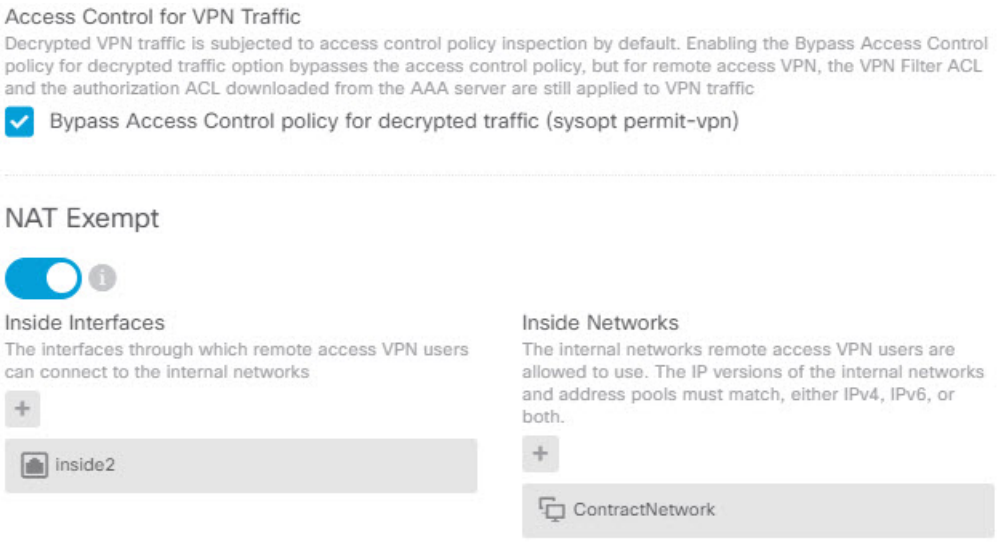
d) Select the group policy you configured for contractors and click **Next**.



e) In the global settings, select the **Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)** option, and configure the **NAT Exempt** options.

For **NAT Exempt**, you need to configure the following options. Note that if you have other connection profiles defined, you need to add to the existing settings, as the configuration applies to all connection profiles.

- **Inside Interfaces**—Select the **inside2** interface. These are the interfaces for the internal networks remote users will be accessing. NAT rules are created for these interfaces.
- **Inside Networks**—Select the ContractNetwork network object. These are the network objects that represent internal networks remote users will be accessing.



f) Upload the AnyConnect Client packages for the platforms you support.

g) Click **Next** and verify the settings.

First, verify that the summary is correct.

Then, click **Instructions** to see what end users need to do to initially install the AnyConnect Client software and test that they can complete a VPN connection. Click **Copy** to copy these instructions to the clipboard, and paste them in a text file or email.

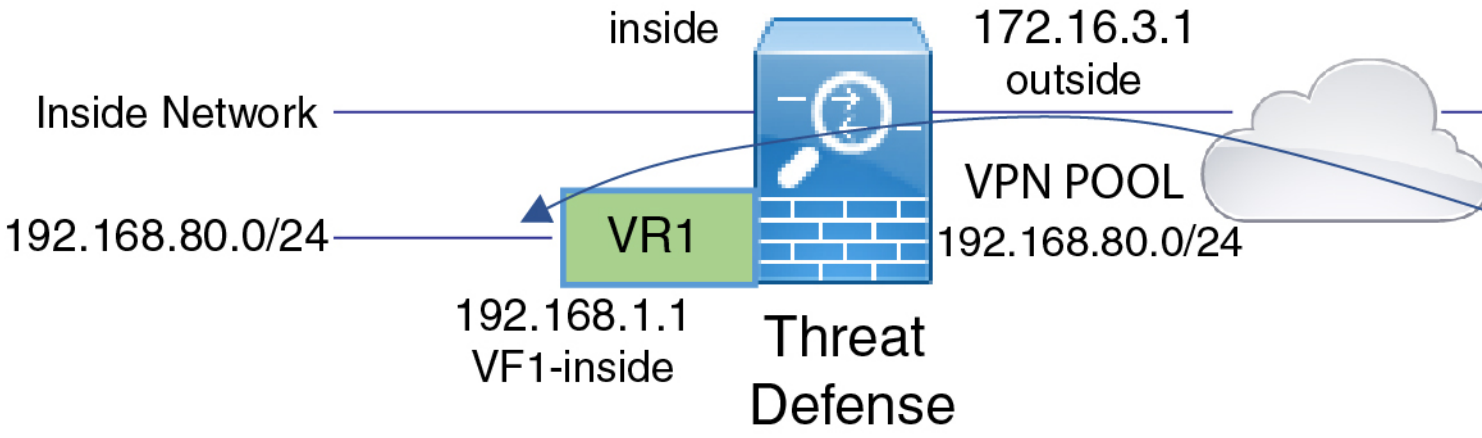
h) Click **Finish**.

How to Allow RA VPN Access to Internal Networks in Different Virtual Routers

If you configure multiple virtual routers on a device, you must configure the RA VPN in the global virtual router. You cannot configure RA VPN on an interface that is assigned to a custom virtual router.

Because the routing tables for virtual routers are separate, you must create static routes if your RA VPN users need to have access to networks that are part of a different virtual router.

Consider the following example. In this case, the RA VPN user connects to the outside interface at 172.16.3.1, and is given an IP address within the pool of 192.168.80.0/24. This user can now access the inside network that is attached to the global virtual router. However, the user cannot reach the 192.168.1.0/24 network that is part of virtual router VR1. To allow traffic flow between the VR1 network and the RA VPN user, you must configure static routes going both ways.



Before you begin

This example assumes that you have already configured the RA VPN, defined the virtual routers, and configured and assigned the interfaces to the appropriate virtual routers.

Procedure

Step 1 Configure the route leak from the Global virtual router to VR1.

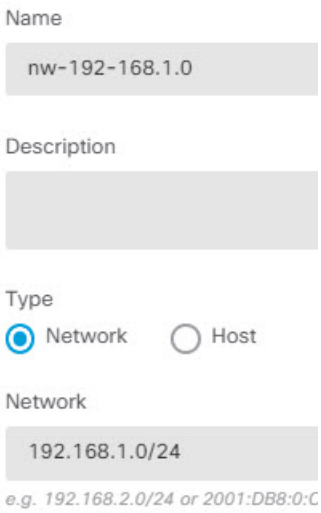
This route allows the AnyConnect Clients assigned IP addresses in the VPN pool to access the 192.168.1.0/24 network in the VR1 virtual router.

a) Choose **Device** > **Routing** > **View Configuration**.

b) Click the view icon ( ) for the Global virtual router.

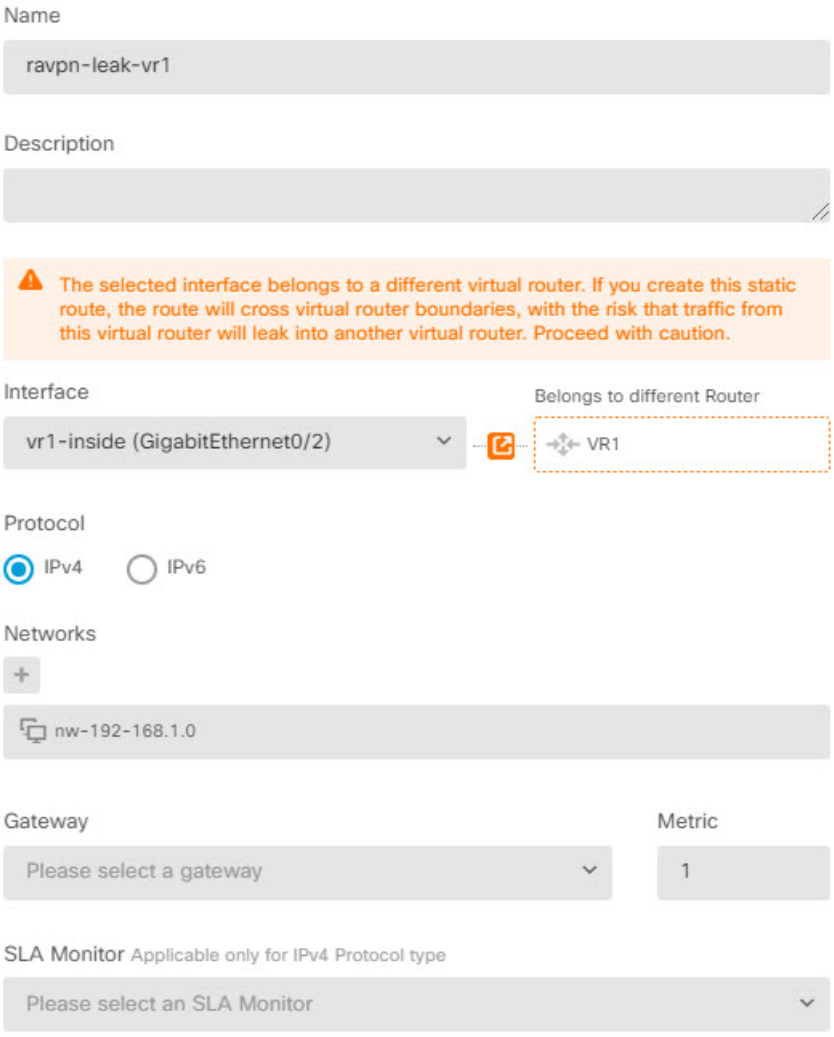
c) On the **Static Routing** tab for the Global router, click + and configure the route:

- **Name**—Any name will do, such as **ravpn-leak-vr1**.
- **Interface**—Select **vr1-inside**.
- **Protocol**—Select **IPv4**.
- **Networks**—Select an object that defines the 192.168.1.0/24 network. Click **Create New Network** to create the object now if necessary.



- **Gateway**—Leave this item blank. When leaking a route into another virtual router, you do not select the gateway address.

The dialog box should look similar to the following:



d) Click **OK**.

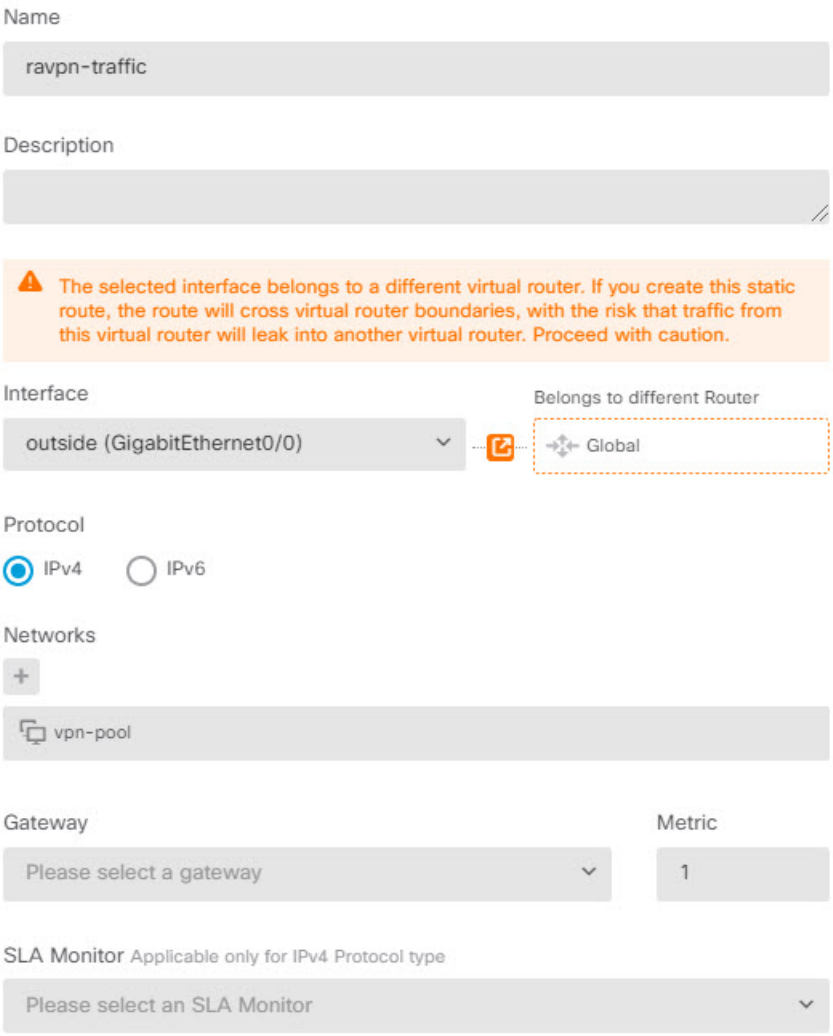
Step 2 Configure the route leak from VR1 to the Global virtual router.

This route allows endpoints on the 192.168.1.0/24 network to initiate connections to the AnyConnect Clients assigned IP addresses in the VPN pool.

a) Choose **VR1** from the virtual routers drop-down list to switch to the VR1 configuration.
b) On the **Static Routing** tab for the VR1 virtual router, click + and configure the route:

- **Name**—Any name will do, such as **ravpn-traffic**.
- **Interface**—Select **outside**.
- **Protocol**—Select **IPv4**.
- **Networks**—Select the object you created for the VPN pool, for example, **vpn-pool**.
- **Gateway**—Leave this item blank. When leaking a route into another virtual router, you do not select the gateway address.

The dialog box should look similar to the following:



c) Click **OK**.

What to do next

If there is overlap between the RA VPN address pool and the IP addresses in the custom virtual router, you must also use static NAT rules on the IP addresses to enable proper routing. However, it is far easier to simply change your RA VPN address pool so that there is no overlap.

How to Customize the AnyConnect Client Icon and Logo

You can customize the icon and logo for the AnyConnect Client app on Windows and Linux client machines. The names of the icons are pre-defined, and there are specific limits to the file type and size for the images you upload.

Although you can use any filename if you deploy your own executable to customize the GUI, this example assumes you are simply swapping icons and logos without deploying a fully-customized framework.

There are a number of images you can replace, and their file names differ based on platform. For complete information on customization options, file names, types, and sizes, please see the chapter on customizing and localizing the AnyConnect Client and installer in the *Cisco AnyConnect Secure Mobility Client Administrator Guide*. For example, the chapter for the 4.8 client is available at:

https://www.cisco.com/c/en/us/td/docs/security/vpn_client/anyconnect/anyconnect48/administration/guide/b_AnyConnect_Administrator_Guide_4-8/customize-localize-anyconnect.html

Before you begin

For the purposes of this example, we will replace the following images for Windows clients. Note that if your image is a different size than the maximum, the system will automatically resize it to the maximum, and stretch the image if necessary.

- app_logo.png

  This application logo image is the application icon, and it can have a maximum size of 128 x 128 pixels.

- company_logo.png

  This company logo image appears in the top-left corner of the tray flyout and Advanced dialogs. The maximum size is 97 x 58 pixels.

- company_logo_alt.png

  The alternative company logo image appears in the bottom-right corner of the About dialog box. The maximum size is 97 x 58 pixels.

To upload these files, you must place them on a server that the Firewall Threat Defense device can access. You can use a TFTP, FTP, HTTP, HTTPS, or SCP server. The URLs to get images from these files can include paths and uesrname/password, as required by your server setup. This example will use HTTPS.

Procedure

Step 1 Upload the image files to each Firewall Threat Defense device that is acting as an RA VPN headend that should use the customized icons and logos.

a) Log into the device CLI using an SSH client.

b) In the CLI, enter the **system support diagnostic-cli** command to enter diagnostic CLI mode.

```
> system support diagnostic-cli
Attaching to Diagnostic CLI ... Press 'Ctrl+a then d' to detach.
Type help or '?' for a list of available commands.

ftdv1>
```

Note
Read the message! You must press **Ctrl+a, then d**, to get out of the diagnostic CLI and back into the normal Firewall Threat Defense CLI mode.

c) Note the command prompt. The normal CLI uses > only, whereas the diagnostic CLI's user EXEC mode uses the hostname plus >. In this example, ftdv1>. You need to get into privileged EXEC mode, which

uses # as the ending character, for example, ftdv1#. If your prompt already has #, skip this step. Otherwise, enter the enable command, and simply press Enter at the password prompt without entering a password.

```
ftdv1> enable
Password:
ftdv1#
```

d) Use the **copy** command to copy each file from the hosting server to the Firewall Threat Defense device's disk0. You can place them in a subdirectory, such as disk0:/anyconnect-images/. You can create a new folder using the **mkdir** command.

For example, if the TFTP server's IP address is 10.7.0.80, and you want to create a new directory, the commands would be similar to the following. Note that responses to the **copy** command are omitted after the first example.

```
ftdv1# mkdir disk0:anyconnect-images

Create directory filename [anyconnect-images]? yes

Created dir disk0:/anyconnect-images

ftdv1# copy /noconfirm https://10.7.0.80/app_logo.png
disk0:/anyconnect-images/app_logo.png

Accessing https://10.7.0.80/app_logo.png...!!!!!!
Writing file disk0:/anyconnect-images/app_logo.png...
!!!!!!
12288 bytes copied in 1.000 secs (12288 bytes/sec)

ftdv1# copy /noconfirm https://10.7.0.80/company_logo.png
disk0:/anyconnect-images/company_logo.png
ftdv1# copy /noconfirm https://10.7.0.80/company_logo_alt.png
disk0:/anyconnect-images/company_logo_alt.png
```

Step 2 Use the **import webvpn** command in the diagnostic CLI to instruct the AnyConnect Client to download these images when installing itself on client machines.

import webvpn AnyConnect-customization type resource platform win name *filename* **disk0:**/*directoryname/filename*

This command is for Windows. For Linux, replace the **win** keyword with **linux** or **linux-64**, as appropriate for your clients.

For example, to import the files uploaded in the previous step, and assuming we are still in the diagnostic CLI:

```
ftdv1# import webvpn AnyConnect-customization type resource platform win
name app_logo.png disk0:/anyconnect-images/app_logo.png

ftdv1# import webvpn AnyConnect-customization type resource platform win
name company_logo.png disk0:/anyconnect-images/company_logo.png

ftdv1# import webvpn AnyConnect-customization type resource platform win
name company_logo_alt.png disk0:/anyconnect-images/company_logo_alt.png
```

Step 3 Verify the configuration:

- To verify the imported files, use the **show import webvpn AnyConnect-customization** command in the diagnostic CLI privileged EXEC mode.

- To verify that the images were downloaded to a client, they should appear when the user runs the client. You can also check the following folder on Windows clients, where %PROGRAMFILES% typically resolves to c:\Program Files.

  %PROGRAMFILES%\Cisco\Cisco AnyConnect Secure Mobility Client\res

What to do next

If you want to return to the default images, use the **revert webvpn** command (in the diagnostic CLI privileged EXEC mode) for each image you customized. The command is:

revert webvpn AnyConnect-customization type resource platform win name *filename*

As with **import webvpn**, replace **win** with **linux** or **linux-64** if you customized those client platforms, and issue the command separately for each image filename you imported. For example:

```
ftdv1# revert webvpn AnyConnect-customization type resource platform win
name app_logo.png

ftdv1# revert webvpn AnyConnect-customization type resource platform win
name company_logo.png

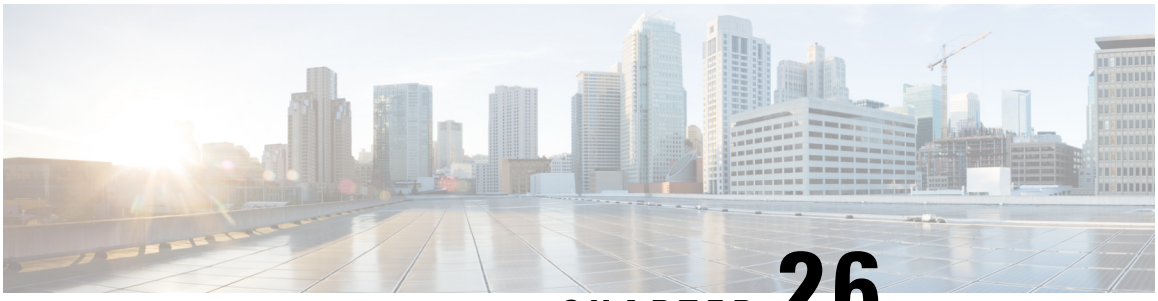
ftdv1# revert webvpn AnyConnect-customization type resource platform win
name company_logo_alt.png
```



PART VII

System Administration

- System Settings, on page 747
- System Management, on page 787



CHAPTER 26

System Settings

The following topics explain how to configure the various system settings that are grouped together on the System Settings page. The settings cover overall system function.

- Configuring Management Access, on page 747
- Configuring System Logging Settings, on page 751
- Configuring DHCP, on page 755
- Configuring Dynamic DNS, on page 759
- Configuring DNS, on page 761
- Configuring the Management Interface, on page 766
- Configuring the Device Hostname, on page 767
- Configuring Network Time Protocol (NTP), on page 768
- Configuring Precision Time Protocol (ISA 3000), on page 769
- Configuring HTTP Proxy for Management Connections, on page 771
- Configuring Cloud Services, on page 772
- Enabling or Disabling Web Analytics, on page 776
- Configuring URL Filtering Preferences, on page 777
- Switch from the Firewall Device Manager to the Firewall Management Center, on page 778
- Switch from the Firewall Management Center to the Firewall Device Manager, on page 783
- Configuring TLS/SSL Cipher Settings, on page 784

Configuring Management Access

Management access refers to the ability to log into the Firewall Threat Defense device for configuration and monitoring purposes. You can configure the following items:

- AAA to identify the identity source to use for authenticating user access. You can use the local user database or an external AAA server. For more information about administrative user management, see Managing Firewall Device Manager and Firewall Threat Defense User Access, on page 808.
- Access control to the management interface and to data interfaces. There are separate access lists for these interfaces. You can decide which IP addresses are allowed for HTTPS (used for the Firewall Device Manager) and SSH (used for CLI). See Configuring the Management Access List, on page 748.
- Management Web Server certificate, which users must accept to connect to the Firewall Device Manager. By uploading a certificate your web browsers already trust, you can avoid users being ask to trust an unknown certificate. See Configuring the Firewall Threat Defense Web Server Certificate, on page 750.

Configuring the Management Access List

By default, you can reach the device's Firewall Device Manager web or CLI interfaces on the management address from any IP address. System access is protected by username/password only. However, you can configure an access list to allow connections from specific IP addresses or subnets only to provide another level of protection.

You can also open data interfaces to allow the Firewall Device Manager or SSH connections to the CLI. You can then manage the device without using the management address. For example, you could allow management access to the outside interface, so that you can configure the device remotely. The username/password protects against unwanted connections. By default, HTTPS management access to data interfaces is enabled on the inside interface but it is disabled on the outside interface. For the Firepower 1010 that has a default "inside" bridge group, this means that you can make the Firewall Device Manager connections through any data interface within the bridge group to the bridge group IP address (default is 192.168.95.1). You can open a management connection only on the interface through which you enter the device.

Caution If you constrain access to specific addresses, you can easily lock yourself out of the system. If you delete access for the IP address that you are currently using, and there is no entry for "any" address, you will lose access to the system when you deploy the policy. Be very careful if you decide to configure the access list.

Before you begin

You cannot configure both the Firewall Device Manager access (HTTPS access) and remote access SSL VPN on the same interface for the same TCP port. For example, if you configure remote access SSL VPN on the outside interface, you cannot also open the outside interface for HTTPS connections on port 443. If you configure both features on the same interface, ensure that you change the HTTPS port for at least one of these services to avoid a conflict.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Management Access** link.

If you are already on the System Settings page, simply click **Management Access** in the table of contents.

You can also configure AAA on this page to allow management access for users defined in an external AAA server. For details, see Managing Firewall Device Manager and Firewall Threat Defense User Access, on page 808.

Step 2 To create rules for the management address:

a) Select the **Management Interface** tab.

The list of rules defines which addresses are allowed access to the indicated port: 443 for the Firewall Device Manager (the HTTPS web interface), 22 for the SSH CLI.

The rules are not an ordered list. If an IP address matches any rule for the requested port, the user is allowed to attempt logging into the device.

Note

To delete a rule, click the trash can icon ( ) for the rule. If you delete all of the rules for a protocol, no one can access the device on that interface using the protocol.

b) Click + and fill in the following options:

- **Protocol**—Select whether the rule is for HTTPS (port 443) or SSH (port 22).
- **IP Address**—Select the network object that defines the IPv4 or IPv6 network or host that should be able to access the system. To specify "any" address, select **any-ipv4** (0.0.0.0/0) and **any-ipv6** (::/0).

c) Click **OK**.

Step 3 To create rules for data interfaces:

a) Select the **Data Interfaces** tab.

The list of rules defines which addresses are allowed access to the indicated port on the interface: 443 for the Firewall Device Manager (the HTTPS web interface), 22 for the SSH CLI.

The rules are not an ordered list. If an IP address matches any rule for the requested port, the user is allowed to attempt logging into the device.

Note

To delete a rule, click the trash can icon ( ) for the rule. If you delete all of the rules for a protocol, no one can access the device on that interface using the protocol.

b) Click + and fill in the following options:

- **Interface**—Select the interface on which you want to allow management access.
- **Protocols**—Select whether the rule is for HTTPS (port 443), SSH (port 22), or both. You cannot configure HTTPS rules for the outside interface if it is used in an remote access VPN connection profile.
- **Allowed Networks**—Select the network objects that define the IPv4 or IPv6 network or host that should be able to access the system. To specify "any" address, select **any-ipv4** (0.0.0.0/0) and **any-ipv6** (::/0).

c) (Optional.) If you want to change the HTTPS Data Port number, click the number and enter a new port. See Configuring the HTTPS Port for Management Access on Data Interfaces, on page 749.

d) Click **OK**.

Configuring the HTTPS Port for Management Access on Data Interfaces

By default, accessing the device for management purposes, either for the Firewall Device Manager or the Firewall Threat Defense API, goes through port TCP/443. You can change the management access port for data interfaces.

If you change the port, users must include the custom port on the URL to access the system. For example, if the data interface is ftd.example.com, and you change the port to 4443, then users must modify the URL to https://ftd.example.com:4443.

All data interfaces will use the same port. You cannot configure different ports per interface.

Note You cannot change the management access port for the management interface. The management interface always uses port 443.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Management Access** link.

If you are already on the System Settings page, simply click **Management Access** in the table of contents.

Step 2 Click the **Data Interfaces** tab.

Step 3 Click the **HTTPS Data Port** number.

Step 4 In the Data Interfaces Setting dialog box, change the **HTTPS Data Port** to the one you want to use.

You cannot specify the following numbers:

- 22, which is used for SSH connections.
- The port used for remote access VPN, if you configured it for any interfaces that you are also allowing for management access. Remote access VPN uses port 443 by default, but you can configure a custom port for it.
- The port used for active authentication in the identity policy, which is 885 by default.

Step 5 Click **OK**.

Configuring the Firewall Threat Defense Web Server Certificate

When you log into the web interface, the system uses a digital certificate to secure communications using HTTPS. The default certificate is not trusted by your browser, so you are shown an Untrusted Authority warning and asked whether you want to trust the certificate. Although users can save the certificate to the Trusted Root Certificate store, you can instead upload a new certificate that browsers are already configured to trust.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Management Access** link.

If you are already on the System Settings page, simply click **Management Access** in the table of contents.

Step 2 Click the **Management Web Server** tab.

Step 3 In **Web Server Certificate**, select the internal certificate to use for securing HTTPS connections to the Firewall Device Manager.

If you have not uploaded or created the certificate, click the **Create New Internal Certificate** link at the bottom of the list and create it now.

The default is the pre-defined DefaultWebserverCertificate object.

Step 4 If the certificate is not self-signed, add all intermediate and root certificates in the full trust chain to the **Trusted Chain** list.

You can add up to 10 certificates in the chain. Click + to add each intermediate certificate, and finally, the root certificate. When you click **Save** (and then **Proceed** on the dialog that warns you that the web server will restart), if a certificate is missing, you will get an error message with the common name of the next certificate in the chain that is missing. You will also get an error if you add a certificate that is not in the chain. Examine these messages carefully to identify the certificate you need to add or remove.

You can upload the certificates from here by clicking **Create New Trusted CA Certificate** after clicking +.

Step 5 Click **Save**.

The change is applied immediately, and the system restarts the web server. You do not need to deploy the configuration.

Wait a few minutes to allow the restart to finish, then refresh your browser.

Configuring System Logging Settings

You can enable system logging (syslog) for Firewall Threat Defense devices. Logging information can help you identify and isolate network or device configuration problems. You can enable syslog for diagnostic logging and for connection-related logging, including access control, intrusion prevention, and file and malware logging.

Diagnostic logging provides syslog messages for events related to device and system health, and the network configuration, that are not related to connections. You configure connection logging within individual access control rules.

Diagnostic logging 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.

For information on these messages, see *Cisco Threat Defense Syslog Messages* at https://www.cisco.com/c/en/us/td/docs/security/firepower/Syslogs/b_fptd_syslog_guide.html.

The following topics explain how to configure the logging of diagnostic and file/malware messages to various output locations.

Syslog message severity levels

The syslog message severity levels define the importance of system messages.

Table 17: Syslog message severity levels

Level Number	Severity Level	Description
0	**emergencies**	You cannot use the system.
1	**alert**	You need to take immediate action.

Level Number	Severity Level	Description
2	**critical**	The system has critical conditions to be addressed.
3	**error**	Error conditions.
4	**warning**	The system has warnings you should review.
5	**notification**	The system reports significant events you should monitor.
6	**informational**	Informational messages only.
7	**debugging**	The system generates only debugging messages Log at this level only temporarily, when debugging issues. This log level can generate many messages, which might affect system performance.



Note Firewall Threat Defense does not generate syslog messages with a severity level of zero (emergencies).

Configure Logging to a Remote Syslog Server

You can configure the system to send syslog messages to an external syslog server. This is the best option for system logging. By using an external server, you can provide more room to hold messages, and use the facilities of the server to view, analyze, and archive messages.

In addition, if you apply file policies to traffic in access control rules, to control file access or malware, or both, you can configure the system to send file event messages to an external syslog server. If you do not configure a syslog server, the events are available in the Firewall Device Manager Event Viewer only.

The following procedure explains how to enable syslog for diagnostic (data) logging and file/malware logging. You can also configure external logging for the following:

- Connection events, by selecting the syslog server on individual access control rules, SSL decryption rules, or Security Intelligence policy settings.
- Intrusion events, by selecting the syslog server in the intrusion policy settings.

Before you begin

The syslog setting for file/malware events is relevant only if you apply file or malware policies, which require the Threat and Malware licenses.

In addition, you must ensure that the **File Events** > **Log Files** option is selected on the access control rules that apply the policies. Otherwise, no events are generated at all, either for syslog or Event Viewer.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Logging Settings** link.

If you are already on the System Settings page, simply click **Logging Settings** in the table of contents

Step 2 Under **Remote Server**, turn the **Data Logging** slider to **On** to enable logging diagnostic data-plane-generated messages to an external syslog server. Then, configure the following options:

- **Syslog Server**—Click + and select one or more syslog server object and click **OK**. If the objects do not exist, click the **Add Syslog Server** link and create them now. For more information, see Configuring Syslog Servers, on page 143.
- **Severity Level for Filtering FXOS Chassis Syslogs**—For certain device models that use FXOS, the severity level for syslog messages generated by the base FXOS platform. This option appears only if it is relevant for your device. Select the severity level. Messages at this level or higher are sent to the syslog server.
- **Message Filtering**—Select one of the following options to control the messages generated for the Firewall Threat Defense operating system.
  - **Severity Level for Filtering All Events**—Select the severity level. Messages at this level or higher are sent to the syslog server.
  - **Custom Logging Filter**—If you want to do additional message filtering, so you get only those messages that interest you, select the event list filter that defines the messages you want to generate. If the filter does not already exist, click **Create New Event List Filter** and create it now. For more information, see Configure Event List Filters, on page 754.

Step 3 Turn the **File/Malware** slider to **On** to enable logging to an external syslog server for file and malware events. Then, configure the options for file/malware logging:

- **Syslog Server**—Select the syslog server object. If the object does not exist, click the **Add Syslog Server** link and create it now.
- **Log at Severity Level**—Select a severity level that should be assigned to the file/malware events. Because all file/malware events are generated at the same severity, no filtering is performed; you will see all events no matter which level you pick. This will be the level shown in the severity field of the message (that is, the x in FTD-x-<message_ID>). File events are message ID 430004, malware events are 430005.

Step 4 Click **Save**.

Configure Logging to the Internal Buffer

You can configure the system to save syslog messages to an internal logging buffer. Use the **show logging** command in the CLI or CLI Console to view the contents of the buffer.

New messages append to the end of the buffer. When the buffer fills up, the system clears the buffer and continues adding messages to it. When the log buffer is full, the system deletes the oldest message to make room in the buffer for new messages.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Logging Settings** link.

If you are already on the System Settings page, simply click **Logging Settings** in the table of contents

Step 2 Turn the **Internal Buffer** slider to **On** to enable the buffer as a logging destination.

Step 3 Configure the options for internal buffer logging:

- **Severity Level for Filtering All Events**—Select the severity level. Messages at this level or higher are sent to the internal buffer.
- **Custom Logging Filter**—(Optional.) If you want to do additional message filtering, so you get only those messages that interest you, select the event list filter that defines the messages you want to generate. If the filter does not already exist, click **Create New Event List Filter** and create it now. For more information, see Configure Event List Filters, on page 754.
- **Buffer Size**—The size of the internal buffer to which syslog messages are saved. When the buffer fills up, it is overwritten. The default is 4096 bytes. The range is 4096 to 52428800.

Step 4 Click **Save**.

Configure Logging to the Console

You can configure the system to send messages to the console. These messages appear when you log into the CLI on the Console port. You can also see these logs in an SSH session to other interfaces (including the management address) by using the **show console-output** command. In addition, you can see these messages in real time in the diagnostic CLI, enter **system support diagnostic-cli** from the main CLI.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Logging Settings** link.

If you are already on the System Settings page, simply click **Logging Settings** in the table of contents

Step 2 Turn the **Console Filter** slider to **On** to enable the console as a logging destination.

Step 3 Select the **Severity** level. Messages at this level or higher are sent to the console.

Step 4 Click **Save**.

Configure Event List Filters

An event list filter 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 a to fine-tune which messages are sent based on a combination of event class, severity, and message identifier (ID).

You would use a filter only if limiting messages by severity level alone is insufficient for your purposes.

The following procedure explains how to create the filter from the **Objects** page. You can also create a filter when you are configuring a logging destination that can use a filter.

Procedure

Step 1 Select **Objects**, then select **Event List Filters** from the table of contents.

Step 2 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 3 Configure the filter properties:

- **Name**—The name of the filter object.
- **Description**—An optional description of the object.
- **Severity and Log Class**—If you want to filter by message class, click +, select a severity level for the class filter and click **OK**. Then, click the drop-down arrow within the severity level, select one or more classes to filter at that severity level, and click **OK**.

  The system will send syslog messages for the specified classes of messages only if they are at that severity level or higher. You can add at most one row for each severity level.

  If you want to filter all classes at a given severity level, leave the Severity list empty and instead select the global severity level for the logging destination when you enable it.
- **Syslog Range/Message ID**—If you want to filter by the syslog message ID, enter a single message ID, or a range of ID numbers for which you want to generate messages. Separate the starting and ending number for a range with a hyphen, for example, 100000-200000. The IDs are 6 digit numbers. For specific message IDs and the related messages, see *Cisco Threat Defense Syslog Messages* at https://www.cisco.com/c/en/us/td/docs/security/firepower/Syslogs/b_fptd_syslog_guide.html.

Step 4 Click **Save**.

You can now select this object in the custom filtering option for logging destinations that allow it. Go to **Device** > **System Settings** > **Logging Settings**.

Configuring DHCP

A DHCP server provides network configuration parameters, such as IP addresses, to DHCP clients. You can either configure DHCP servers on interfaces to provide configuration parameters to DHCP clients on the attached network, or enable DHCP relay on the interfaces to forward requests to an external DHCP server that is operating on another device in the network.

These features are mutually exclusive: you can configure one or the other, but not both.

Configuring the DHCP Server

A DHCP server provides network configuration parameters, such as IP addresses, to DHCP clients. You can configure a DHCP server on an interface to provide configuration parameters to DHCP clients on the attached network.

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 does not support BOOTP requests.



Note Do not configure a DHCP server on a network that already has a DHCP server operating on it. The two servers will conflict and results will be unpredictable.

Before you begin

DHCP clients must be on the same network as the interface on which the server is enabled. That is, there cannot be an intervening router between the server and client, although there can be a switch.

If you must support multiple networks, and do not want to configure a DHCP server on each interface, you can instead configure DHCP relay to forward DHCP requests from one network to a DHCP server that resides on a different network. In this case, the DHCP server must reside on a different device in the network: you cannot configure a DHCP server on one interface and DHCP relay on another interface on the same device. When using DHCP relay, ensure that you configure the DHCP server with address pools for each network address space that the DHCP server will manage.

To configure DHCP relay, see Configuring DHCP Relay, on page 758.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **DHCP Server / Relay** link.

If you are already on the System Settings page, simply click **DHCP** > **DHCP Server** in the table of contents.

The page has two tabs. Initially, the **Configuration** tab shows the global parameters.

The **DHCP Servers** tab shows the interfaces on which you have configured DHCP server, whether the server is enabled, and the address pool for the server.

Step 2 On the **Configuration** tab, configure auto-configuration and global settings.

DHCP auto configuration enables the DHCP Server to provide DHCP clients with DNS server, domain name, and WINS server information obtained from a DHCP client that is running on the specified interface. Typically, you would use auto-configuration if you are obtaining an address using DHCP on the outside interface, but you could choose any interface that obtains its address through DHCP. If you cannot use auto-configuration, you can manually define the required options.

a) Click **Enable Auto Configuration** > **On** (the slider should be on the right) if you want to use auto-configuration, and then select the interface that is obtaining its address through DHCP in **From Interface**.

If you configure virtual routers, you can use DHCP server auto configuration on an interface in the global virtual router only. Auto configuration is not supported for interfaces assigned to a user-defined virtual router.

b) If you do not enable auto-configuration, or if you want to override any of the automatically configured settings, configure the following global options. These settings will be sent to DHCP clients on all interfaces that host DHCP server.

- **Primary WINS IP Address**, **Secondary WINS IP Address**—The addresses of the Windows Internet Name Service (WINS) servers clients should use for NetBIOS name resolution.
- **Primary DNS IP Address**, **Secondary DNS IP Address**—The addresses of the Domain Name System (DNS) servers clients should use for domain name resolution. Click **Use OpenDNS** if you want to configure the OpenDNS public DNS servers. Clicking the button loads the appropriate IP addresses into the fields.

c) Click **Save**.

Step 3 Click the **DHCP Servers** tab and configure the servers.

a) Do one of the following:

- To configure DHCP server for an interface that is not already listed, click +.
- To edit an existing DHCP server, click the edit icon () for the server.

To delete a server, click the trash can icon () for the server.

b) Configure the server properties:

- **Enable DHCP Server**—Whether to enable the server. You can configure a server but keep it disabled until you are ready to use it.
- **Interface**—Select the interface on which you will provide DHCP addresses to clients. The interface must have a static IP address; you cannot be using DHCP to obtain the interface address if you want to run a DHCP server on the interface. For bridge groups, you configure the DHCP server on the Bridge Virtual Interface (BVI), not the member interfaces, and the server operates on all member interfaces.

  You cannot configure DHCP server on the Diagnostic interface; configure it on the Management interface instead, on the **Device** > **System Settings** > **Management Interface** page.
- **Address Pool**—The range of IP addresses from lowest to highest that the server is allowed to provide to clients that request an address. Specify the start and end address for the pool, separated by a hyphen. For example, 10.100.10.12-10.100.10.250.

  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, the broadcast address, or the subnet network address.

  The size of the address pool is limited to 256 addresses per pool. If the address pool range is larger than 253 addresses, the netmask of the Firewall Threat Defense interface cannot be a Class C address (for example, 255.255.255.0) and needs to be something larger, for example, 255.255.254.0.

c) Click **OK**.

Configuring DHCP Relay

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 Firewall Threat Defense device because it does not forward broadcast traffic. The DHCP relay agent lets you configure the interface of the Firewall Threat Defense device that is receiving the broadcasts to forward DHCP requests to a DHCP server that is available through another interface.

Thus, clients on subnets that do not host a DHCP server can still get IP address leases from a DHCP server that resides on a different subnet.

Before you begin

- Configure the DHCP server with address pools for each of the subnets you are adding. For example, if you enable the DHCP relay client on an interface with the 192.168.1.1/24 address, to support clients on the 192.168.1.0/24 network, the DHCP server must be able to supply IP addresses on the 192.168.1.0/24 subnet, for example, 192.168.1.2-192.168.1.254.
- Create host network objects for each of the DHCP servers, specifying the IP address of the server.
- Ensure that you have removed or disabled all servers on the **DHCP** > **DHCP Servers** page. You cannot host a DHCP server on any interface with DHCP relay enabled on an interface, even if they are different interfaces.
- Interface limitations—An interface must have a name to be used for either server or agent. In addition:
  - The interface cannot be a member of a routing ECMP traffic zone.
  - The interface cannot obtain its address using DHCP.
  - You can configure both DHCP server and DHCP relay on physical interfaces, subinterfaces, VLAN interfaces, and EtherChannels (but not their members).
  - You can also configure DHCP relay server on virtual tunnel interfaces (VTI).
  - Neither service supports the management interface, or bridge groups and their members.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **DHCP Server / Relay** link, then click **DHCP** > **DHCP Relay** in the table of contents.

If you are already on the System Settings page, simply click **DHCP** > **DHCP Relay** in the table of contents.

Step 2 (Optional.) Adjust the **IPv4 Relay Timeout** and **IPv6 Relay Timeout** settings as needed.

These timeouts set the number of seconds that are allowed for DHCP relay address negotiation for the given IP version. The default is 60 seconds (1 minute), but you can set a different timeout from 1-3600 seconds. Longer timeouts might be appropriate if there is significant lag between the subnet and the DHCP server.

Step 3 Configure the **DHCP Relay Servers**.

The DHCP relay servers are the DHCP servers in the network that should service DHCP relay requests. These DHCP servers reside on different devices in the network from the device you are configuring.

a) Click **+**, select a host network object that has the IP address of a DHCP server, and click **OK**.

If the object does not yet exist, click **Create New Network** and create it now. If you no longer want to use a DHCP server you had added, click the **X** on the right of the server's entry to delete it.

b) Click the DHCP server entry you added, and select the interface through which the DHCP server can be reached.

Step 4 Configure the DHCP Relay Agents.

The DHCP relay agents run on the interfaces. They forward DHCP requests from clients on their network segment to the DHCP servers, then return the responses to the client.

a) Click **+**, select the interfaces that should run the DHCP relay agent, and click **OK**.

If you no longer want run the DHCP relay agent on an interface, click the **X** on the right of the server's entry to delete it. Optionally, you can simply disable all DHCP relay services without removing the interface from the table.

b) Click the interface entry you added, select the DHCP services you want the agent to provide, and click **OK**.

- **Enable IPv4**—Forward IPv4 address requests to the DHCP server. If you do not select this option, any IPv4 address requests are ignored, and the client cannot obtain an IPv4 address.
- **Set Route** (IPv4 only)—Change the first default router address in the packet sent from the DHCP server to the address of the Firewall Threat Defense device interface that is running the DHCP relay agent. This action allows the client to set its default route to point to the Firewall Threat Defense device even if the DHCP server specifies a different router. If there is no default router option in the packet, the DHCP relay agent adds one containing the interface address.
- **Enable IPv6**—Forward IPv6 address requests to the DHCP server. If you do not select this option, any IPv6 address requests are ignored, and the client cannot obtain an IPv6 address.

Step 5 Click **Save**.

Configuring Dynamic DNS

You can configure the system to use the web update method to send Dynamic Domain Name System (DDNS) changes to Dynamic DNS services. These services then update the DNS server to use the new IP address associated with a fully-qualified domain name (FQDN). Thus, when users try to access the system using a hostname, DNS will resolve the name to the correct IP address.

Using DDNS can help ensure that the FQDNs defined for the interfaces on the system always resolve to the correct IP address. This is especially important if you configure an interface to get its address using DHCP. But there is also value in using it for static IP addresses, to ensure the DNS server has the correct addresses, and that it can be easily updated if you change the static address.

You can configure DDNS to use a select group of DDNS service providers, or use the custom option to direct updates to any other DDNS provider that supports web updates. The FQDNs you specify for interfaces should be registered with these service providers.

Note You can use the Firewall Device Manager to configure web update DDNS only. You cannot configure DDNS for the method defined in IETF RFC 2136.

Before you begin

The system must have a trusted CA certificate that will validate the provider's certificate, or the DDNS connection will not succeed. You can download the certificates from the service provider's site. Please ensure the appropriate certificate is uploaded and deployed. Also ensure that you set the **Validation Usage** for the uploaded certificate to include **SSL Server**. See Uploading Trusted CA Certificates, on page 154.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **DDNS Service** link.

If you are already on the System Settings page, simply click **DDNS Service** in the table of contents.

The page shows a list of DDNS update methods, including the service provider, interface, fully-qualified domain name (FQDN) for the interface, and how often the DNS server will be updated for changes to the FQDN's IP address. You can click the **Show Status** link for an entry to check whether it is working correctly.

Step 2 Do one of the following:

- To create a new Dynamic DNS update method, click **+** or the **Create DDNS Service** button.
- To edit an existing Dynamic DNS update method, click the edit icon ( ) for the method.

To delete a method, click the trash can icon ( ) for the method.

Step 3 Configure Dynamic DNS Service properties:

- **Name**—A name for the service.
- **Web Type Update**—Select the types of addresses to be updated, based on what is supported by your DDNS service provider. The default is to update **All Addresses**, both IPv4 and IPv6. You can instead update the **IPv4 Address**, **IPv4 and One IPv6 Address**, **One IPv6 Address**, **All IPv6 Addresses**.

  Please note the following for IPv6 addresses:

  - Global addresses only are updated. The link local address is never updated.
  - Because the Firewall Device Manager allows you to configure a single IPv6 address per interface, in practice, only one IPv6 address will ever be updated.

- **Service Provider**—Select the service provider that will receive and process the dynamic DNS updates. You can use the following service providers.
  - **No-IP**—The No-IP DDNS service provider, https://www.noip.com/.

- **Dynamic DNS**—The Oracle Dynamic DNS service provider, https://account.dyn.com/.
- **Google**—The Google Domains service provider, https://domains.google.com.
- **Custom URL**—Any other DDNS service provider. You will need to enter the URL required by your selected provider, including username and password, into the **Web URL** field. The DDNS service should abide by the standards described at https://help.dyn.com/remote-access-api/.

- **Username**, **Password** (non-Custom URL methods)—The username and password, defined on the service provider's platform, to use when sending dynamic DNS updates.

  Note:

  - The username cannot include spaces, or the @ and : characters, because they would act as delimiters.
  - The password cannot include spaces, or the @ character, because it acts as delimiter. Any : character after the first : and before @ is considered part of the password.

- **Web URL** (Custom URL method)—If you selected a custom URL as the service provider, enter the URL for your dynamic DNS service. The URL must be in the following format, limited to 511 characters:

  http(s)://username:password@provider-domain/xyz?hostname=<h>&myip=<a>

  https://username:password@domain-provider/xyz?hostname=%3Ch%3E&myip=%3Ca%3E

- **Interfaces and Fully-Qualified Domain Name**—Select the interfaces whose DNS records you want updated with this service provider, then enter the fully-qualified domain name for each interface. For example, interface.example.com. Interfaces are restricted as follows:
  - You can select named physical and subinterfaces only.
  - You cannot select the following types of interface: management, BVI/EtherChannel or its members, VLANs, virtual tunnel interfaces (VTI).
  - A given interface can be selected in one DDNS update method only. You can select all interfaces that should use a service provider in the same DDNS update object.

- **Update Interval**—How often the dynamic DNS update should be sent. The default is **On Change**, which sends an update whenever the IP address for the interface changes. Alternatively, you can select **Hourly**, **Daily**, or **Monthly**. For daily and monthly, also configure the time of the day, and for monthly, the day of the month, to send updates.

Step 4 Click **OK**.

Configuring DNS

The Domain Name System (DNS) servers are used to resolve hostnames to IP addresses. You configure DNS servers during initial system setup, and these servers are applied to the data and management interfaces. You can change them after setup, and use separate sets of servers for the data and management interfaces.

At minimum, you must configure DNS for the management interface. You must also configure DNS for the data interfaces if you want to use FQDN-based access control rules, or if you want to use hostnames in CLI commands such as **ping**.

Configuring DNS is a two-step process: you configure DNS groups, then you configure DNS on the interfaces.

The following topics explain the process in more detail.

Configuring DNS Groups

DNS groups define a list of DNS servers and some associated attributes. You can configure DNS separately on the management and data interfaces. DNS servers are needed to resolve fully-qualified domain names (FQDN), such as www.example.com, to IP addresses.

After you complete the device setup wizard, you will have one or both of the following system-defined DNS groups:

- CiscoUmbrellaDNSServerGroup—This group includes the IP addresses of the DNS servers available with Cisco Umbrella. If you selected these servers during initial setup, this is the only system-defined group. You cannot change the name or server list in this group, but you can edit the other properties.
- CustomDNSServerGroup—If you do not select the Umbrella servers during device setup, the system creates this group with your list of servers. You can edit any property in this group.

Procedure

Step 1 Select **Objects**, then select **DNS Groups** from the table of contents.

Step 2 Do one of the following:

- To create a group, click the **Add Group** ( ) button.
- To edit a group, click the edit icon ( ) for the group.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 3 Configure the following properties:

- **Name**—The name of the DNS server group. The name DefaultDNS is reserved: you cannot use it.
- **DNS IP Addresses**—Enter the IP address of a DNS server. Click **Add Another DNS IP Address** to configure more than one server. If you want to remove a server address, click the delete icon ( ) for the address.

  The list is in priority order: the first server in the list is always used, and subsequent servers are used only if a response is not received from the servers above it. You can configure up to 3 servers.
- **Domain Search Name**—Enter the domain name for your network, e.g. example.com. This domain is added to hostnames that are not fully-qualified, for example, serverA instead of serverA.example.com. The name must be shorter than 63 characters to use the group for data interfaces.
- **Retries**—The number of times, from 0 to 10, to retry the list of DNS servers when the system does not receive a response. The default is 2. This setting applies to DNS groups used on the data interfaces only.
- **Timeout**—The number of seconds, from 1 to 30, to wait before trying the next DNS server. The default is 2 seconds. Each time the system retries the list of servers, this timeout doubles. This setting applies to DNS groups used on the data interfaces only.

Step 4 Click **OK**.

Configuring DNS for Data and Management Traffic

The Domain Name System (DNS) servers are used to resolve hostnames to IP addresses. There are two DNS server settings that apply to different types of traffic: data and special management traffic. Data traffic includes any services that use FQDNs for which a DNS lookup is necessary, such as Access Control Rules and Remote Access VPN. Special management traffic includes traffic originating on the Management interface such as Smart Licensing and database updates.

If you use the CLI setup wizard, you configure the management DNS servers during initial system configuration. You can also set the data and management DNS servers in the Firewall Device Manager setup wizard. You can change the DNS servers defaults using the following procedure.

You can also change the management DNS configuration in the CLI using the **configure network dns servers** and **configure network dns searchdomains** commands. If the data and management interfaces are using the same DNS group, the group is updated and on your next deployment, the changes are also applied to the data interfaces.

To determine the correct interface for DNS server communications, the Firewall Threat Defense uses a routing lookup, but which routing table is used depends on the interfaces for which you enable DNS. See the interface settings below for more information.

If you have problems with DNS resolution, see:

- Troubleshooting General DNS Problems, on page 764
- Troubleshooting DNS for the Management Interface, on page 822

Before you begin

- Ensure you have created a DNS server group. For instructions, see Configuring DNS Groups, on page 762.
- Ensure that the Firewall Threat Defense device has appropriate static or dynamic routes to access the DNS servers.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **DNS Server** link.

If you are already on the **System Settings** page, click **DNS Server** in the table of contents.

Step 2 Configure DNS for the **Data Interface**.

a) Enable DNS lookups on all interfaces or on specific interfaces. These choices also affect which routing tables are used.

Note that enabling DNS lookups on an interface is not the same as specifying the source interface for lookups. The device always uses a route lookup to determine the source interface.

- **ANY** (do not choose any interfaces)—Enables DNS lookups on all interfaces, including Management and management-only interfaces. The device checks the data routing table, and if no route is found, falls back to the management-only routing table.
- Interfaces selected but not the Diagnostic interface or a management-only interface—Enables DNS lookups on the specified interfaces. The device checks the data routing table only.
- Interfaces selected plus the Diagnostic interface or a management-only interface—Enables DNS lookups on the specified interfaces. The device checks the data routing table, and if no route is found, falls back to the management-only routing table.
- Only the Diagnostic interface or a management-only interface selected—Enables DNS lookups on Diagnostic or a management-only interface. The device checks only the management-only routing table.

b) Select the **DNS Group** that defines the servers to use on the data interfaces. If the group does not exist yet, click **Create New DNS Group** and create it now. Select **None** if you want to prevent lookups on the data interfaces.

c) (Optional.) Configure the **FQDN DNS Settings** if you use FQDN network objects in access control rules.

These options are used when resolving FQDN objects only, and are ignored for any other type of DNS resolution.

- **Poll Time**—The time, in minutes, of the polling cycle used to resolve FQDN network objects to IP addresses. FQDN objects are resolved only if they are used in the access control policy. The timer determines the maximum time between resolutions; the DNS entry's time-to-live (TTL) value is also used to determine when to update the IP address resolution, so individual FQDNs might be resolved more frequently than the polling cycle. The default is 240 (four hours). The range is 1 to 65535 minutes.
- **Expiry**—The number of minutes after a DNS entry expires (that is, the TTL obtained from the DNS server has passed) that the entry is removed from the DNS lookup table. Removing an entry requires that the table be recompiled, so frequent removals can increase the processing load on the device. Because some DNS entries can have very short TTL (as short as three seconds), you can use this setting to virtually extend the TTL. The default is 1 minute (that is, the entry is removed one minute after the TTL has passed). The range is 1 to 65535 minutes.

d) Click **Save**. You must also deploy the configuration to apply the changes to the device.

Step 3 Configure DNS for the **Management Interface**.

a) Select the **DNS Group** that defines the servers to use on the Management interface. If the group does not exist yet, click **Create New DNS Group** and create it now.

b) Click **Save**. You must deploy changes to update the management DNS servers.

Troubleshooting General DNS Problems

You must separately configure DNS servers for the Management and data interfaces. Some features do name resolution through one or the other type of interface, but not both. Sometimes, a given feature will use different resolution methods depending on how you use it.

For example, the **ping** *hostname* and **ping interface** *interface_name hostname* commands uses the data interface DNS servers to resolve the name, whereas the **ping system** *hostname* command uses the Management

interface DNS servers. This makes it possible for you to test connectivity through specific interfaces and through the routing table.

Keep this in mind when you are troubleshooting problems with hostname lookup.

For troubleshooting DNS for the Management interface, also see Troubleshooting DNS for the Management Interface, on page 822.

When You Get No Name Resolution

Following are some troubleshooting tips if name resolution is simply not happening.

- Verify that you have configured DNS servers for both the management and data interfaces. For data interfaces, use Any for the interface. Specify interfaces explicitly only if you do not want to allow DNS on some interfaces.
- If you are using the diagnostic interface for lookups on data interfaces, verify that you configured an IP address on the interface. Lookups require that the interface has an IP address.
- You cannot reach the DNS server through the Diagnostic interface or through a management-only interface, because the route lookup finds a match in the data routing table so there is no fall back to the management-only routing table. If you want to use the Diagnostic interface, make sure that is the only interface selected.
- Ping the IP address of each DNS server to verify that it is reachable. Use the **system** and **interface** keywords to test specific interfaces. If ping is unsuccessful, check your static routes and gateways. You might need to add static routes for the servers.
- If ping is successful, but name resolution is failing, check your access control rules. Verify that you are allowing DNS traffic (UDP/53) for the interfaces through which the servers are reachable. It is also possible that this traffic is getting blocked by a device that is between your system and the DNS server, so you might need to use different DNS servers.
- If ping works, there are adequate routes, and access control rules are not the problem, consider that the DNS server might not have a mapping for the FQDN. You might need to use different servers.

When You Get Wrong Name Resolution

If you are getting name resolution, but the IP address for a name is not current, there might be a caching issue. This problem would affect data-interface based features only, such as FQDN network objects used in access control rules.

The system has a local cache of DNS information obtained from previous lookups. When a new lookup is required, the system first looks in the local cache. If the local cache has the information, the resulting IP address is returned. If the local cache can not resolve the request, a DNS query is sent to the DNS servers. If an external DNS server resolves the request, the resulting IP address is stored in the local cache with its corresponding hostname.

Each lookup has a time to live value, which is defined by the DNS server, and expires from the cache automatically. In addition, the system periodically refreshes the value for FQDNs that are used in access control rules. At minimum, this refresh happens at the poll time interval (by default, every 4 hours), but it can be more frequent based on the entry's time to live value.

Use the **show dns-hosts** and **show dns** commands to check the local cache. If the IP addresses for an FQDN are wrong, you can use the **dns update** [**host** *hostname*] command to force the system to refresh the information. If you use the command without specifying a host, all hostnames are refreshed.

You can remove cached information using the **clear dns** [**host** *fqdn*] and **clear dns-hosts cache** commands.

Configuring the Management Interface

The Management interface is a virtual interface attached to the physical Management port. Note that the physical interface also includes the Diagnostic virtual interface, which you can configure on the **Interfaces** page with other physical interfaces. See Management Interface, on page 236 for more information about the Diagnostic interface.

The management interface has two uses:

- You can open web and SSH connections to the IP address and configure the device through the interface.
- The system obtains smart licensing and database updates through this IP address.

If you use the CLI setup wizard, you configure the management address and gateway for the device during initial system configuration. If you use the Firewall Device Manager setup wizard, the management address and gateway remain the defaults.

If necessary, you can change these addresses through the Firewall Device Manager. You can also change the management address and gateway in the CLI using the **configure network ipv4 manual** and **configure network ipv6 manual** commands. To restore the default management interface settings, use the **configure network** {**ipv4** | **ipv6**} **dhcp-dp-route** command.

You can define static addresses, or obtain an address through DHCP if another device on the management network is acting as a DHCP server. For most platforms, the Management interface obtains an IP address from DHCP by default.



Caution If you change the address to which you are currently connected, you will lose access to the Firewall Device Manager (or the CLI) when you save the changes, as they are applied immediately. You will need to reconnect to the device. Ensure that the new address is valid and available on the management network.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Management Interface** link.

If you are already on the **System Settings** page, click **Management Interface** in the table of contents

Step 2 Choose how you want to define the management gateway.

The gateway determines how the system can reach the internet to obtain smart licenses, database updates (such as VDB, rule, Geolocation, URL), and to reach the management DNS and NTP servers. Choose from these options:

Static IP Options:

- **Use the Data Interfaces as the Gateway**—Select this option if you do not have a separate management network connected to the Management interface. Traffic is routed to the internet based on the routing table, typically going through the outside interface. This option is not supported on the Firewall Threat Defense Virtual devices.

- **Use Unique Gateways for the Management Interface**—Specify unique gateways (below) for IPv4 and IPv6 if you have a separate management network connected to the Management interface.

DHCP IP Options:

- **Use Unique Gateways for the Management Interface with Fallback to Data Interfaces**—If the DHCP server provides a gateway, the system routes management traffic through the Management interface to the gateway. If the DHCP server does not provide a gateway, then the system routes management traffic based on the data interface routing table, typically sending traffic through the outside interface. This option is not supported on the Firewall Threat Defense Virtual devices.
- **Use Unique Gateways for the Management Interface (no Fallback)**—The system routes management traffic through the Management interface to the gateway provided by the DHCP server. If the DHCP server does not provide a gateway, the system will only be able to reach local hosts on the Management interface. To route through data interfaces, choose the Fallback option.

Step 3 Configure the management address, subnet mask or IPv6 prefix, and gateway (if necessary) for IPv4, IPv6, or both.

You must configure at least one set of properties. Leave one set blank to disable that addressing method.

Select **Type** > **DHCP** to obtain the address and gateway through DHCP or IPv6 auto configuration.

Step 4 (Optional.) If you configure a static IPv4 address, configure a DHCP server on the interface.

If you configure a DHCP server on the management interface, clients on the management network can obtain their address from the DHCP pool. This option is not supported on the Firewall Threat Defense Virtual devices.

a) Click **Enable DHCP Server** > **On**.
b) Enter the **Address Pool** for the server.

The address pool is the range of IP addresses from lowest to highest that the server is allowed to provide to clients that request an address. The range of IP addresses must be on the same subnet as the management address and cannot include: the IP address of the interface itself, the broadcast address, or the subnet network address. Specify the start and end address for the pool, separated by a hyphen. For example, 192.168.45.46-192.168.45.254.

Step 5 Click **Save**, read the warning, and click **OK**.

Configuring the Device Hostname

You can change the device hostname.

You can also change the hostname in the CLI using the **configure network hostname** command.



Caution If you change the hostname when connected to the system using the hostname, you will lose access to the Firewall Device Manager when you save the changes, as they are applied immediately. You will need to reconnect to the device.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Hostname** link.

If you are already on the System Settings page, simply click **Hostname** in the table of contents

Step 2 Enter a new hostname.

Step 3 Click **Save**.

The hostname change is immediately applied for some system processes. However, you must deploy changes to complete the update so that the same name is used by all system processes.

Configuring Network Time Protocol (NTP)

You must configure Network Time Protocol (NTP) servers to define the time on the system. You configure NTP servers during initial system setup, but you can change them using the following procedure. If you have problems with the NTP connection, see Troubleshooting NTP, on page 821.

The Firewall Threat Defense device supports NTPv4.

Note For the Firepower 4100/9300, you do not set NTP through the Firewall Device Manager. Configure NTP in FXOS.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Time Services** link.

If you are already on the System Settings page, simply click **Time Services** in the table of contents

Step 2 In **NTP Time Server**, select whether you want to use your own or Cisco's time servers.

- **Default NTP Servers**—If you select this option, the server list shows the server names that are used for NTP.
- **User-Defined NTP Servers**—If you select this option, enter the fully qualified domain name or IPv4 or IPv6 address of the NTP server you want to use. For example, ntp1.example.com or 10.100.10.10. You can add up to 3 NTP servers.

Step 3 Click **Save**.

Configuring Precision Time Protocol (ISA 3000)

The Precision Time Protocol (PTP) is a time-synchronization protocol developed to synchronize the clocks of various devices in a packet-based network. These device clocks are generally of varying precision and stability. The protocol is designed specifically for industrial, networked measurement and control systems, and is optimal for use in distributed systems because it requires minimal bandwidth and little processing overhead.

A PTP system is a distributed, networked system consisting of a combination of PTP and non-PTP devices. PTP devices include ordinary clocks, boundary clocks and transparent clocks. Non-PTP devices include network switches, routers and other infrastructure devices.

You can configure the Firewall Threat Defense device to be a transparent clock. The Firewall Threat Defense device does not synchronize its clock with the PTP clocks. The Firewall Threat Defense device will use the PTP default profile, as defined on the PTP clocks.

When you configure the PTP devices, you define a domain number for the devices that are meant to function together. Thus, you can configure multiple PTP domains, and then configure each non-PTP device to use the PTP clocks for one specific domain.

Before you begin

Determine the domain number configured on the PTP clocks that the device should use. Also, determine the interfaces through which the system can reach the PTP clocks in the domain.

Following are guidelines for configuring PTP:

- This feature is only available on the Cisco ISA 3000 appliance.
- Cisco PTP supports multicast PTP messages only.
- PTP is available only for IPv4 networks, not for IPv6 networks.
- PTP configuration is supported on physical Ethernet data interfaces, whether routed or bridge group members. It is not supported on the management interface, subinterfaces, EtherChannels, Bridge Virtual Interfaces (BVI), or any other virtual interfaces.
- PTP flows on VLAN subinterfaces are supported, assuming the appropriate PTP configuration is present on the parent interface.
- You must ensure that PTP packets are allowed to flow through the device. PTP traffic is identified by UDP destination ports 319 and 320, and destination IP address 224.0.1.129, so any access control rule that allows this traffic should work.
- When PTP packets flow between routed interfaces, you must enable multicast routing and each interface should join the 224.0.1.129 IGMP multicast group. When PTP packets flow between interfaces in the same bridge group, you do not need to enable multicast routing and configure the IGMP group.

Procedure

Step 1 Verify the configuration of the PTP clock-facing interfaces.

The default configuration places all interfaces in the same bridge group, but you can remove interfaces from the bridge group. It is important that you determine whether the interfaces are routed, or bridge group members, because you must configure them differently with respect to multicast IGMP groups.

The following procedure explains how to determine which interfaces are part of the bridge group. Check whether the interfaces you are configuring for PTP are bridge group members.

a) Click **View All Interfaces** in **Device** > **Interfaces**.
b) Find the interfaces in the list, and check the Mode column. BridgeGroupMember means it is part of a bridge group, otherwise it should be Routed.

Step 2 Click **Device**, then click the **System Settings** > **Time Services** link.

If you are already on the **System Settings** page, simply click **Time Services** in the table of contents

Step 3 Configure the PTP settings:

- **Domain Number**—The domain number that is configured on the PTP devices in your network, from 0-255. Packets received on a different domain are treated like regular multicast packets and will not undergo any PTP processing.
- **Clock Mode**—Select **EndToEndTransparent**. You can operate the device as a PTP transparent clock only.

  You can alternatively select **Forward**, but this is essentially the same as not configuring PTP. The domain number is ignored. PTP packets pass through the device based on the routing table for multicast traffic. This is the default PTP configuration.
- **Interfaces**—Select all of the interfaces through which the system can connect to the PTP clock in your network. PTP is enabled on these interfaces only.

Step 4 Click **Save**.

Step 5 If any of the interfaces you selected are routed interfaces, that is, they are not bridge group members, you need to use FlexConfig to enable multicast routing and to join the routed interfaces to the correct IGMP group.

Do not complete this step if all selected interfaces are bridge group members. You will get a deployment failure if you try to configure IGMP on a bridge group member.

a) Click **View Configuration** in **Device** > **Advanced Configuration**.
b) Click **FlexConfig** > **FlexConfig Objects** in the Advanced Configuration table of contents.
c) Create the object needed to enable multicast routing and to configure the IGMP join for the routed interfaces.

   Following would be the base template for the object. In this example, GigabitEthernet1/2 is the one routed interface on which you are enabling PTP. Change the interface hardware name as appropriate, and if you have more than one routed interface, repeat the **interface** and **igmp** commands for each additional interface.

   The **igmp** command joins the 224.0.1.129 IGMP group. This is the correct IP address for all interfaces regardless of network address.

```
multicast-routing
interface GigabitEthernet1/2
 igmp join-group 224.0.1.129
```

   The negate template would look like the following:

```
no multicast-routing
```

```
interface GigabitEthernet1/2
 no igmp join-group 224.0.1.129
```

d) Click **FlexConfig Policy** in the table of contents, add this object to the FlexConfig policy, and click **Save**.

Verify that the preview shows the expected commands from your object.

What to do next

After you deploy changes, you can verify the PTP settings. From the Firewall Device Manager CLI Console, or an SSH or Console session, issue the various **show ptp** commands. For example, if you configured PTP for domain 10 for GigabitEthernet1/2 only, the output might look like the following:

```
> show ptp clock
 PTP CLOCK INFO
  PTP Device Type: End to End Transparent Clock
  Operation mode: One Step
  Clock Identity: 34:62:88:FF:FE:1:73:81
  Clock Domain: 10
  Number of PTP ports: 4
> show ptp port
 PTP PORT DATASET: GigabitEthernet1/1
  Port identity: Clock Identity: 34:62:88:FF:FE:1:73:81
  Port identity: Port Number: 1
  PTP version: 2
  Port state: Disabled

 PTP PORT DATASET: GigabitEthernet1/2
  Port identity: Clock Identity: 34:62:88:FF:FE:1:73:81
  Port identity: Port Number: 2
  PTP version: 2
  Port state: Enabled

 PTP PORT DATASET: GigabitEthernet1/3
  Port identity: Clock Identity: 34:62:88:FF:FE:1:73:81
  Port identity: Port Number: 3
  PTP version: 2
  Port state: Disabled

 PTP PORT DATASET: GigabitEthernet1/4
  Port identity: Clock Identity: 34:62:88:FF:FE:1:73:81
  Port identity: Port Number: 4
  PTP version: 2
  Port state: Disabled
```

Configuring HTTP Proxy for Management Connections

If there is not a direct connection between the system and the Internet, you can set up an HTTP proxy for the management interface. The system will then use the proxy for all management connections, including connections to the Firewall Device Manager and from the system to Cisco for downloading database updates.

You can also configure an HTTP proxy in the Firewall Threat Defense CLI using the **configure network http-proxy** command.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **HTTP Proxy** link.

If you are already on the **System Settings** page, simply click **HTTP Proxy** in the table of contents

Step 2 Click the toggle to enable the proxy, then configure the proxy settings:

- **HTTP Proxy**—The IP address of the proxy server.
- **Port**—The port number the proxy server is configured to listen to for HTTP connections.
- **Use Proxy Authentication**—Select this option if the server is configured to require authentication for proxied connections. If you select this option, also enter the **Username** and **Password** of an account that can log into the proxy server.

Step 3 Click **Save**, then confirm that you want to make the change.

Your changes are applied immediately. A deployment job is not needed.

Because you are changing how the system completes management connections, you will lose your connection to the Firewall Device Manager. Wait a few minutes for the change to be complete, then refresh your browser window and log in again.

Configuring Cloud Services

You can enroll in Cloud Services so that you can use various cloud-based applications, such as Security Cloud Control, Cisco Threat Response, and the Cisco Success Network.

Once registered in the cloud, the page will show registration status and the type of tenancy, and the account name under which the device is registered.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Cloud Services** link.

If you are already on the **System Settings** page, simply click **Cloud Services** in the table of contents.

If your device is not registered, this page shows enrollment methods for registering with the Cisco cloud. After you register with the cloud, you will be able to enable or disable individual cloud services.

Step 2 To register with the Cisco cloud (in evaluation mode or after unregistering from Cloud Services), select one of the following options:

- **Security**/Security Cloud Control **Account**—You can use one of the following methods:
  - **Auto-enroll with Tenancy from Security Cloud Control** (Firepower 1000, 2100, Secure Firewall 3100 only). You can use auto-enrollment instead of obtaining a registration key. First, go to Security Cloud Control and add the device using the device's serial number. Then, in the Firewall Device

Manager, select this check box, and initiate enrollment. Get the serial number from the device chassis or packing slip. For FXOS, you can go into the FXOS CLI and use the **show chassis detail** command to retrieve the correct serial number, labeled Serial (SN). Note that the Firewall Threat Defense command **show serial-number** provides a different serial number, which is not recommended for Security Cloud Control registration.

- Log into your Security Cloud Control or other security account and generate a registration key. Then return to this page, select your **Cloud Services Region**, and paste in your **Registration Key**.

  You can also at this time enable **Security Cloud Control** and **Cisco Success Network**. These are enabled by default.

- **Smart License**—(Only if you will not use Security Cloud Control.) Click the link to go to the Smart Licensing page and register with CSSM. Registering with CSSM also registers the device with Cloud Services.

  Note
  If you unregistered from Cloud Services, then the Smart License approach to registration has some additional steps. In this case, select your **Cloud Services Region**, then click **Register**. Read the disclosure and click **Accept**.

Step 3 Once you have registered for Cloud Services, you can enable or disable features as needed. See the following topics:

- Enabling or Disabling Security Cloud Control (Legacy Device Manager Mode)
- Connecting to the Cisco Success Network, on page 774
- Sending Events to the Cisco Cloud, on page 775
- Unregistering from Cloud Services, on page 776

Enabling or Disabling Security Cloud Control (Legacy Device Manager Mode)

Note This section only applies to the legacy device manager mode in Security Cloud Control, not the cloud-delivered management center.

If you enrolled in Cloud Services using a registration key from Security Cloud Control, as recommended in Configuring Cloud Services, on page 772, the device is already registered with Security Cloud Control. Afterward, you can disable or re-enable the connection as desired.

If the device is registered to Cloud Services using Smart Licensing, you will have problems if you enable Security Cloud Control: the device will not show up in the Security Cloud Control inventory. We strongly recommend that you first unregister the device from the Cloud Services; select **Unregister Cloud Services** from the gear (⚙) drop-down list. After you unregister, get a registration token from Security Cloud Control and re-register using the token and your security account as explained in Configuring Cloud Services, on page 772.

For more information about how cloud management works, refer to the Security Cloud Control portal (http://www.cisco.com/go/cdo) or ask the reseller or partner with whom you are working.

Before you begin

If you intend to configure high availability, you must register both devices that you will use in the high availability group.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Cloud Services** link.

If you are already on the System Settings page, simply click **Cloud Services** in the table of contents.

Step 2 Click the **Enable**/**Disable** button for the Security Cloud Control feature to change the setting as appropriate.

Connecting to the Cisco Success Network

When you register the device, you decide whether to enable the connection to the Cisco Success Network. See Registering the Device, on page 87.

By enabling Cisco Success Network, you are providing usage information and statistics to Cisco that are essential for Cisco to provide you with technical support. This information also allows Cisco to improve the product and to make you aware of unused available features so that you can maximize the value of the product in your network.

When you enable the connection, your device establishes a secure connection to the Cisco Cloud so that your device can participate in additional service offerings from Cisco such as technical support services, cloud management and monitoring services. Your device will establish and maintain this secure connection at all times. For information on completely disconnecting from the cloud, see Unregistering from Cloud Services, on page 776.

After you have registered the device, you can change the Cisco Success Network setting.

Note When the system sends data to Cisco, the task list shows a Telemetry Job.

Before you begin

To enable Cisco Success Network the device must be enrolled with the cloud. To enroll the device, either register the device with Cisco Smart Software Manager (on the Smart Licensing page), electing the Cisco Success Network option during registration, or enroll with Security Cloud Control by entering a registration key (legacy device manager mode in Security Cloud Control only).

Note If you enable Cisco Success Network on the active unit in a high availability group, you are also enabling the connection on the standby unit.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Cloud Services** link.

If you are already on the System Settings page, simply click **Cloud Services** in the table of contents.

Step 2 Click the **Enable**/**Disable** control for the Cisco Success Network feature to change the setting as appropriate.

You can click the **sample data** link to see the type of information that is sent to Cisco.

When enabling the connection, read the disclosure and click **Accept**.

Sending Events to the Cisco Cloud

You can send events to the Cisco cloud server. From there, various Cisco cloud services can access the events. You can then use these cloud applications to analyze the events and to evaluate threats that the device might have encountered.

The cloud tools determine whether the events you send are used. Consult the tool's documentation, or examine the event data, to ensure you are not sending unused events to the cloud, wasting both your bandwidth and storage space. Keep in mind that the tools pull the events from the same source, so your selection should reflect all the tools you use, not just the most restrictive tool. For example:

- The Security Analytics and Logging tool in Security Cloud Control can make use of all connection events.
- Threat Reponse uses high priority connection events only, so there is no need to send all connection events to the cloud. In addition, it will use only the Security Intelligence high-priority events.

Before you begin

You must enroll the device with Cloud Services before you can enable this service.

You can connect to Threat Response at https://visibility.amp.cisco.com/ in the US region, https://visibility.eu.amp.cisco.com in the EU region, and https://visibility.apjc.amp.cisco.com in the APJC region. You can watch videos about the use and benefits of the application on YouTube at http://cs.co/CTRvideos. For more information, see *Cisco Secure Firewall Threat Defense and SecureX Threat Response Integration guide*, which you can find at https://www.cisco.com/c/en/us/support/security/defense-center/products-installation-and-configuration-guides-list.html.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Cloud Services** link.

If you are already on the System Settings page, simply click **Cloud Services** in the table of contents.

Step 2 Click the **Enable**/**Disable** control for the **Send Events to the Cisco Cloud** option to change the setting as appropriate.

Step 3 When you are enabling the service, you are prompted to select the events to send to the cloud. Later, you can change these selections by clicking **Edit** next to the list of selected events. Select the types of events to send and click **OK**.

- **File/Malware**—For any file policies you have applied in any access control rule.
- **Intrusion**—For any intrusion policies you have applied in any access control rule.
- **Connection**—For access control rules where you have enabled logging. When you select this option, you can also elect to send All Connection Events, or only send the High Priority connection events. High-priority connection events are those related to connections that trigger intrusion, file, or malware events, or that match Security Intelligence blocking policies.

Unregistering from Cloud Services

If you no longer want to use any cloud services, you can unregister the device from the cloud. You might want to unregister when you are removing the device from service or otherwise disposing of it. If you need to change your cloud services region, you unregister, then select the new region when you re-register.

Unregistering from the cloud using this procedure has no impact on Smart Licensing registration.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Cloud Services** link.

If you are already on the System Settings page, simply click **Cloud Services** in the table of contents.

Step 2 Select **Unregister Cloud Services** from the gear ( ) drop-down list.

Step 3 Read the warning and click **Unregister**.

Any cloud services you had enabled will be automatically disabled, and your ability to enable them again will be removed. However, you will now see the controls for registering with the cloud, and you can re-register.

Enabling or Disabling Web Analytics

Enabling web analytics provides anonymous product usage information to Cisco based on page hits. The information includes pages viewed, the time spent on a page, browser versions, product version, device hostname, and so forth. This information can help Cisco determine feature usage patterns and help Cisco improve the product. All usage data is anonymous and no sensitive data is transmitted.

Web analytics is enabled by default.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Web Analytics** link.

If you are already on the System Settings page, simply click **Web Analytics** in the table of contents.

Step 2 Click the **Enable**/**Disable** control for the **Web Analytics** feature to change the setting as appropriate.

Configuring URL Filtering Preferences

The system obtains the URL category and reputation database from Cisco Collective Security Intelligence (CSI, Cisco Talos Intelligence Group (Talos)). These preferences control database updates and how the system handles URLs with unknown category or reputation. You must enable the URL filtering license to set these preferences.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **URL Filtering Preferences** link.

If you are already on the System Settings page, simply click **URL Filtering Preferences** in the table of contents

Step 2 Configure the following options:

- **Enable Automatic Updates**—Allows the system to automatically check for and download updated URL data, which includes category and reputation information. The system checks for updates every 30 minutes, although the data is typically updated once per day. The default is to enable updates. If you deselect this option, and you are using category and reputation filtering, periodically enable it to get new URL data.
- **Query Cisco CSI for Unknown URLs**—Whether to check with Cisco CSI for updated information for URLs that do not have category and reputation data in the local URL filtering database. If the lookup returns this information within a reasonable time limit, it is used when selecting access rules based on URL conditions. Otherwise, the URL matches the Uncategorized category. Selecting this option is important for lower-end systems, which install a smaller URL database due to memory limitations.
- **URL Time to Live** (available if you select **Query Cisco CSI for Unknown URLs**)—How long to cache the category and reputation lookup values for a given URL. When the time to live expires, the next attempted access of the URL results in a fresh category/reputation lookup. A shorter time results in more accurate URL filtering, a longer time results in better performance for unknown URLs. You can set the TTL to 2, 4, 8, 12, 24, or 48 hours, one week, or Never (the default).

Step 3 As needed, you can **Check the Category for a URL**.

You can check on the category and reputation for a particular URL. Enter the URL in the **URL to Check** box and click **Go**. You will be taken to an external website to see the results. If you disagree with a categorization, click the **Submit a URL Category Dispute** link and let us know.

Step 4 Click **Save**.

Switch from the Firewall Device Manager to the Firewall Management Center

You can configure the Firewall Threat Defense device to connect to the Firewall Management Center for management if you want to switch from the Firewall Device Manager.

When you perform the Firewall Management Center setup using the Firewall Device Manager, *all* interface configuration completed in the Firewall Device Manager is retained when you switch to the Firewall Management Center for management, in addition to the Management interface and the manager access settings. Note that other default configuration settings, such as the access control policy or security zones, are not retained. When you use the Firewall Threat Defense CLI for initial setup for the Firewall Management Center, only the Management interface and the manager access settings are retained (for example, the default inside interface configuration is not retained).

After you switch to the Firewall Management Center, you can no longer use the Firewall Device Manager to manage the Firewall Threat Defense device.

Before you begin

If the firewall is configured for high availability, you must first break the high availability configuration using the Firewall Device Manager (if possible) or the **configure high-availability disable** command. Ideally, break high availability from the active unit.

Procedure

Step 1 If you registered the firewall to the Cisco Smart Software Manager, you must unregister before switching managers. See Unregistering the Device, on page 90.

Unregistering the firewall frees the base license and all feature licenses. If you do not unregister the firewall, those licenses remain assigned to the firewall in Cisco Smart Software Manager.

Step 2 (Might be required) Configure the Management interface. See Configuring the Management Interface, on page 766.

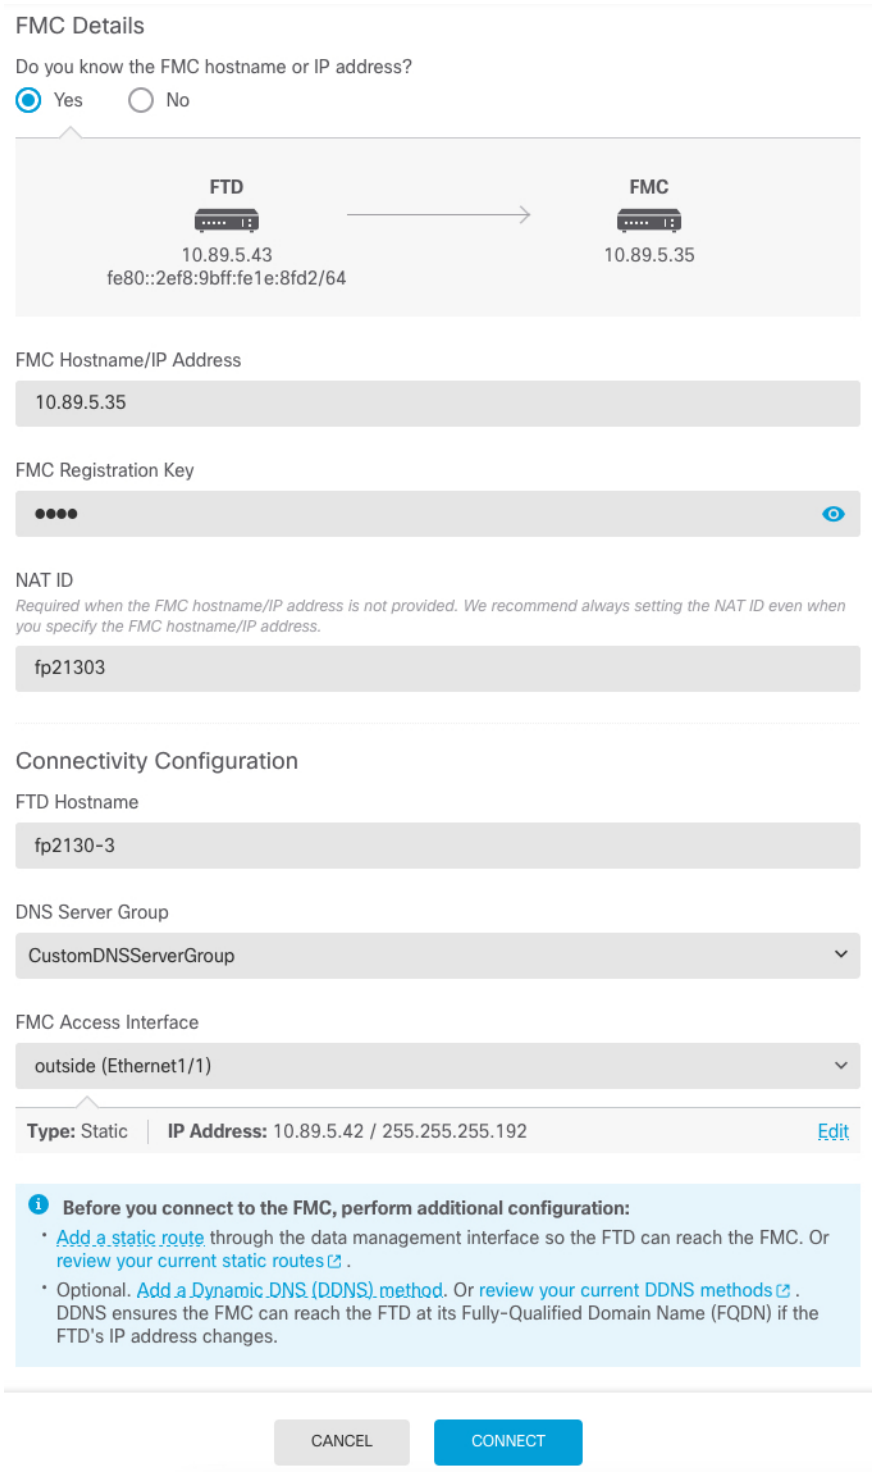
You may need to change the Management interface configuration, even if you intend to use a data interface for manager access. You will have to reconnect to the Firewall Device Manager if you were using the Management interface for the Firewall Device Manager connection.

- Data interface for manager access—The Management interface must have the gateway set to data interfaces. By default, the Management interface receives an IP address and gateway from DHCP. If you do not receive a gateway from DHCP (for example, you did not connect this interface to a network), then the gateway will default to data interfaces, and you do not need to configure anything. If you did receive a gateway from DHCP, then you need to instead configure this interface with a static IP address and set the gateway to data interfaces.
- Management interface for manager access—If you want to configure a static IP address, be sure to also set the default gateway to be a unique gateway instead of the data interfaces. If you use DHCP, then you do not need to configure anything assuming you successfully get the gateway from DHCP.

Step 3 Choose **Device** > **System Settings** > **Management Center**, and click **Proceed** to set up the Firewall Management Center management.

Step 4 Configure the **FMC Details**.

Figure 45: FMC Details



a) For **Do you know the FMC hostname or IP address?**, click **Yes** if you can reach the Firewall Management Center using an IP address or hostname, or **No** if the Firewall Management CenterSecurity Cloud Control is behind NAT or does not have a public IP address or hostname.

At least one of the devices, either the Firewall Management Center or the Firewall Threat Defense device, must have a reachable IP address to establish the two-way, TLS-1.3-encrypted communication channel between the two devices.

b) If you chose **Yes**, then enter the **FMC Hostname/IP Address**.

c) Specify the **FMC Registration Key**.

This key is a one-time registration key of your choice that you will also specify on the Firewall Management Center when you register the Firewall Threat Defense device. The registration key must be between 2 and 36 characters. Valid characters include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-). This ID can be used for multiple devices registering to the Firewall Management Center.

a) Specify a **NAT ID**.

This ID is a unique, one-time string of your choice that you will also specify on the Firewall Management Center. The NAT ID must be between 2 and 36 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 Firewall Management Center. The NAT ID is used in combination with the IP address to verify that the connection is coming from the correct device; only after authentication of the IP address/NAT ID will the registration key be checked. We recommended that you always use the NAT ID even when it is optional, but it is required if:

- You set the Firewall Management Center IP address to **DONTRESOLVE**.
- When adding the device on the Firewall Management Center, you do not specify a reachable device IP address or hostname.
- You use the data interface for management, even if you specify IP addresses on both sides.
- The Firewall Management Center uses multiple management interfaces.

Step 5 Configure the **Connectivity Configuration**.

a) Specify the **FTD Hostname**.

If you use a data interface for the **FMC Access Interface** access, then this FQDN will be used for this interface.

b) Specify the **DNS Server Group**.

Choose an existing group, or create a new one. The default DNS group is called **CiscoUmbrellaDNSServerGroup**, which includes the OpenDNS servers.

If you intend to choose a data interface for the **FMC Access Interface**, then this setting sets the *data* interface DNS server. The Management DNS server that you set with the setup wizard is used for management traffic. The data DNS server is used for DDNS (if configured) or for security policies applied to this interface. You are likely to choose the same DNS server group that you used for Management, because both management and data traffic reach the DNS server through the outside interface.

On the Firewall Management Center, the data interface DNS servers are configured in the Platform Settings policy that you assign to this Firewall Threat Defense device. When you add the Firewall Threat Defense device to the Firewall Management Center, the local setting is maintained, and the DNS servers are *not* added to a Platform Settings policy. However, if you later assign a Platform Settings policy to the Firewall Threat Defense device that includes a DNS configuration, then that configuration will overwrite the local

setting. We suggest that you actively configure the DNS Platform Settings to match this setting to bring the Firewall Management Center and the Firewall Threat Defense device into sync.

Also, local DNS servers are only retained by the Firewall Management Center if the DNS servers were discovered at initial registration.

If you intend to choose the Management interface for the **FMC Access Interface**, then this setting configures the Management DNS server.

c) For the **FMC Access Interface**, choose any configured interface.

You can change the manager interface after you register the Firewall Threat Defense device to the Firewall Management Center, to either the Management interface or another data interface.

Step 6 (Optional) If you chose a data interface, and it was not the outside interface, then add a default route.

You will see a message telling you to check that you have a default route through the interface. If you chose outside, you already configured this route as part of the setup wizard. If you chose a different interface, then you need to manually configure a default route before you connect to the Firewall Management Center. See Configuring Static Routes, on page 312 for more information about configuring static routes.

If you chose the Management interface, then you need to configure the gateway to be a unique gateway before you can proceed on this screen. See Configuring the Management Interface, on page 766.

Step 7 (Optional) If you chose a data interface, click **Add a Dynamic DNS (DDNS) method**.

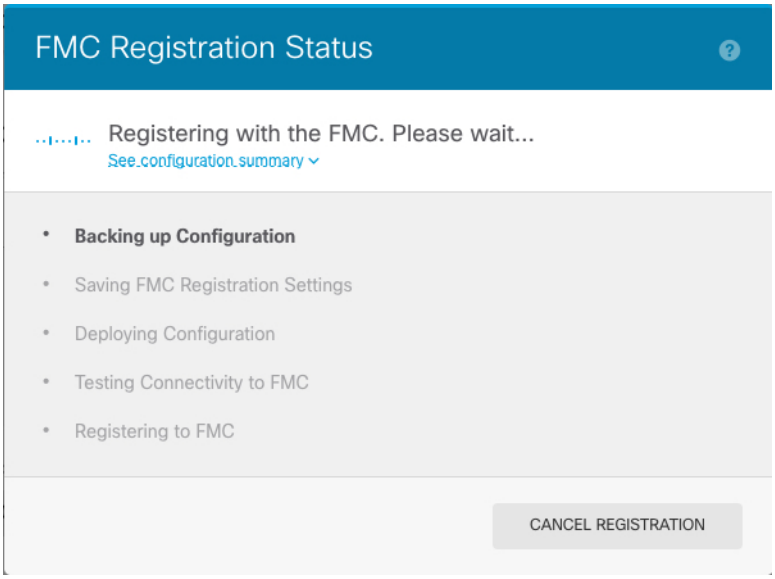
DDNS ensures the Firewall Management Center can reach the Firewall Threat Defense device at its Fully-Qualified Domain Name (FQDN) if the IP address changes. See **Device** > **System Settings** > **DDNS Service** to configure DDNS.

If you configure DDNS before you add the Firewall Threat Defense device to the Firewall Management Center, the Firewall Threat Defense device automatically adds certificates for all of the major CAs from the Cisco Trusted Root CA bundle so that the Firewall Threat Defense device can validate the DDNS server certificate for the HTTPS connection. Firewall Threat Defense supports any DDNS server that uses the DynDNS Remote API specification (https://help.dyn.com/remote-access-api/).

DDNS is not supported when using the Management interface for manager access.

Step 8 Click **Connect**. The **FMC Registration Status** dialog box shows the current status of the switch to the Firewall Management Center. After the **Saving FMC Registration Settings** step, go to the Firewall Management Center, and add the firewall.

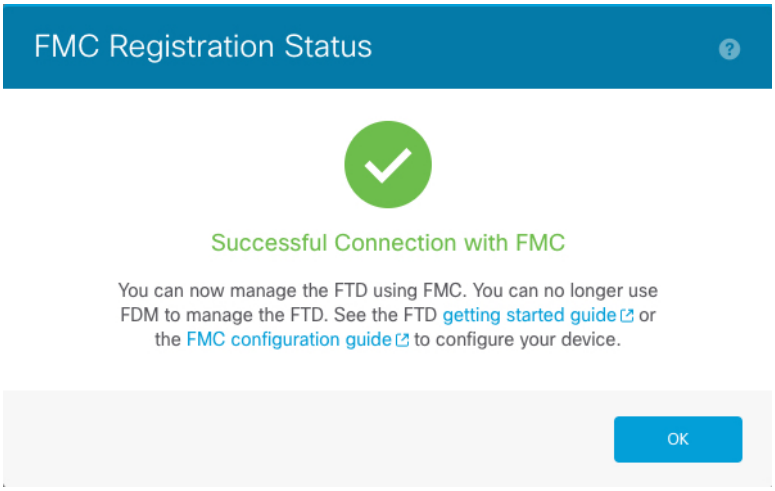
Figure 46: FMC Registration Status



If you want to cancel the switch to the Firewall Management Center, click **Cancel Registration**. Otherwise, do not close the Firewall Device Manager browser window until after the **Saving FMC Registration Settings** step. If you do, the process will be paused, and will only resume when you reconnect to the Firewall Device Manager.

If you remain connected to the Firewall Device Manager after the **Saving FMC Registration Settings** step, you will eventually see the **Successful Connection with FMC** dialog box, after which you will be disconnected from the Firewall Device Manager.

Figure 47: Successful Connection with FMC



Switch from the Firewall Management Center to the Firewall Device Manager

You can configure the Firewall Threat Defense device currently being managed by the Firewall Management Center to use the Firewall Device Manager instead.

You can switch from the Firewall Management Center to the Firewall Device Manager without reinstalling the software. Before switching from the Firewall Management Center to the Firewall Device Manager, verify that the Firewall Device Manager meets all of your configuration requirements. If you want to switch from the Firewall Device Manager to the Firewall Management Center, see Switch from the Firewall Device Manager to the Firewall Management Center, on page 778.



Caution Switching to the Firewall Device Manager erases the device configuration and returns the system to the default configuration. However, the Management IP address and hostname are preserved.

Procedure

Step 1 In the Firewall Management Center, delete the firewall from the **Devices** > **Device Management** page.

Step 2 Connect to the Firewall Threat Defense CLI using SSH or the console port. For SSH, open a connection to the **management IP address**, and log into the Firewall Threat Defense CLI with the **admin** username (or any other user with admin privileges).

The console port defaults to the FXOS CLI. Connect to the Firewall Threat Defense CLI using the **connect ftd** command. The SSH session connects directly to the Firewall Threat Defense CLI.

If you cannot connect to the management IP address, do one of the following:

- Ensure that the Management physical port is wired to a functioning network.
- Ensure that the management IP address and gateway are configured for the management network. Use the **configure network ipv4/ipv6 manual** command.

Step 3 Verify you are currently in remote management mode.

show managers

Example:

```
> show managers
Type                      : Manager
Host                      : 10.89.5.35
Registration              : Completed
```

Step 4 Delete the remote manager and go into no manager mode.

configure manager delete

You cannot go directly from remote management to local management.

Example:

```
> configure manager delete
Deleting task list
Manager successfully deleted.

>
> show managers
No managers configured.
```

Step 5 Configure the local manager.

configure manager local

You can now use a web browser to open the local manager at **https://***management-IP-address*.

Example:

```
> configure manager local
Deleting task list

> show managers
Managed locally.
```

Configuring TLS/SSL Cipher Settings

The SSL cipher settings control which TLS versions and encryption cipher suites are allowed for TLS/SSL connections to the device.

Normally, the cipher suite you configure should have more than one available encryption cipher suite. The system will determine the highest TLS version that both the client and Firewall Threat Defense device support, then pick a cipher suite that both support that is compatible with the TLS version. The system will select the strongest TLS version and cipher suite supported by both endpoints to ensure the most secure connection possible among the ciphers you allow.

Before you begin

By default, the system uses the DefaultSSLCipher object to define the allowed cipher suites. You can instead choose one of the other pre-defined cipher objects, or create your own custom object. Ideally, create a single object that includes all and only the TLS versions and ciphers that you want to allow. To create your own object, see Configure TLS/SSL Cipher Objects, on page 785.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **SSL Settings** link.

Step 2 Configure the **SSL Settings**:

These settings control the ciphers clients are allowed to use when establishing remote access VPN connections.

- **Ciphers**—Select the SSL Cipher objects that define the TLS versions and encryption algorithms that are allowed.

  Click **Create New Cipher** at the bottom of the list if you need to create an object now.

- **Ephemeral Diffie-Hellman Group**—The DH group to use for ephemeral encryption algorithms. For an explanation of the DH groups, see Decide which Diffie-Hellman modulus group to use, on page 632. The default is 14.

- **Elliptical Curve DH Group**—The DH group to use for elliptical curve encryption algorithms. The default is 19.

Step 3 Click **Save**.

Configure TLS/SSL Cipher Objects

SSL Cipher objects define a combination of security level, TLS/DTLS protocol versions, and encryption algorithms that can be used when establishing an SSL connection to an Firewall Threat Defense device. Use these objects in **Device** > **System Settings** > **SSL Settings** to define the security requirements for users who make SSL connections to the box.

The TLS versions and ciphers that you can select is controlled by your Smart License account. If you satisfy export compliance requirements, you can select any combination of options. If your license is not export compliant, you are limited to TLSv1.0 and DES-CDC-SHA, which are the lowest security options. Evaluation mode is considered a non-compliant mode, so your options are limited until you license the system.

The system includes several pre-defined objects. You need to create new objects only if the pre-defined objects do not fit your security requirements. The objects are:

- DefaultSSLCipher—This is a low-security level group. It is the default used in the SSL settings, to ensure that as many clients as possible can complete connections to the system. It includes all protocol versions and ciphers supported by the system.
- CiscoRecommendedCipher—This is a high-security level group, which includes only the most secure ciphers and TLS version. This group provides the highest security, but you need to ensure that your clients can use the matching ciphers. There is a greater likelihood that some clients will be unable to complete connections due to cipher miss-match problems.

Procedure

Step 1 Select **Objects**, then select **SSL Ciphers** from the table of contents.

Step 2 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon () for the object.

To delete an unreferenced object, click the trash can icon () for the object.

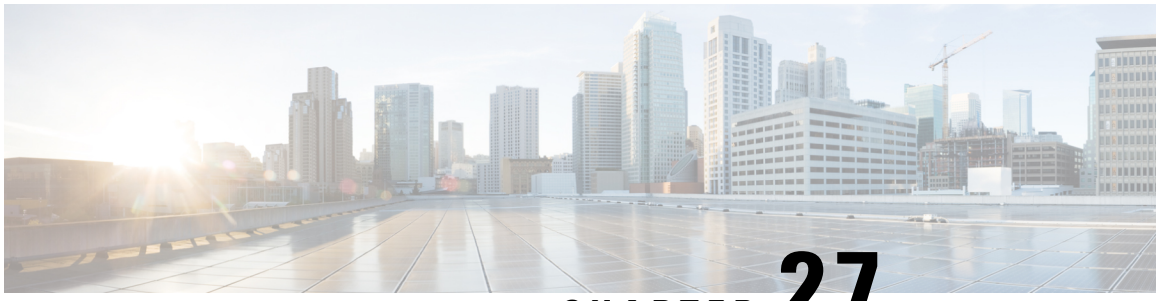
Step 3 Enter a **Name** for the object and optionally, a description.

Step 4 Configure the following options:

- **Security Level**—The relative security level for the object. Note that if you edit the protocol versions or cipher suite list after selecting a security level, the actual level of security provided by the object might not match the security level. Choose one of the following:
  - **All**—Include all TLS levels and cipher suites in the object, from low to high security.
  - **Low**—Includes all the TLS versions and ciphers, which allows users to complete connections with the least secure ciphers. For a non-export compliant license, this includes TLSv1.0 and DES-CBC-SHA.
  - **Medium**—Includes all of the TLS versions, but removes some relatively insecure ciphers. There is only a minimal difference between this option and the Low/All option. You cannot use this option with non-export compliant licenses.
  - **High**—Allows the latest DTLS and TLS versions only, and the ciphers that work with these versions. This option limits connections to the currently most-secure ciphers available. You cannot use this option with non-export compliant licenses.
  - **Custom**—Select this option if you want to select TLS versions and ciphers individually. The options you select will dictate whether you are defining a high or low security encryption setting. Although there are no defaults for a custom object, if you selected another level before selecting custom, the previously displayed options remain selected for your convenience.
- **Protocol Versions**—The TLS/DTLS versions that a client is allowed to use when establishing a TLS/SSL connection to the Firewall Threat Defense device. For a custom object, select the versions you want to support. For other security levels, ideally you should not edit the list, but you can add or remove versions as desired.
- **Applicable Cipher Suites**—The encryption algorithms that the client can use. Click **+** to add new suites; click **x** on a suite to remove it.

  Your selection of protocol version controls which suites are available in this list. If you change the protocol versions, any selected suite that no longer works with the selected versions is flagged: you must remove these, or add back the required protocol version.

Step 5 Click **OK**.



CHAPTER 27

System Management

The following topics explain how to perform system management tasks such as updating system databases and backing up and restoring the system.

- Installing Software Updates, on page 787
- Backing Up and Restoring the System, on page 796
- Auditing and Change Management, on page 802
- Exporting the Device Configuration, on page 808
- Managing Firewall Device Manager and Firewall Threat Defense User Access, on page 808
- Rebooting or Shutting Down the System, on page 815
- Troubleshooting the System, on page 816
- Hardware Management Tasks, on page 827

Installing Software Updates

You can install updates to the system databases and to the system software. The following topics explain how to install these updates.

Updating System Databases and Feeds

The system uses several databases and Security Intelligence feeds to provide advanced services. Cisco provides updates to these databases and feeds so that your security policies use the latest information available.

Overview of System Database and Feed Updates

Firewall Threat Defense uses the following databases and feeds to provide advanced services.

Intrusion rules

As new vulnerabilities become known, the Cisco Talos Intelligence Group (Talos) releases intrusion rule updates that you can import. These updates affect intrusion rules, preprocessor rules, and the policies that use the rules.

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.

For changes made by an intrusion rule update to take effect, you must redeploy the configuration.

Intrusion rule updates may be large, so import rules during periods of low network use. On slow networks, an update attempt might fail, and you will need to retry.

Geolocation database (GeoDB)

The Cisco Geolocation Database (GeoDB) is a database of geographical data (such as country, city, coordinates) associated with routable IP addresses.

GeoDB updates provide updated information on physical locations that your system can associate with detected routable IP addresses. You can use geolocation data as a condition in access control rules.

The time needed to update the GeoDB depends on your appliance; the installation usually takes 30 to 40 minutes. 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.

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 firewall system correlates the fingerprints with the vulnerabilities to help you determine whether a particular host increases your risk of network compromise. The Cisco Talos Intelligence Group (Talos) issues periodic updates to the VDB.

The time it takes to update vulnerability mappings depends on the number of hosts in your network map. You may want to schedule the update during low system usage times to minimize the impact of any system downtime. As a rule of thumb, divide the number of hosts on your network by 1000 to determine the approximate number of minutes to perform the update.

After you update the VDB, you must redeploy configurations before updated application detectors and operating system fingerprints can take effect.

Cisco Talos Intelligence Group (Talos) Security Intelligence Feeds

Talos provides access to regularly updated intelligence feeds for use in Security Intelligence policies. Sites representing security threats such as malware, spam, botnets, and phishing appear and disappear faster than you can update and deploy custom configurations. These feeds contain addresses and URLs for known threats. When the system updates a feed, you do not have to redeploy. The new lists are used for evaluating subsequent connections.

URL Category/Reputation Database

The system obtains the URL category and reputation database from Cisco Collective Security Intelligence (CSI). If you configure URL filtering access control rules that filter on category and reputation, requested URLs are matched against the database. You can configure database updates and some other URL filtering preferences on **System Settings** > **URL Filtering Preferences**. You cannot manage URL category/reputation database updates the same way you manage updates for the other system databases.

Updating System Databases

You can manually retrieve and apply system database updates at your convenience. Updates are retrieved from the Cisco support site. Thus, there must be a path to the internet from the system's management address.

Alternatively, you can retrieve the update packages from the internet yourself, then upload them from your workstation. This method is primarily meant for air-gapped networks, where there is no path to the internet for retrieving the updates from Cisco. Download the updates from software.cisco.com from the same folders where you would download system software upgrades.

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 additional contextual data associated with routable IP addresses. The Firewall Device Manager does not and has never used the information in the IP package. This split saves significant disk space in locally managed Firewall Threat Defense deployments. If you are getting the GeoDB from Cisco yourself, make sure you get the country code package, which has the same file name as the old all-in-one package: Cisco_GEODB_Update-*date-build*.

You can also set up a regular schedule to retrieve and apply database updates. Because these updates can be large, schedule them for times of low network activity.

Note While a database update is in progress, you might find that the user interface is sluggish to respond to your actions.

Before you begin

To avoid any potential impact to pending changes, deploy the configuration to the device before manually updating these databases.

Please be aware that VDB and URL category updates can remove applications or categories. You need to update any access control or SSL decryption rules that use these deprecated items before you can deploy changes.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Updates summary.

This opens the Updates page. Information on the page shows the current version for each database and the last date and time each database was updated.

Step 2 To manually update a database, click one of the following options in the section for that database:

- **Update from Cloud**—To have the Firewall Device Manager retrieve the update package from Cisco. This is the easiest and most reliable method, but there must be a path to the internet to use it.
- **(down arrow)** > ***option***—To select the update package from your workstation or a drive connected to your workstation. The option will be one of the following:
  - **Select File**—Select a VDB or Geolocation package. You must select a newer version than the one installed. You cannot downgrade these packages.
  - **Update to Newer Version**—Select an Intrusion Rule package that is newer than the one that is currently installed.
  - **Downgrade to Older Version**—Select an Intrusion Rule package that is older than the one that is currently installed.

Rule and VDB updates require a configuration deployment to make them active. When you update from the cloud, you are asked whether you want to deploy now; click **Yes**. If you click **No**, remember to initiate a deployment job at your earliest convenience.

If you upload your own file, you must always deploy the changes manually.

Note
When manually uploading an intrusion rule package, make sure you upload the right package type for your Snort version, SRU for Snort 2, LSP for Snort 3. You can upload a package for the non-active Snort version, but it will not be activated unless you switch versions. For information on switching Snort versions, see Switching Between Snort 2 and Snort 3, on page 516.

Step 3 (Optional) To set up a regular database update schedule:

a) Click the **Configure** link in the section for the desired database. If there is already a schedule, click **Edit**.

The update schedules for the databases are separate. You must define the schedules separately.

b) Set the update start time:

- The frequency of the update (Daily, Weekly, or Monthly).
- For weekly or monthly, the days of the week or month you want the update to occur.
- The time you want the update to start. The time you specify is adjusted for Daylight Savings Time, so it will move an hour forward or backward whenever the time is adjusted in your area. You must edit the schedule at the time change if you want to keep this exact time throughout the year.

c) For Rule or VDB updates, select the **Automatically Deploy the Update** checkbox if you want the system to deploy the configuration whenever the database is updated.

The update is not effective until it is deployed. The automatic deployment also deploys any other configuration changes that are not yet deployed.

d) Click **Save**.

Note
If you want to remove a recurring schedule, click the **Edit** link to open the scheduling dialog box, then click the **Remove** button.

Updating Cisco Security Intelligence Feeds

Cisco Talos Intelligence Group (Talos) provides access to regularly updated Security 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. When the system updates a feed, you do not have to redeploy. The new lists are used for evaluating subsequent connections.

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.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Updates summary.

This opens the Updates page. Information on the page shows the current version for the **Security Intelligence Feeds** and the last date and time the feeds were updated.

Step 2 To manually update the feeds, click **Update Now** in the **Security Intelligence Feeds** group.

If you manually update the feeds on one unit in a high availability group, you need to also manually update them on the other unit to ensure consistency.

Step 3 (Optional.) To configure a regular update frequency:

a) Click the **Configure** link in the section for Cisco Feeds. If there is already a schedule, click **Edit**.

b) Select the desired frequency.

The default is **Hourly**. You can also set a **Daily** update (specify the time of day) or **Weekly** update (select the days of the week and time of day). The time you specify is adjusted for Daylight Savings Time, so it will move an hour forward or backward whenever the time is adjusted in your area. You must edit the schedule at the time change if you want to keep this exact time throughout the year.

Click **Delete** to prevent automatic updates.

c) Click **OK**.

Upgrading Firewall Threat Defense

Use this procedure to upgrade a standalone Firewall Threat Defense device. If you need to update FXOS, do that first. To upgrade high availability threat defense, see Upgrading High Availability Firewall Threat Defense, on page 221.

Caution Traffic is dropped while you upgrade. Even if the system appears inactive or unresponsive, do not manually reboot or shut down during upgrade; you could place the system in an unusable state and require a reimage. You can manually cancel failed or in-progress major and maintenance upgrades, and retry failed upgrades. If you continue to have issues, contact Cisco TAC.

For details on these and other issues you may encounter during upgrade, see Troubleshooting Threat Defense Upgrades, on page 795.

Before you begin

Complete upgrade planning. Make sure your deployment is healthy and successfully communicating.

Note Upgrade planning starts with reading the Secure Firewall Threat Defense release notes. It then includes taking backups, obtaining upgrade packages, and performing associated upgrades (such as FXOS for the Firepower 4100/9300). It also includes checks for necessary configuration changes, readiness checks, disk space checks, and checks for both running and scheduled tasks. For details, see the Secure Firewall Threat Defense upgrade guides for Firewall Device Manager for your version.

Procedure

Step 1 Select **Device**, then click **View Configuration** in the Updates panel.
The System Upgrade panel shows the currently running software version and any upgrade package that you have already uploaded.

Step 2 Upload the upgrade package.

You can upload one package only. If you upload a new package, it replaces the old one. Make sure you have the correct package for your target version and device model. Click **Browse** or **Replace File** to begin the upload.

When the upload completes, the system displays a confirmation dialog box. Before you click **OK**, optionally select **Run Upgrade Immediately** to choose rollback options and upgrade now. If you upgrade now, it is especially important to have completed as much of the pre-upgrade checklist as possible (see the next step).

Step 3 Perform final pre-upgrade checks, including the readiness check.

Revisit the pre-upgrade checklist. Make sure you have completed all relevant tasks, especially the final checks. If you do not run the readiness check manually, it runs when you initiate the upgrade. If the readiness check fails, the upgrade is canceled. For more information, see Running an Upgrade Readiness Check, on page 793.

Step 4 Click **Upgrade Now** to start the upgrade.

a) Choose rollback options.

You can **Automatically cancel on upgrade failure and roll back to the previous version**. With this option enabled, the device automatically returns to its pre-upgrade state upon major or maintenance upgrade failure. Disable this option if you want to be able to manually cancel or retry a failed upgrade.

b) Click **Continue** to upgrade and reboot the device.

You are automatically logged off and taken to a status page where you can monitor the upgrade until the device reboots. The page also includes an option to cancel the in-progress installation. If you disabled automatic rollback and the upgrade fails, you can manually cancel or retry the upgrade.

Traffic is dropped while you upgrade. For the ISA 3000 only, if you configured hardware bypass for power failure, traffic is dropped during the upgrade but is passed without inspection while the device completes its post-upgrade reboot.

Step 5 Log back in when you can and verify upgrade success.

The Device Summary page shows the currently running software version.

Step 6 Complete post-upgrade tasks.

a) Update system databases. If you do not have automatic updates configured for intrusion rules, VDB, and GeoDB, update them now.
b) Complete any other required post-upgrade configuration changes.
c) Deploy.

Running an Upgrade Readiness Check

Before the system installs an upgrade, it runs a readiness check to ensure the upgrade is valid for the system, and to check other items that sometimes prevent a successful upgrade. If the readiness check fails, you should fix the problems before trying the installation again. If the check has failed, you will be prompted about the failure the next time you try the installation, and you are given the option to force the installation if you want to.

You can also manually run the readiness check prior to initiating the upgrade, as described in this procedure.

Before you begin

Upload the upgrade package you want to check.

Procedure

Step 1 Select **Device**, then click **View Configuration** in the Updates summary.

The **System Upgrade** section shows the currently running software version and any update that you have already uploaded.

Step 2 Look at the **Readiness Check** section.

- If the upgrade check has not been performed yet, click the **Run Upgrade Readiness Check** link. The progress of the check is shown in this area. It should take about 20 seconds to complete the process.
- If the upgrade check has already been run, this section indicates whether the check succeeded or failed. For failed checks, click **See Details** to view more information about the readiness check. After fixing problems, run the check again.

Step 3 If the readiness check fails, you should resolve the issues before you install the upgrade. The detailed information includes help on how to fix indicated problems. For a failed script, click the **Show Recovery Message** link to see the information.

Following are some typical problems:

- FXOS version incompatibility—On systems where you install FXOS upgrades separately, such as the Firepower 4100/9300, an upgrade package might require a different minimum FXOS version than the Firewall Threat Defense software version you are currently running. In this case, you must first upgrade FXOS before you can upgrade the Firewall Threat Defense software.
- Unsupported device model—The upgrade package cannot be installed on this device. You might have uploaded the wrong package, or the device is an older model that is simply no longer supported in the new Firewall Threat Defense software version. Please check device compatibility and upload a supported package, if one is available.
- Insufficient disk space—If not enough space is available, try deleting unneeded files, such as system backups. Delete only those files you have created.

Monitoring Firewall Threat Defense Upgrades

When you start the Firewall Threat Defense upgrade, you are automatically logged off and taken to a status page where you can monitor overall upgrade progress. The page also includes an option to cancel the in-progress installation. If you disabled automatic rollback and the upgrade fails, the page allows you to manually cancel or retry the upgrade.

You can also SSH to the device and use the CLI: **show upgrade status**. Add the **continuous** keyword to view log entries as they are made, and **detail** to see detailed information. Add both keywords to get continuous detailed information.

After the upgrade completes, you lose access to the status page and the CLI when the device reboots.

Canceling or Retrying Firewall Threat Defense Upgrades

Use the upgrade status page or the CLI to manually cancel failed or in-progress major or maintenance upgrades, and to retry failed upgrades:

- Upgrade status page: Click **Cancel Upgrade** to cancel an in-process upgrade. If the upgrade fails, you can click **Cancel Upgrade** to stop the job and to return to the state of the device prior to the upgrade, or click **Continue** to retry the upgrade.
- CLI: Use **upgrade cancel** to cancel an in-process upgrade. If the upgrade fails, you can use **upgrade cancel** to stop the job and to return to the state of the device prior to the upgrade, or use **upgrade retry** to retry the upgrade.



Note By default, Firewall Threat Defense automatically reverts to its pre-upgrade state upon upgrade failure ("auto-cancel"). To be able to manually cancel or retry a failed upgrade, disable the auto-cancel option when you initiate the upgrade. In a high availability deployment, auto-cancel applies to each device individually. That is, if the upgrade fails on one device, only that device is reverted.

Cancel and retry are not supported for patches. For information on reverting a successful upgrade, see Reverting Firewall Threat Defense, on page 794.

Reverting Firewall Threat Defense

If a major or maintenance upgrade succeeds but the system does not function to your expectations, you can revert. Reverting Firewall Threat Defense returns the software to its state just before the last major or maintenance upgrade; post-upgrade configuration changes are not retained. Reverting after patching necessarily removes patches as well. Note that you cannot revert individual patches or hotfixes.

The following procedure explains how to revert from Firewall Device Manager. If you cannot get into Firewall Device Manager, you can revert from the Firewall Threat Defense command line in an SSH session using the **upgrade revert** command. You can use the **show upgrade revert-info** command to see what version the system will revert to.

Before you begin

If the unit is part of a high availability pair, you must revert both units. Ideally, initiate the revert on both units at the same time so that the configuration can be reverted without failover issues. Open sessions with both

units and verify that revert will be possible on each, then start the processes. Note that traffic will be interrupted during the revert, so do this at off hours if at all possible.

For the Firepower 4100/9300 chassis, major Firewall Threat Defense versions have a specially qualified and recommended companion FXOS version. This means that after you revert the Firewall Threat Defense software, you might be running a non-recommended version of FXOS (too new). Although newer versions of FXOS are backwards compatible with older the Firewall Threat Defense versions, we do perform enhanced testing for the recommended combinations. You cannot downgrade FXOS, so if you find yourself in this situation, and you want to run a recommended combination, you will need to reimage the device.

Procedure

Step 1 Select **Device**, then click **View Configuration** in the **Updates** summary.

Step 2 In the **System Upgrade** section, click the **Revert Upgrade** link.

You are presented with a confirmation dialog box that shows the current version and the version to which the system will revert. If there is no available version to revert to, there will not be a **Revert Upgrade** link.

Step 3 If you are comfortable with the target version (and one is available), click **Revert**.

After you revert, you must re-register the device with the Smart Software Manager.

Troubleshooting Threat Defense Upgrades

These issues can occur when you are upgrading any device, whether standalone or in a high availability pair. To troubleshoot issues specific to high availability upgrades, see Troubleshooting High Availability Threat Defense Upgrades, on page 222.

Upgrade package errors.

To find the correct upgrade package, select or search for your model on the Cisco Support & Download site, then browse to the software download page for the appropriate version. Available upgrade packages are listed along with installation packages, hotfixes, and other applicable downloads. Upgrade package file names reflect the platform, package type (upgrade, patch, hotfix), software version, and build.

Upgrade packages from Version 6.2.1+ are signed, and terminate in .sh.REL.tar. If you are upgrading from an older version, download the package that terminates in .sh instead. The Cisco Support & Download site indicates the correct package for your version. Do not untar signed upgrade packages. Do not rename upgrade packages or transfer them by email.

Cannot reach the device at all during upgrade.

Devices stop passing traffic during the upgrade or if the upgrade fails. Before you upgrade, make sure traffic from your location does not have to traverse the device itself to access the device's management interface.

Device appears inactive or is unresponsive during upgrade.

You can manually cancel in-progress major and maintenance upgrades; see Canceling or Retrying Firewall Threat Defense Upgrades, on page 794. If the device is unresponsive, or if you cannot cancel the upgrade, contact Cisco TAC.

Caution Even if the system appears inactive, do *not* manually reboot or shut down during upgrade. You could place the system in an unusable state and require a reimage.

Upgrade is successful but the system does not function to your expectations.

First, make sure that cached information gets refreshed. Do not simply refresh the browser window to log back in. Instead, delete any "extra" path from the URL and reconnect to the home page; for example, http://threat-defense.example.com/.

If you continue to have issues and need to return to an earlier major or maintenance release, you may be able to revert; see Reverting Firewall Threat Defense, on page 794. If you cannot revert, you must reimage.

Upgrade fails.

When you initiate a major or maintenance upgrade, you use the **Automatically cancel on upgrade failure...** (auto-cancel) option to choose what happens if upgrade fails, as follows:

- Auto-cancel enabled (default): If upgrade fails, the upgrade cancels and the device automatically reverts to its pre-upgrade state. Correct any issues and try again later.
- Auto-cancel disabled: If upgrade fails, the device remains as it is. Correct the issues and retry immediately, or manually cancel the upgrade and try again later.

For more information, see Canceling or Retrying Firewall Threat Defense Upgrades, on page 794. If you cannot retry or cancel, or if you continue to have issues, contact Cisco TAC.

Reimaging the Device

Reimaging a device involves wiping out the device configuration and installing a fresh software image. The intention of reimaging is to have a clean installation with a factory default configuration.

You would reimage the device in these circumstances:

- You want to convert the system from ASA Software to Firewall Threat Defense Software. You cannot upgrade a device running an ASA image to one running a Firewall Threat Defense image.
- The device is running a pre-6.1.0 image, and you want to upgrade to 6.1 or a later image and configure the device using the Firewall Device Manager. You cannot use the Firewall Management Center to upgrade a pre-6.1 device and then switch to local management.
- The device is not functioning correctly and all attempts at fixing the configuration have failed.

For information on how to reimage a device, see *Reimage the Cisco ASA or Threat Defense Device* or the *Threat Defense Quick Start* guide for your device model. These guides are available at http://www.cisco.com/c/en/us/support/security/firepower-ngfw/products-installation-guides-list.html.

Backing Up and Restoring the System

You can back up the system configuration so that you can restore the device if the configuration becomes corrupted due to subsequent miss-configuration or physical mishap.

You can restore a backup onto a replacement device only if the two devices are the same model and are running the same version of the software (including the build number, not just the same point release). Do not use the backup and restore process to copy configurations between appliances. A backup file contains information that uniquely identifies an appliance, so that it cannot be shared in this manner.

Note The backup does not include the management IP address configuration. Thus, when you recover a backup file, the management address is not replaced from the backup copy. This ensures that any changes you made to the address are preserved, and also makes it possible to restore the configuration on a different device on a different network segment. The backup also does not include licensing or cloud registration information, so whatever license or cloud registration state exists at the time of restore is retained.

Backups include the configuration only, and not the system software. If you need to completely reimage the device, you need to reinstall the software, then you can upload a backup and recover the configuration.

The configuration database is locked during backup. You cannot make configuration changes during a backup, although you can view policies, dashboards, and so forth. During a restore, the system is completely unavailable.

The table on the Backup and Restore page lists all existing backup copies that are available on the system, including the file name of the backup, the date and time it was created, and the file size. The type of backup (manual, scheduled, or recurring) is based on how you directed the system to create that backup copy.

Tip Backup copies are created on the system itself. You must manually download backup copies and store them on secure servers to ensure that you have the backup copies you need for disaster recovery. The system maintains up to 3 backup copies on the device. New backups replace the oldest backup.

The following topics explain how to manage backup and restore operations.

Backing Up the System Immediately

You can start a backup whenever you want.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Backup and Restore summary.

This opens the Backup and Restore page. The table lists all existing backup copies that are available on the system.

Step 2 Click **Manual Backup** > **Back Up Now**.

Step 3 Enter a name for the backup and optionally a description.

If you decide you want to perform the backup at a future time rather than immediately, you can click **Schedule** instead.

Step 4 (Optional.) Select the **Encrypt File** option to encrypt the backup file.

If you select the option, you must enter the **Password** (and **Confirm Password**) that will be required to restore the backup file.

Step 5 (ISA 3000 only.) Select the **Location of Backup Files**.

You can create the backup on the **Local Hard Disk** or on the **SD Card**. The benefit of using the SD card is that you can use the card to recover the configuration to a replacement device.

Step 6 Click **Back Up Now**.

The system starts the backup process. When the backup is complete, the backup file will appear in the table. You can then download the backup copy to your system and store it elsewhere, if desired.

You can leave the Backup and Restore page after initiating the backup. However, the system will likely be sluggish, and you should consider pausing your work to allow the backup to complete.

In addition, the system will acquire locks on the configuration database during parts or all of the backup, which can prevent you from making changes for the duration of the backup process.

Backing Up the System at a Scheduled Time

You can set up a scheduled backup to back up the system at a specific future date and time. A scheduled backup is a one-time occurrence. If you want to create a backup schedule to regularly create backups, configure a recurring backup instead of a scheduled backup.



Note If you want to delete the schedule for a future backup, edit the schedule and click **Remove**.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Backup and Restore summary.

Step 2 Click **Scheduled Backup** > **Schedule a Backup**.

If you already have a scheduled backup, click **Scheduled Backup** > **Edit** .

Step 3 Enter a name for the backup and optionally a description.

Step 4 Select the date and time for the backup.

Step 5 (Optional.) Select the **Encrypt File** option to encrypt the backup file.

If you select the option, you must enter the **Password** (and **Confirm Password**) that will be required to restore the backup file.

Step 6 (ISA 3000 only.) Select the **Location of Backup Files**.

You can create the backup on the **Local Hard Disk** or on the **SD Card**. The benefit of using the SD card is that you can use the card to recover the configuration to a replacement device.

Step 7 Click **Schedule**.

When the selected date and time arrives, the system takes a backup. When completed, the backup copy is listed in the table of backups.

Setting Up a Recurring Backup Schedule

You can set up a recurring backup to back up the system on a regular schedule. For example, you could take a backup every Friday at midnight. A recurring backup schedule helps ensure that you always have a set of recent backups.



Note If you want to delete a recurring schedule, edit the schedule and click **Remove**.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Backup and Restore summary.

Step 2 Click **Recurring Backup** > **Configure**.

If you already have a recurring backup configured, click **Recurring Backup** > **Edit**.

Step 3 Enter a name for the backup and optionally a description.

Step 4 Select the **Frequency** and the related schedule:

- **Daily**—Select the time of day. A backup is taken every day at the scheduled time.
- **Weekly**—Select the days of the week and the time of day. A backup is taken on each day you select at the scheduled time. For example, you could schedule backups for every Monday, Wednesday, and Friday at 23:00 hours (11 PM).
- **Monthly**—Select the days of the month and the time of day. A backup is taken on each day you select at the scheduled time. For example, you could schedule backups for the first (1), fifteenth (15), and twenty-eighth (28) days of the month at 23:00 hours (11 PM).

The time you specify is adjusted for Daylight Savings Time, so it will move an hour forward or backward whenever the time is adjusted in your area. You must edit the schedule at the time change if you want to keep this exact time throughout the year.

Step 5 (Optional.) Select the **Encrypt File** option to encrypt the backup file.

If you select the option, you must enter the **Password** (and **Confirm Password**) that will be required to restore the backup file.

Step 6 (ISA 3000 only.) Select the **Location of Backup Files**.

You can create the backup on the **Local Hard Disk** or on the **SD Card**. The benefit of using the SD card is that you can use the card to recover the configuration to a replacement device.

Step 7 Click **Save**.

When the selected dates and times arrive, the system takes a backup. When completed, the backup copy is listed in the table of backups.

The recurring schedule continues to take backups until you change or remove it.

Restoring a Backup

You can restore backups as needed so long as the device is running the same software version (including build number) as it was running when you took the backup. You can restore a backup onto a replacement device only if the two devices are the same model and are running the same version of the software (including build number).

However, you cannot restore a backup if the device is part of a high availability pair. You must first break HA from the **Device** > **High Availability** page, then you can restore the backup. If the backup includes the HA configuration, the device will rejoin the HA group. Do not restore the same backup on both units, because they would then both go active. Instead, restore the backup on the unit you want to go active first, then restore the equivalent backup on the other unit.

If the backup copy you want to restore is not already on the device, you must upload the backup first before restoring it.

During a restore, the system is completely unavailable.



Note The backup does not include the management IP address configuration. Thus, when you recover a backup file, the management address is not replaced from the backup copy. This ensures that any changes you made to the address are preserved, and also makes it possible to restore the configuration on a different device on a different network segment. The backup also does not include licensing or cloud registration information, so whatever license or cloud registration state exists at the time of restore is retained.

Before you begin

If you are restoring the backup on a different system, for example, when replacing a device, best practice is to first register the device and enable any optional licenses required by the features configured in the backup file. The backup file does not include license or cloud services information, so whatever license changes or cloud registrations you make prior to restore are retained.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Backup and Restore summary.

This opens the Backup and Restore page. The table lists all existing backup copies that are available on the system.

Step 2 If the backup copy that you want to restore is not in the list of available backups, click **Upload** > **Browse** and upload the backup copy.

Step 3 Click the restore icon ( ) for the file.

You are asked to confirm the restore. By default, the backup copy will be deleted after the restore, but you can select **Do not remove the backup after restoring** to keep it before proceeding with the restore.

If the backup file was encrypted, you must enter the **Password** that is required to open the file and unencrypt it.

The system will reboot after restore completes.

Note
After the system reboots, it automatically checks for Vulnerability Database (VDB), Geolocation, and Rules database updates, and downloads them if needed. Because these updates can be large, the initial attempt might fail. Please check the task list, and if a download failed, manually download an update as described in Updating System Databases, on page 788. The system also redeploys policies. Any subsequent deployment will fail until the update is successful.

Step 4 If necessary, click **Device** > **Smart License** > **View Configuration**, re-register the device, and re-enable the required optional licenses.

The backup does not include license or cloud registration information. Thus, if you restore a backup to a new system, for example, when replacing a device, and the system is in evaluation mode, you need to register it and enable any licenses you need. If you registered the device and enabled licenses prior to the restore, no additional changes are required.

If you are simply restoring a previous backup to the same system, you should not need to make any changes to the licenses or cloud registration. However, verify that all optional licenses needed are enabled, as the backup could possible include features that require licenses you disabled after the backup was created.

Managing Backup Files

As you create new backups, the backup files are listed on the Backup and Restore page. Backup copies are not retained indefinitely: as disk space usage on the device reaches the maximum threshold, older backup copies are deleted to make room for newer ones. In addition, when you install any upgrade other than a hot fix, all backup files are deleted. Thus, you should regularly manage the backup files to ensure that you have the specific backup copies you most want to keep.

You can do the following to manage your backup copies:

- Download files to secure storage—To download a backup file to your workstation, click the download icon ( ) for the file. You can then move the file to your secure file storage.
- Upload a backup file to the system—If you want to restore a backup copy that is no longer available on the device, click **Upload** > **Browse File** and upload it from your workstation. You can then restore it.

Note Uploaded files may be renamed to match the original filename. Also, if there are more than 3 backup copies already on the system, the oldest one will be deleted to make room for the uploaded file. You cannot upload files that were created by an older software version.

- Restore a backup—To restore a backup copy, click the restore icon ( ) for the file. The system is unavailable during restore, and will reboot after restore completes. You should deploy the configuration after the system is up and running.

- Delete a backup file—If you no longer want a particular backup, click the delete icon ( ) for the file. You are asked to confirm the deletion. Once deleted, you cannot recover the backup file.

Auditing and Change Management

You can view status information about system events and actions that users have taken. This information can help you audit the system and ensure that the system is being managed properly.

Click **Device** > **Device Administration** > **Audit Log** to see the audit log. In addition, you can find system management information by clicking the **Task List** or **Deployment** icon buttons in the upper right corner.

The following topics cover some of the main concepts and tasks for system auditing and change management.

Audit Events

The audit log can include the following types of event:

Custom Feed Update Event, Custom Feed Update Failed

These events indicate a successfully completed or failed update to a custom Security Intelligence feed. The details include who started the update and information about the feed that was being updated.

Custom Rules File Import Summary Event

These events indicate that you imported a file that contained one or more custom intrusion rule. The event includes a summary of the number of rules added, updated, and deleted, and a differences view that shows details about the imported rules.

Deployment Completed, Deployment Failed: *job name* or *entity name*

These events indicate a successfully completed or failed deployment job. The details include who started the job and information about the job entity. Failed jobs include the error message related to the failure.

The details also include a **Differences View** tab, which shows the changes that were deployed to the device in the job. This is the combination of all the Entity change events for the deployed entities.

To filter on these events, simply click the **Deployment History** pre-defined filter. Note that the event type for these events is Deployment Event, you cannot filter on completed or failed events only.

The event name includes the user-defined job name (if you configure one), or "User (*username*) Triggered Deployment." There are also "Device Setup Automatic Deployment" and "Device Setup Automatic Deployment (Final Step)" jobs that occur during the device setup wizard.

Entity Created, Entity Updated, Entity Deleted: *entity name* (*entity type*)

These events indicate that a change was made to the identified entity or object. The entity details include who made the change, as well as the entity name, type, and ID. You can filter on these items. The details also include a **Differences View** tab, which shows the changes that were made to the object.

HA Action Event

These events relate to actions on the high availability configuration, either actions that you initiated, or actions that the system initiated. HA Action Event is the event type, but the event names are one of the following:

- **HA Suspended**—You intentionally suspended HA on the system.

- **HA Resumed**—You intentionally resumed HA on the system.
- **HA Reset**—You intentionally reset HA on the system.
- **HA Failover: Unit Switched Modes**—Either you intentionally switched modes, or the system failed over due to health metric violations. The message indicates that the active peer became standby, or the standby peer became active.

High Availability Sync Completed

The active unit synchronized the configuration with the standby unit. The event includes the change information for the previous version compared to the synchronized version.

Interface List Scanned

This event indicates that you scanned for changes in the interface inventory.

Pending Changes Discarded

This event indicates that you deleted all pending changes. Any changes indicated in Entity Created, Entity Updated, and Entity Deleted events between this event and the previous Deployment Completed event are removed, and the state of the affected objects is returned to the last deployed version.

Rules Update Event

When running Snort 3, this event from the LSPUpdateServer entity shows detailed information about the intrusion rules that were added, removed, or changed when a new intrusion rules package was downloaded and installed. The event is limited to 100 rules, so if more than 100 are added, removed, or changed, the event will not have complete information.

Task Started, Task Completed, Task Failed

The task events indicate the start and end of a job initiated either by the system or a user. These two events are consolidated into a single task in the task list, which you can see by clicking the **Task List** button in the upper right corner.



Tasks include actions such as deployment jobs and manual or scheduled database updates. Any item in the task list will correspond to two task events in the audit log, an indication of the start of the task, and either a successful completion or a failure.

User Logged In, User Logged Out: *username*

These events show the time and source IP address for the user logging into and out of the Firewall Device Manager. The User Logged Out event occurs for both active log outs and automatic log outs due to idle time being exceeded.

These events do not relate to RA VPN users establishing connections with the device. They also do not include log in/log out to the device CLI.

Viewing and Analyzing the Audit Log

The audit log includes information about system-initiated and user-initiated events such as deployment jobs, database updates, and login/logout of the Firewall Device Manager.

For an explanation of the types of event you can see in the log, see Audit Events, on page 802.

Procedure

Step 1 Click **Device**, then click the **Device Administration** > **View Configuration** link.

Step 2 Click **Audit Log** in the table of contents if it is not already selected.

Events are grouped by date, and within a day, by time, with the most recent date/time at the top of the list. Initially, each event is collapsed, so you see only the time, event name, user who initiated the event, and source IP address of the user. "System" for user and IP address means that the device itself initiated the event.

You can do the following:

- Click > next to the event name to open it and see the event details. Click the icon again to close the event. Many events have a simple list of event attributes, such as event type, user name, source IP address, and so forth. However, Entity and Deployment events have two tabs:
  - **Summary** shows the basic event attributes.
  - **Differences View** shows a comparison of the existing "deployed" configuration with the changes made as part of the event. For deployment jobs, this view can be long and require scrolling. It sums up all differences from the Entity event changes that were part of the deployment job.
- Select a different time range from the drop-down list to the right of the filter field. The default is to view events from the past 2 weeks, but you can change that to the last 24 hours, 7 days, month, or 6 months. Click **Custom** to specify an exact range by entering the start and end date and time.
- Click any link in the log to add a search filter for that item. The list updates so that only those events that include the item are shown. You can also simply click in the **Filter** box and build a filter directly. There are some pre-defined filters beneath the filter box that you can click to load the related filter criteria. For detailed information on filtering the events, see Filtering the Audit Log, on page 804.
- Reload the browser page to refresh the log with the latest events.

Filtering the Audit Log

You can apply a filter to the audit log to narrow your view to certain types of message only. Each element in the filter is an exact, complete match. For example, "User = admin" shows only those events initiated by the user with the name **admin**.

You can use the following techniques, alone or in combination, to build a filter. The list is automatically updated each time you add a filter element.

Click a Predefined Filter

Beneath the **Filter** field are the predefined filters. Simply click a link to load the filter. You are asked for confirmation. If you already have a filter applied, it is replaced; it is not added to.

Clicking Highlighted Items

The easiest way to build a filter is to click on items in the log table or event details that contain the values on which you intend to filter. Clicking an item updates the **Filter** field with a correctly-formulated element

for that value and element combination. However, using this technique requires that the existing list of events contains the desired values.

If you can add a filter element for an item, the item is underlined when you mouse over it and you see the command **Click to Add to Filter**.

Selecting Atomic Elements

You can also build a filter by clicking in the **Filter** field and selecting the desired atomic element from the drop-down list, typing in the match value after the equal sign, then pressing Enter. You can filter on the following elements. Note that not all elements are relevant for every event type.

- **Event Type**—This is usually but not always the same as the event name (without variable qualifiers like entity name or user). For deployment events, the event type is Deployment Event. For an explanation of the event types, see Audit Events, on page 802.
- **User**—The name of the user who initiated the event. The system user is spelled in all capitals: SYSTEM.
- **Source IP**—The IP address from which the user initiated the event. The source IP address for system-initiated events is SYSTEM.
- **Entity ID**—The UUID for the entity or object, which is a long unreadable string such as 8e7021b4-2e1e-11e8-9e5d-0fc002c5f931. Normally, to use this filter you either need to click an entity ID in an event's details, or retrieve the necessary ID through a relevant GET call using the REST API.
- **Entity Name**—The name of the entity or object. For user-created entities, this is typically the name you gave the object, for example, InsideNetwork for a network object. For system-generated entities, or in some cases user-defined entities, this is a predefined but intelligible name, for example, "User (admin) Triggered Deployment" for deployment jobs you do not explicitly name.
- **Entity Type**—The kind of entity or object. These are predefined but intelligible names, such as Network Object. You can find entity types in the API Explorer by looking at the relevant object model for the "type" value. The API types are normally all lower-case with no spaces. If you type them in exactly as shown in the model, the string changes to a more readable format when you press Enter. Typing in either form works. To open API Explorer, click the more options button (⋮) and choose **API Explorer**.

Rules for Complex Audit Log Filters

When building a complex filter that contains more than one atomic element, keep the following rules in mind:

- Elements of the same type have an OR relationship between all values for that type. For example, including "User = admin" and "User = SYSTEM" matches events that were initiated by either user.
- Elements of different types have an AND relationship. For example, including "Event Type = Entity Updated" and "User = SYSTEM" shows only those events where the system updated an entity rather than an active user.
- You cannot use wildcards, regular expressions, partial matches, or simple text string matches.

Examining Deployment and Entity Change History

Deployment and entity events include a **Differences View** tab in the event details. This tab shows a color-coded comparison of the old configuration with the changes.

- For deployment jobs, this is a comparison of the configuration that was running on the device prior to deployment to the changes that were actually deployed.
- For entity events, these are the configuration changes made to the previous version of the object. The previous version might be the version actually on the device, or it might be a change to an object that has not yet been deployed.

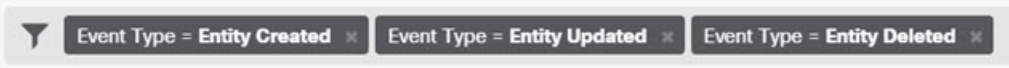
Procedure

Step 1 Click **Device**, then click the **Device Administration** > **View Configuration** link.

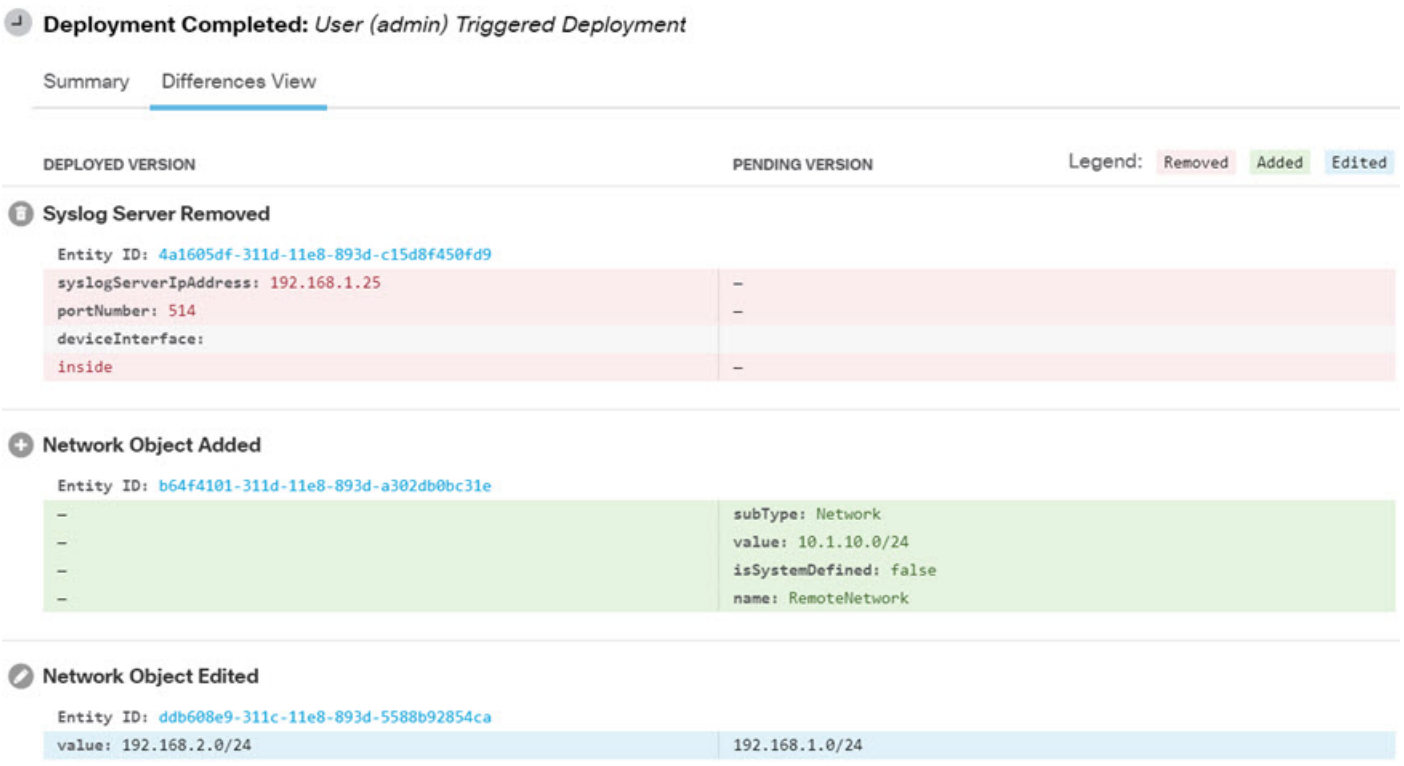
Step 2 Click **Audit Log** in the table of contents if it is not already selected.

Step 3 (Optional.) Filter the messages:

- Deployment events—Click the **Deployment History** predefined filter under the filter box.
- Entity change events—Manually create a filter using the Event Type element for the type of change that interests you. To see all entity changes, include three specifications for Entity Created, Entity Updated, and Entity Deleted. The filter would look like the following:



Step 4 Open the event and click the **Differences View** tab.



The changes are color coded, and the heading indicates the type of object and whether it was Added (Created), Removed (Deleted), or Edited (Updated). Edited objects show only those attributes that were changed or deleted from the object. In deployment jobs, there are separate headings for each entity changed. The heading indicates the entity type for the object.

Discarding All Pending Changes

If you are unsatisfied with a set of configuration changes that have not yet been deployed, you can discard all pending changes. This returns all features to the state that exists on the device. You can then start again on your configuration changes.

Procedure

Step 1 Click the **Deploy Changes** icon in the upper right of the web page.

The icon is highlighted with a dot when there are pending changes.



Step 2 Click **More Options** > **Discard All**.

Step 3 Click **OK** in the confirmation dialog.

The system discards the changes, and you will see a message that there are no pending changes when the process completes. The system adds a Pending Changes Discarded event to the audit log.

Exporting the Device Configuration

You can export a copy of the currently-deployed configuration in JSON format. You can use the file for archival or record-keeping purposes. Any sensitive data, such as passwords and secret keys, is masked.

You cannot import the file into this or another device. This ability is not a replacement for backing up the system.

You must have completed at least one successful deployment job before you can download the configuration.

Procedure

Step 1 Choose **Device**, then click **View Configuration** in the **Device Administration** group.

Step 2 Click **Download Configuration** in the table of contents.

Step 3 Click **Get Device Configuration** to start a job that creates the file.

If you have previously created a file, you will see a download button and the message **File is ready to download**, with the creation date for the file.

Depending on the size of the configuration, it can take several minutes to generate the file. Check the task list or audit log, or return to this page periodically, until the Export Config job completes and the file is generated.

Step 4 When the file is generated, return to this page and click the **Download the Configuration File** button ( ) to save the file to your workstation.

Managing Firewall Device Manager and Firewall Threat Defense User Access

You can configure an external authentication and authorization source for users to log into Firewall Threat Defense (HTTPS access). You can use an external server in addition to, or instead of, the local user database and the system-defined **admin** user. Note that you cannot create additional local user accounts for Firewall Device Manager access.

Although you can have multiple external Firewall Device Manager user accounts that can change the configuration, these changes are not tracked by user. When one user deploys changes, changes made by all users are deployed. There is no locking: that is, more than one user might attempt to update the same object at the same time, which will result in only one user successfully saving the change. You also cannot discard changes based on user.

You can have 5 concurrent user sessions. If a sixth user logs in, the oldest user session is automatically logged off. There is also an idle timeout, which logs inactive users out after 20 minutes.

You can also configure external authentication and authorization for SSH access to the Firewall Threat Defense CLI. The local database is always checked before using the external source, so you can create additional local users for failsafe access. Do not create duplicate users in both the local and external source. Except for the **admin** user, there is no crossover between the CLI and Firewall Device Manager users: the user accounts are completely separate.



Note When using external servers, you can control access by user to subsets of your devices by either setting up separate AAA server groups, or by creating authentication/authorization policies within the AAA servers that allow the user access to certain Firewall Threat Defense device IP addresses only.

The following topics explain how to configure and manage Firewall Device Manager user access and CLI user access.

Configuring External Authorization (AAA) for the Firewall Device Manager (HTTPS) Users

You can provide HTTPS access to the Firewall Device Manager from an external AAA server. By enabling AAA authentication and authorization, you can provide different levels of access rights, and not have every user log in through the local **admin** account.

These external users are also authorized for the Firewall Threat Defense API and the API Explorer.

You can provide role-based access control (RBAC) by setting up authorization for management users in the AAA server. The levels differ depending on server type. When a user logs into Firewall Device Manager, the username and role are shown in the upper right of the page. After you set up the accounts correctly on the AAA server, you can enable it for administrative access using this procedure.

RADIUS User Authorization

To provide role-based access control (RBAC), update the user accounts on your RADIUS server to define the **cisco-av-pair** attribute (in ISE, but the attribute is spelled Cisco-AVPair in Free RADIUS; check your system for correct spelling). This attribute must be defined correctly on a user account, or the user is denied access to the Firewall Device Manager. Following are the supported values for the **cisco-av-pair** attribute:

- **fdm.userrole.authority.admin** provides full Administrator access. These users can do all actions that the local **admin** user can do.
- **fdm.userrole.authority.rw** provides read-write access. These users can do everything a read-only user can do, and also edit and deploy the configuration. The only restrictions are for system-critical actions, which include installing upgrades, creating and restoring backups, and ending the sessions of the Firewall Device Manager users.
- **fdm.userrole.authority.ro** provides read-only access. These users can view dashboards and the configuration, but cannot make any changes. If the user tries to make a change, the error message explains that this is due to lack of permission.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Management Access** link.

If you are already on the System Settings page, simply click **Management Access** in the table of contents.

Step 2 Click the **AAA Configuration** tab if it is not already selected.

Step 3 Configure the **HTTPS Connection** options:

- **Server Group for Management/REST API**—Select the server group (for external authentication/authorization) or local user database (LocalIdentitySource), that you want to use as a primary authentication source.

  If the server group does not yet exist, click the link to create a new one and create it now. For RADIUS, you will also need to create server objects for each server, to add them to the group, but you can do this while defining the server group. For more information on RADIUS, see RADIUS Servers and Groups, on page 165.

- **Authentication with LOCAL** (RADIUS.)—If you select an external RADIUS server group, you can specify how to use the local identity source, which contains the local **admin** user account. Select one of the following:

  - **Before External Server**—The system checks the username and password against the local source first.
  - **After External Server**—The local source is checked only if the external source is unavailable.
  - **Never**—(Not recommended.) The local source is never used, so you cannot log in as the admin user.

    Caution
    If you select **Never**, you will not be able to log into the Firewall Device Manager using the **admin** account. You will be locked out of the system if the AAA server becomes unavailable, or if you miss-configure the accounts in the AAA server.

Step 4 Click **Save**.

Configuring External Authorization (AAA) for the Firewall Threat Defense CLI (SSH) Users

You can provide SSH access to the Firewall Threat Defense CLI from an external RADIUS server. By enabling RADIUS authentication and authorization, you can provide different levels of access rights from a single authentication source, rather than define separate local user accounts on each device.

These SSH external users are **not** authorized for the Firewall Threat Defense API and the API Explorer. The mechanism you use to define authorization for SSH is different from the one required for HTTPS access. However, you can configure the same RADIUS user with both SSH and HTTPS authorization criteria, so that a given user can access the system through both protocols.

To provide role-based access control (RBAC) for SSH access, update the user accounts on your RADIUS server to define the **Service-Type** attribute. This attribute must be defined on a user account, or the user is denied SSH access to the device. Following are the supported values for the **Service-Type** attribute:

- **Administrative (6)**—Provides **config** access authorization to the CLI. These users can use all commands in the CLI.

- **NAS Prompt (7)** or any level other than 6—Provides **basic** access authorization to the CLI. These users can use read-only commands, such as **show** commands, for monitoring and troubleshooting purposes.

After you set up the accounts correctly on the RADIUS server, you can enable it for SSH administrative access using this procedure.



Note Do not create duplicate users in both the local and external source. If you do create duplicate usernames, ensure that they have the same authorization rights. You cannot log in using the password of the external version of the user account when the authorization rights differ in the local user account; you can log in using the local password only. If the rights are the same, the password you use determines if you are logged in as the external or the local user, assuming the passwords are different. Even though the local database is checked first, if a username exists in the local database but the password is incorrect, the external server is checked and if the password is correct for the external source, the login will succeed.

Before you begin

Please inform externally-defined users of the following behavior to set their expectations appropriately:

- The first time an external user logs in, the Firewall Threat Defense creates the required structures but cannot simultaneously create the user session. The user simply needs to authenticate again to start the session. The user will see a message similar to the following: "New external username identified. Please log in again to start a session."
- Similarly, if the user's authorization as defined in the Service-Type changed since the last login, the user will need to re-authenticate. The user will see a message similar to the following: "Your authorization privilege has changed. Please log in again to start a session."

If external users can successfully log into the GUI but not the CLI, there might be a problem with the secret key for the AAA servers. Ensure that the keys do not include & or \ characters. Although these might work with GUI logins, they do not work with SSH logins.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Management Access** link.

If you are already on the System Settings page, simply click **Management Access** in the table of contents.

Step 2 Click the **AAA Configuration** tab if it is not already selected.

Step 3 Configure the **SSH Connection** options:

- **Server Group**—Select the RADIUS server group or local user database (LocalIdentitySource), that you want to use as a primary authentication source. You must select a RADIUS server group to use external authorization.

  If the server group does not yet exist, click the **Create New RADIUS Server Group** link and create it now. You will also need to create RADIUS server objects for each server, to add them to the group, but you can do this while defining the server group. For more information on RADIUS, see RADIUS Servers and Groups, on page 165.

Note that SSH connections use the first 2 servers in the group only. If you use a group with 3 or more servers, the additional servers are never tried. In addition, the **Dead Time** and **Maximum Failed Attempts** group attributes are not used.

- **Authentication with LOCAL**—If you select an external server group, you can specify how to use the local identity source. For SSH access, the local database is always checked before the external server.

Step 4 Click **Save**.

Managing the Firewall Device Manager User Sessions

Choose **Monitoring** > **Sessions** to see a list of users who are currently logged into the Firewall Device Manager. The list shows how long each user has been logged in for the current session.

If the same username appears more than once, it means that the user has opened sessions from different source addresses. Sessions are tracked separately based on username and source address, each session with its own unique time stamp.

The system allows 5 concurrent user sessions. If a sixth user logs in, the oldest current session is automatically logged out. In addition, idle users are automatically logged out after 20 minutes of inactivity.

If the Firewall Device Manager user types in the wrong password and fails to log in on 3 consecutive attempts, the user's account is locked for 5 minutes. The user must wait before trying to log in again. There is no way to unlock the Firewall Device Manager user account, nor can you adjust the retry count or lock timeout. (Note that for SSH users, you can adjust these settings and unlock the account.)

If necessary, you can end a user session by clicking the delete icon () for the session. If you delete your own session, you are also logged out. There is no lockout period if you end a session: the user can immediately log back in.

Enabling the Firewall Device Manager Access on a Standby HA Unit for External Users

If you configure external authorization for the Firewall Device Manager users, those users can log into both the active and standby unit of a high availability pair. However, to successfully log into the standby unit for the first time requires a few extra steps compared to logging into the active unit.

After an external user logs into the active unit for the first time, the system creates an object that defines the user and the user’s access rights. An admin or read-write user must then deploy the configuration from the active unit for the user object to appear on the standby unit.

Only after the deployment and subsequent configuration synchronization completes successfully can the external user log into the standby unit.

Admin and read-write users can deploy changes after logging into the active unit. However, read-only users cannot deploy the configuration, and must ask a user who has the appropriate rights to deploy the configuration.

Creating Local User Accounts for the Firewall Threat Defense CLI

You can create users for CLI access on Firewall 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 Firewall 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. Allowed commands are dig, ping, and traceroute.

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

The user is prompted to change the password on first login. The user can initiate subsequent password changes 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.

Rebooting or Shutting Down the System

If necessary, you can reboot or shut down the system.

In addition to the procedure below, you can also perform these tasks through an SSH session or the Firewall Device Manager CLI Console using the **reboot** or **shutdown** commands.

Note The ISA 3000 does not support shutting down; you can only reboot the system.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Reboot/Shutdown** > link.

If you are already on the System Settings page, simply click **Reboot/Shutdown** in the table of contents

Step 2 Click the button that performs the function you need.

- **Reboot**—If you believe the system is not performing correctly and other efforts to resolve the problem have failed, you can reboot the device. In addition, there might be a few procedures that ask you to reboot the device to reload the system software.
- **Shut Down**—Shut down the system to turn off power in a controlled fashion. Use shutdown when you intend to remove the device from the network, for example, to replace it. After shutting down the device, you can turn it back on only from the hardware On/Off switch.

Step 3 Wait until the action completes.

If you have a console connection to the firewall, monitor the system prompts as the firewall shuts down. You will see the following prompt:

```
System is stopped.
It is safe to power off now.
Do you want to reboot instead? [y/N]
```

You cannot perform other actions in the Firewall Device Manager or the CLI while the system is rebooting or shutting down.

During reboot, the Firewall Device Manager page should refresh when the reboot is complete and bring you to the login page. If you try refreshing the page before the reboot is complete, the web browser might return 503 or 404 errors, based on the operational state of the Firewall Device Manager web server at that point in time.

For shutdown, the system will eventually not be able to respond at all and you will get 404 errors. This is the expected result, because you are completely turning off the system.

Troubleshooting the System

The following topics address some system-level troubleshooting tasks and capabilities. For information on troubleshooting a specific feature, such as access control, see the chapter for the feature.

Pinging Addresses to Test Connectivity

Ping is a simple command that lets you determine if a particular address is alive and responsive. This means that basic connectivity is working. However, other policies running on a device could prevent specific types of traffic from successfully getting through a device. You can use **ping** by opening the CLI console or by logging into the device CLI.



Note Because the system has multiple interfaces, you can control the interface used for pinging an address. You must ensure that you are using the right command, so that you are testing the connectivity that matters. For example, the system must be able to reach the Cisco license server through the virtual Management interface, so you must use the **ping system** command to test the connection. If you use **ping**, you are testing whether an address can be reached through the data interfaces, which might not give you the same result.

The normal ping uses ICMP packets to test the connection. If your network prohibits ICMP, you can use a TCP ping instead (for data interface pings only).

You can ping either an IP address or a fully-qualified hostname (FQDN). For a ping to work on an FQDN, the DNS servers configured for either the management or data interfaces must successfully return an IP address. You must configure DNS servers separately for management and data interfaces. If you do not have DNS servers configured for a specific interface, use the **dig** command to look up the IP address of a given FQDN.

Following are the main options for pinging network addresses.

Pinging an address through the virtual Management interface

Use the **ping system** command.

ping system *host*

The host can be an IP address or fully-qualified domain name (FQDN), such as www.example.com. Unlike pings through the data interfaces, there is no default count for system pings. The ping continues until you stop it using Ctrl+c. For example:

```
> ping system www.cisco.com
PING origin-www.cisco.COM (72.163.4.161) 56(84) bytes of data.
64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=1 ttl=242 time=10.6 ms
64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=2 ttl=242 time=8.13 ms
64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=3 ttl=242 time=8.51 ms
64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=4 ttl=242 time=8.40 ms
^C
--- origin-www.cisco.COM ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3003ms
rtt min/avg/max/mdev = 8.139/8.927/10.650/1.003 ms
>
```

Pinging an address through a data interface using the routing table

Use the **ping** command. Without specifying an interface, you are testing whether the system can generically find a route to the host. Because this is how the system normally routes traffic, this is typically what you want to test.

ping *host*

For example:

```
> ping 171.69.38.1
Sending 5, 100-byte ICMP Echos to 171.69.38.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/10 ms
```



Note You can specify the timeout, repeat count, packet size, and even the data pattern to send. Use the help indicator, ?, in the CLI to see the available options.

Pinging an address through a specific data interface

Use the **ping interface** *if_name* command if you want to test connectivity through a specific data interface. You can also specify the diagnostic interface using this command, but not the virtual management interface.

ping interface *if_name host*

For example:

```
> ping interface inside 171.69.38.1
Sending 5, 100-byte ICMP Echos to 171.69.38.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/10 ms
```

Pinging an address through a data interface using TCP ping

Use the **ping tcp** command. A TCP ping sends SYN packets and considers the ping successful if the destination sends a SYN-ACK packet.

ping tcp [**interface** *if_name*] *host port*

You must specify the host and TCP port.

You can optionally specify the interface, which is the source interface of the ping, not the interface through which to send the pings. This type of ping always uses the routing table.

A TCP ping sends SYN packets and considers the ping successful if the destination sends a SYN-ACK packet. For example:

```
> ping tcp 10.0.0.1 21
Type escape sequence to abort.
No source specified. Pinging from identity interface.
Sending 5 TCP SYN requests to 10.0.0.1 port 21
from 10.0.0.10, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Note You can also specify the timeout, repeat count, and the source address for the TCP ping. Use the help indicator, ?, in the CLI to see the available options.

Tracing Routes to Hosts

If you are having problems sending traffic to an IP address, you can trace the route to the host to determine if there is a problem on the network path. A traceroute works by sending UDP packets on an invalid port, or ICMPv6 echoes, to a destination. The routers along the way to the destination respond with an ICMP Time Exceeded Message, and report that error to traceroute. Each node receives three packets, so you get three chances per node to get an informative result. You can use **traceroute** by opening the CLI console or by logging into the device CLI.

Note There are separate commands for tracing a route through a data interface (**traceroute**) or through the virtual management interface (**traceroute system**). Ensure that you use the right command.

The following table describes the possible result per packet as displayed in the output.

Output Symbol	Description
*	No response was received for the probe within the timeout period.
nn msec	For each node, the round-trip time (in milliseconds) for the specified number of probes.
!N.	ICMP network unreachable.
!H	ICMP host unreachable.
!P	ICMP protocol unreachable.
!A	ICMP administratively prohibited.
?	Unknown ICMP error.

Tracing a route through the virtual management interface

Use the **traceroute system** command.

traceroute system *destination*

The host can be an IPv4/IPv6 address or fully-qualified domain name (FQDN), such as www.example.com. For example:

```
> traceroute system www.example.com
traceroute to www.example.com (172.163.4.161), 30 hops max, 60 byte packets
 1  192.168.0.254 (192.168.0.254)  0.213 ms  0.310 ms  0.328 ms
 2  10.88.127.1 (10.88.127.1)  0.677 ms  0.739 ms  0.899 ms
 3  lab-gw1.example.com (10.89.128.25)  0.638 ms  0.856 ms  0.864 ms
 4  04-bb-gw1.example.com (10.152.240.65)  1.169 ms  1.355 ms  1.409 ms
 5  wan-gw1.example.com (10.152.240.33)  0.712 ms  0.722 ms  0.790 ms
```

```
 6  wag-gw1.example.com (10.152.240.73)  13.868 ms  10.760 ms  11.187 ms
 7  rbb-gw2.example.com (172.30.4.85)  7.202 ms  7.301 ms  7.101 ms
 8  rbb-gw1.example.com (172.30.4.77)  8.162 ms  8.225 ms  8.373 ms
 9  sbb-gw1.example.com (172.16.16.210)  7.396 ms  7.548 ms  7.653 ms
10  corp-gw2.example.com (172.16.16.58)  7.413 ms  7.310 ms  7.431 ms
11  dmzbb-gw2.example.com (172.16.0.78)  7.835 ms  7.705 ms  7.702 ms
12  dmzdcc-gw2.example.com (172.16.0.190)  8.126 ms  8.193 ms  11.559 ms
13  dcz05n-gw1.example.com (172.16.2.106)  11.729 ms  11.728 ms  11.939 ms
14  www1.example.com (172.16.4.161)  11.645 ms  7.958 ms  7.936 ms
```

Tracing a route through a data interface

Use the **traceroute** command.

traceroute *destination*

The host can be an IPv4/IPv6 address or fully-qualified domain name (FQDN), such as www.example.com, if you configure DNS servers for the data interfaces. If you do not have DNS servers configured for a specific interface, use the **dig** command to look up the IP address of a given FQDN. For example:

```
> traceroute 209.165.200.225
Tracing the route to 209.165.200.225
 1  10.83.194.1 0 msec 10 msec 0 msec
 2  10.83.193.65 0 msec 0 msec 0 msec
 3  10.88.193.101 0 msec 10 msec 0 msec
 4  10.88.193.97 0 msec 0 msec 10 msec
 5  10.88.239.9 0 msec 10 msec 0 msec
 6  10.88.238.65 10 msec 10 msec 0 msec
 7 172.16.7.221 70 msec 70 msec 80 msec
 8 209.165.200.225 70 msec 70 msec 70 msec
```

Note You can specify the timeout, time to live, number of packets per node, and even the IP address or interface to use as the source of the traceroute. Use the help indicator, ?, in the CLI to see the available options.

Making the Device Appear on Traceroutes

By default, the Firewall Threat Defense device does not appear on traceroutes as a hop. To make it appear, you need to decrement the time-to-live on packets that pass through the device, and increase the rate limit on ICMP unreachable messages. To accomplish this, you must create a FlexConfig object that configures the required service policy rule and other options.

For a detailed discussion of service policies and traffic classes, see the *Cisco ASA Series Firewall Configuration Guide* available from https://www.cisco.com/c/en/us/support/security/asa-firepower-services/products-installation-and-configuration-guides-list.html.

Note If you decrement time to live, packets with a TTL of 1 will be dropped, but a connection will be opened for the session on the assumption that the connection might contain packets with a greater TTL. Note that some packets, such as OSPF hello packets, are sent with TTL = 1, so decrementing time to live can have unexpected consequences. Keep these considerations in mind when defining your traffic class.

Procedure

Step 1 Click **View Configuration** in **Device** > **Advanced Configuration**.

Step 2 Click **FlexConfig** > **FlexConfig Objects** in the Advanced Configuration table of contents.

Step 3 Create the object to decrement TTL.

a) Click the **+** button to create a new object.

b) Enter a name for the object. For example, **Decrement_TTL**.

c) In the **Template** editor, enter the following lines, including indentations.

```
icmp unreachable rate-limit 50 burst-size 1
policy-map global_policy
 class class-default
  set connection decrement-ttl
```

d) In the **Negate Template** editor, enter the lines required to undo this configuration.

Just as you need to include the parent commands to enter the correct sub-mode for a command to enable it, you also need to include those commands in the negate template.

The negate template will be applied if you remove this object from the FlexConfig policy (after having deployed it successfully), and also during an unsuccessful deployment (to reset the configuration to its previous condition).

Thus, for this example, the negate template would be the following:

```
no icmp unreachable rate-limit 50 burst-size 1
policy-map global_policy
 class class-default
  no set connection decrement-ttl
```

e) Click **OK** to save the object.

Step 4 Add the objects to the FlexConfig policy.

Only those objects selected in the FlexConfig policy get deployed.

a) Click **FlexConfig Policy** in the table of contents.

b) Click **+** in the Group List.

c) Select the Decrement_TTL object and click **OK**.

The preview should update with the commands in the template. Verify you are seeing the expected commands.

d) Click **Save**.

You can now deploy the policy.

Troubleshooting NTP

The system relies on accurate and consistent time to function correctly and to ensure that events and other data points are handled accurately. You must configure at least one, but ideally three, Network Time Protocol (NTP) servers to ensure the system always has reliable time information.

The device summary connection diagram (click **Device** in the main menu) shows the status of the connection to the NTP server. If the status is yellow or orange, then there is an issue with the connection to the configured servers. If the connection problem persists (it is not just a momentary issue), try the following.

- First, ensure that you have at least three NTP servers configured on **Device** > **System Settings** > **NTP**. Although this is not a requirement, reliability is greatly enhanced if you have at least three NTP servers.
- Ensure that there is a network path between the management interface IP address (defined on **Device** > **System Settings** > **Management Interface**) and the NTP servers.
  - If the management interface gateway is the data interfaces, you can configure static routes to the NTP servers on **Device** > **Routing** if the default route is not adequate.
  - If you set an explicit management interface gateway, log into the device CLI and use the **ping system** command to test whether there is a network path to each NTP server.
- Log into the device CLI and check the status of the NTP servers with the following commands.
  - **show ntp**—This command shows basic information about the NTP servers and their availability. However, the connectivity status in the Firewall Device Manager uses additional information to indicate the status, so there can be inconsistency in what this command shows and what the connectivity status diagram shows. You can also issue this command from the CLI console.
  - **system support ntp**—This command includes the output of **show ntp** plus the output of the standard NTP command **ntpq**, which is documented with the NTP protocol. Use this command if you need to confirm NTP synchronization.

    Look for the section "Results of 'ntpq -pn.' For example, you might see something like the following:

```
Results of 'ntpq -pn'
remote                   : +216.229.0.50
refid                    : 129.7.1.66
st                       : 2
t                        : u
when                     : 704
poll                     : 1024
reach                    : 377
delay                    : 90.455
offset                   : 2.954
jitter                   : 2.473
```

    In this example, the + before the NTP server address indicates that it is a potential candidate. An asterisk here, *, indicates the current time source peer.

    The NTP daemon (NTPD) uses a sliding window of eight samples from each one of the peers and picks out one sample, then the clock selection determines the true chimers and the false tickers. NTPD then determines the round-trip distance (the offset of a candidate must not be over one-half the round trip delay). If connection delays, packet loss, or server issues cause one or all the candidates to be rejected, you would see long delays in the synchronization. The adjustment also occurs over a very long period of time: the clock offset and oscillator errors must be resolved by the clock discipline algorithm and this can take hours.

Note If the refid is .LOCL., this indicates the peer is an undisciplined local clock, that is, it is using its local clock only to set the time. The Firewall Device Manager always marks the NTP connection yellow (not synchronized) if the selected peer is .LOCL. Normally, NTP does not select a .LOCL. candidate if a better one is available, which is why you should configure at least three servers.

Troubleshooting DNS for the Management Interface

You must configure at least one DNS server for use by the Management interface. The server is needed for cloud connections to services such as smart licensing, database updates (such as GeoDB, rules, and VDB), and any other activity that needs domain name resolution.

Configuring a DNS server is rather trivial. You simply enter the IP addresses of the DNS servers you use when you initially configure the device. You can later change them on the **Device** > **System Settings** > **DNS Server** page.

However, the system can fail to resolve fully-qualified domain names (FQDN) due to network connectivity issues or problems with the DNS server itself. If you find the system cannot use your DNS servers, consider the following actions to identify and resolve the problem. Also see Troubleshooting General DNS Problems, on page 764.

Procedure

Step 1 Determine if you have a problem.

a) Use SSH to log into the device CLI.

b) Enter **ping system www.cisco.com**. If you get an "unknown host" message like the following, then the system could not resolve the domain name. If the ping is successful, then you are done: DNS is working. (Press Ctrl+C to stop the ping.)

```
> ping system www.cisco.com
ping: unknown host www.cisco.com
```

Note
It is critical that you include the **system** keyword in the **ping** command. The **system** keyword sends the ping through the management IP address, which is the only interface that uses the management DNS server. Pinging www.cisco.com is also a good option, because you need a route to that server for smart licensing and updates.

Step 2 Verify the configuration for the management interface.

a) Click **Device** > **System Settings** > **Management Interface**, and verify the following. If you make changes, the changes are applied immediately on clicking **Save**. If you change the Management address, you will need to reconnect and log back in.

- The gateway IP address is correct for the Management network. If you using the data interfaces as the gateway, subsequent steps will verify that configuration.

- If you are not using the data interfaces as a gateway, verify that the Management IP address/subnet mask and the gateway IP address are on the same subnet.

b) Click **Device** > **System Settings** > **DNS Server** and verify that the right DNS servers are configured.

If you are deploying the device on your network edge, your service provider might have specific requirements about the DNS server you can use.

c) If you are using the data interfaces as the gateway, verify that you have the required routes.

You need a default route for 0.0.0.0. You might need additional routes if the DNS server is not available through the gateway for the default route. There are two basic situations:

- If you are using DHCP to obtain an address for the outside interface, and you selected the **Obtain Default Route using DHCP** option, the default route is not visible in the Firewall Device Manager. From SSH, enter **show route** and verify that there is a route for 0.0.0.0. Because this is the default configuration for the outside interface, this is a likely situation that you might encounter. (Go to **Device** > **Interfaces** to view the configuration of the outside interface.)
- If you are using a static IP address on the outside interface, or you are not obtaining the default route from DHCP, then open **Device** > **Routing**. Verify that the correct gateway is being used for the default route.

If the DNS server cannot be reached through the default route, you must define a static route to it on the **Routing** page. Note that you should not add routes for directly connected networks, that is, networks that are connected directly to any of the system's data interfaces, as the system can route to those networks automatically.

Also verify that there are no static routes that are misdirecting traffic to the server out the wrong interface.

d) If the deployment button indicates that there are undeployed changes, deploy them now and wait for deployment to complete.



e) Retest **ping system www.cisco.com**. If you still have problems, continue with the next step.

Step 3 In the SSH session, enter **dig www.cisco.com**.

- If **dig** indicates that it got a response from the DNS server, but the server could not find the name, it means that DNS is configured correctly, but the DNS server you are using does not have an address for the FQDN. This error is indicated by the NXDOMAIN status. The response would look similar to the following:

```
> dig www.cisco.com

; <<>> DiG 9.11.4 <<>> www.cisco.com
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NXDOMAIN, id: 43246
;; flags: qr rd ra; QUERY: 1, ANSWER: 0, AUTHORITY: 1, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 1280
; COOKIE: 78b1c6b2b3ef5b689fc2f65260db9e9b36a7d9fefb301943 (good)
;; QUESTION SECTION:
;www.cisco.com.                 IN      A
```

```
;; AUTHORITY SECTION:
.                       3600    IN      SOA     a.root-servers.net.
nstld.verisign-grs.com. 2021062901 1800 900 604800 86400

;; Query time: 13 msec
;; SERVER: 10.163.47.11#53(10.163.47.11)
;; WHEN: Tue Jun 29 22:28:43 UTC 2021
;; MSG SIZE  rcvd: 145
```

Resolution: In this case, you need to configure a different DNS server, or get the one you have updated so it can resolve the FQDNs you need resolved. Work with your network administrator or ISP to get the IP address of a DNS server that will work for your network.

- If the command times out, then the system cannot reach your DNS servers, or all of the DNS servers are currently down and not responding (which is less likely). Continue with the next step.

Step 4 Use the **traceroute system** *DNS_server_ip_address* command to trace the route to the DNS server.

For example, if the DNS server is 10.100.10.1, enter:

```
> traceroute system 10.100.10.1
```

Following are the possible results:

- The traceroute completes and reaches the DNS server. In this case, there is in fact a route to the DNS server and the system can reach it. Thus, there is no routing problem. However, for some reason, DNS requests to this server are not getting a response.

  Resolution: There is a possibility that a router or firewall along the path is dropping UDP/53 traffic, which is the port used for DNS. You might try a DNS server along a different network path. This is a difficult problem to resolve, as you will need to determine which node is blocking traffic, and work with the system administrator to get the access rules changed.

- The traceroute cannot reach even one node, which would look like the following:

```
> traceroute system 10.100.10.1
traceroute to 10.100.10.1 (10.100.10.1), 30 hops max, 60 byte packets
 1  * * *
 2  * * *
 3  * * *
(and so forth)
```

  Resolution: In this case, the routing problem is within your system. Try doing a **ping system** for the gateway IP address. Re-verify the configuration of the management interface as mentioned in earlier steps, and ensure that you have the required gateways and routes configured.

- The traceroute makes it through a few nodes before it can no longer resolve the route, which would look like the following:

```
> traceroute system 10.100.10.1
traceroute to 10.100.10.1 (10.100.10.1), 30 hops max, 60 byte packets
 1  192.168.0.254 (192.168.0.254)  0.475 ms  0.532 ms  0.542 ms
 2  10.88.127.1 (10.88.127.1)  0.803 ms  1.434 ms  1.443 ms
 3  site04-lab-gw1.example.com (10.89.128.25)  1.390 ms  1.399 ms  1.435 ms
 4  * * *
 5  * * *
```

```
6  * * *
```

Resolution: In this case, routing breaks down at the last node. You might need to work with the system administrator to get correct routes installed in that node. However, if there is intentionally no route to the DNS server through the node, you need to change your gateway, or create your own static route, to point to a router that can route traffic to the DNS server.

Analyzing CPU and Memory Usage

To view system-level information about CPU and memory usage, select **Monitoring** > **System** and look for the CPU and Memory bar graphs. These graphs show information collected through the CLI using the **show cpu system** and **show memory system** commands.

If you open the CLI console or log into the CLI, you can use additional versions of these commands to view other information. Typically, you would look at this information only if you are having persistent problems with usage, or at the direction of the Cisco Technical Assistance Center (TAC). Much of the detailed information is complex and requires TAC interpretation.

Following are some highlights of what you can examine. You can find more detailed information about these commands in Cisco Firepower Threat Defense Command Reference at http://www.cisco.com/c/en/us/td/docs/security/firepower/command_ref/b_Command_Reference_for_Firepower_Threat_Defense.html.

- **show cpu** displays data plane CPU utilization.
- **show cpu core** displays usage for each CPU core separately.
- **show cpu detailed** displays additional per-core and overall data plane CPU usage.
- **show memory** displays data plane memory usage.

Note Some of the keywords (not mentioned above) require that you first set up profiling or other features using the **cpu** or **memory** commands. Use these features at the direction of TAC only.

Viewing Logs

The system logs information for a wide variety of actions. You can use the **system support view-files** command to open a system log. Use this command while working with the Cisco Technical Assistance Center (TAC) so that they can help you interpret the output and to select the appropriate log to view.

The command presents a menu for selecting a log. Use the following commands to navigate the wizard:

- To change to a sub-directory, type in the name of the directory and press Enter.
- To select a file to view, enter **s** at the prompt. You are then prompted for a file name. You must type the complete name, and capitalization matters. The file list shows you the size of the log, which you might consider before opening very large logs.
- Press the space bar when you see --More-- to see the next page of log entries; press Enter to see just the next log entry. When you reach the end of the log, you are taken to the main menu. The --More-- line

shows you the size of the log and how much of it you have viewed. **Use Ctrl+C to close the log and exit the command if you do not want to page through the entire log.**

- Type **b** to go up one level in the structure to the menu.

If you want to leave the log open so you can see new messages as they are added, use the **tail-logs** command instead of **system support view-files**.

The following example shows how view the cisco/audit.log file, which tracks attempts to log into the system. The file listing starts with directories at the top, then a list of files in the current directory.

```
> system support view-files

===View Logs===

============================
Directory: /ngfw/var/log
----------sub-dirs----------
cisco
mojo
removed_packages
setup
connector
sf
scripts
packages
removed_scripts
httpd
-----------files------------
2016-10-14 18:12:04.514783 | 5371       | SMART_STATUS_sda.log
2016-10-14 18:12:04.524783 | 353        | SMART_STATUS_sdb.log
2016-10-11 21:32:23.848733 | 326517     | action_queue.log
2016-10-06 16:00:56.620019 | 1018       | br1.down.log

<list abbreviated>


([b] to go back or [s] to select a file to view, [Ctrl+C] to exit)
Type a sub-dir name to list its contents: cisco

============================
Directory: /ngfw/var/log/cisco
-----------files------------
2017-02-13 22:44:42.394907 | 472        | audit.log
2017-02-13 23:40:30.858198 | 903615     | ev_stats.log.0
2017-02-09 18:14:26.870361 | 0          | ev_stats.log.0.lck
2017-02-13 05:24:00.682601 | 1024338    | ev_stats.log.1
2017-02-12 08:41:00.478103 | 1024338    | ev_stats.log.2
2017-02-11 11:58:00.260805 | 1024218    | ev_stats.log.3
2017-02-09 18:12:13.828607 | 95848      | firstboot.ngfw-onbox.log
2017-02-13 23:40:00.240359 | 6523160    | ngfw-onbox.log

([b] to go back or [s] to select a file to view, [Ctrl+C] to exit)
Type a sub-dir name to list its contents: s

Type the name of the file to view ([b] to go back, [Ctrl+C] to exit)
> audit.log
2017-02-09 18:59:26 - SubSystem:LOGIN, User:admin, IP:10.24.42.205, Message:Login successful,

2017-02-13 17:59:28 - SubSystem:LOGIN, User:admin, IP:10.24.111.72, Message:Login successful,

2017-02-13 22:44:36 - SubSystem:LOGIN, User:admin, IP:10.24.111.72, Message:Login failed,
```

```
2017-02-13 22:44:42 - SubSystem:LOGIN, User:admin, IP:10.24.111.72, Message:Login successful,

2017-02-13 22:44:42 - SubSystem:LOGIN, User:admin, IP:10.24.111.72, Message:Unlocked account.,

<remaining log truncated>
```

Creating a Troubleshooting File

Cisco Technical Assistance Center (TAC) personnel might ask you to submit system log information when you submit a problem report. This information assists them with diagnosing the problem. You do not need to submit a diagnostics file unless asked to do so.

The following procedure explains how to create and download the diagnostics file.

Procedure

Step 1 Click **Device**.

Step 2 Under **Troubleshooting**, click **Request File to be Created** or **Re-Request File to be Created** (if you have created one before).

The system starts generating the diagnostic file. You can go to other pages and return here to check status. When the file is ready, the date and time of the file creation is shown along with a download button.

Step 3 When the file is ready, click the download button.

The file is downloaded to your workstation using your browser's standard download method.

Hardware Management Tasks

The following topics explain how to do some hardware maintenance tasks. See the various hardware guides for more complete information and information about other hardware management tasks.

Replacing an ISA 3000 Device

The ISA 3000 has an SD card that you can remove and insert in another ISA 3000 device. If you create system backups on the SD card, you can use this capability to easily replace a device. You simply take the SD card from the failing device and insert it in the new device. The backups are then available for you to restore.

To ensure that you have the necessary backups, configure the backup job to create the backup on the SD card.

Hot swap an SSD

If you have two SSDs, they form a RAID when you boot up. You can perform the following tasks at the Firewall Threat Defense CLI while the firewall is powered up:

- Hot swap one of the SSDs—If an SSD is faulty, you can replace it. Note that if you only have one SSD, you cannot remove it while the firewall is powered on.
- Remove one of the SSDs—If you have two SSDs, you can remove one.
- Add a second SSD—If you have one SSD, you can add a second SSD and form a RAID.

  When you use only one SSD in Secure Firewall devices, the **Drive State** will show as **degraded**.



Caution Do not remove an SSD without first removing it from the RAID using this procedure. You can cause data loss.

Procedure

Step 1 Remove one of the SSDs.

a) Remove the SSD from the RAID.

configure raid remove-secure local-disk {**1** | **2**}

The **remove-secure** keyword removes the SSD from the RAID, disables the self-encrypting disk feature, and performs a secure erase of the SSD. If you only want to remove the SSD from the RAID and want to keep the data intact, you can use the **remove** keyword.

Example:

```
> configure raid remove-secure local-disk 2
```

b) Monitor the RAID status until the SSD no longer shows in the inventory.

show raid

After the SSD is removed from the RAID, the **Operability** and **Drive State** will show as **degraded**. The second drive will no longer be listed as a member disk.

Example:

```
> show raid
Virtual Drive
ID:                     1
Size (MB):              858306
Operability:            operable
Presence:               equipped
Lifecycle:              available
Drive State:            optimal
Type:                   raid
Level:                  raid1
Max Disks:              2
Meta Version:           1.0
Array State:            active
Sync Action:            idle
Sync Completed:         unknown
Degraded:               0
Sync Speed:             none
```

```
RAID member Disk:
Device Name:                nvme0n1
Disk State:                 in-sync
Disk Slot:                  1
Read Errors:                0
Recovery Start:             none
Bad Blocks:
Unacknowledged Bad Blocks:

Device Name:                nvme1n1
Disk State:                 in-sync
Disk Slot:                  2
Read Errors:                0
Recovery Start:             none
Bad Blocks:
Unacknowledged Bad Blocks:

> show raid
Virtual Drive
ID:                         1
Size (MB):                  858306
Operability:                degraded
Presence:                   equipped
Lifecycle:                  available
Drive State:                degraded
Type:                       raid
Level:                      raid1
Max Disks:                  2
Meta Version:               1.0
Array State:                active
Sync Action:                idle
Sync Completed:             unknown
Degraded:                   1
Sync Speed:                 none

RAID member Disk:
Device Name:                nvme0n1
Disk State:                 in-sync
Disk Slot:                  1
Read Errors:                0
Recovery Start:             none
Bad Blocks:
Unacknowledged Bad Blocks:
```

c) Physically remove the SSD from the chassis.

Step 2 Add an SSD.

a) Physically add the SSD to the empty slot.

b) Add the SSD to the RAID.

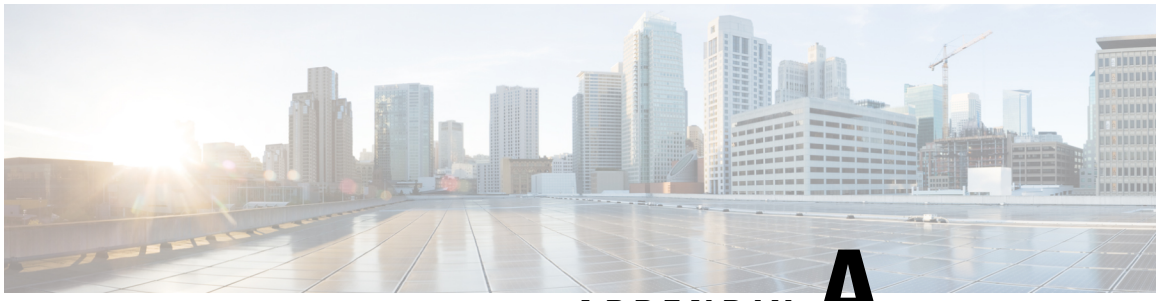
configure raid add local-disk {**1** | **2**}

It can take several hours to complete syncing the new SSD to the RAID, during which the firewall is completely operational. You can even reboot, and the sync will continue after it powers up. Use the **show raid** command to show the status.

If you install an SSD that was previously used on another system, and is still locked, enter the following command:

configure raid add local-disk {**1** | **2**} *psid*

The *psid* is printed on the label attached to the back of the SSD. Alternatively, you can reboot the system, and the SSD will be reformatted and added to the RAID.



APPENDIX A

Advanced Configuration

Some device features are configured using ASA configuration commands. Although the Firewall Device Manager can configure many command-based features, it does not support all of them. If you need to use some of these ASA features that are not otherwise supported in the Firewall Device Manager, you can use Smart CLI or FlexConfig to manually configure the features.

The following topics explain this type of advanced configuration in more detail.

- About Smart CLI and FlexConfig, on page 831
- Guidelines and Limitations for Smart CLI and FlexConfig, on page 840
- Configuring Smart CLI Objects, on page 840
- Configuring the FlexConfig Policy, on page 841
- Troubleshooting the FlexConfig Policy, on page 853
- Examples for FlexConfig, on page 854

About Smart CLI and FlexConfig

Firewall Threat Defense uses ASA configuration commands to implement some features, but not all features. There is no unique set of the Firewall Threat Defense configuration commands.

You can configure features using the CLI using the following methods:

- **Smart CLI**—(Preferred method.) A Smart CLI template is a pre-defined template for a particular feature. All of the commands needed for the feature are provided, and you simply need to select values for variables. The system validates your selection, so that you are more likely to configure a feature correctly. If a Smart CLI template exists for the feature you want, you must use this method.
- **FlexConfig**—The FlexConfig policy is a collection of FlexConfig objects. The FlexConfig objects are more free-form than Smart CLI templates, and the system does no CLI, variable, or data validation. You must know ASA configuration commands and follow the ASA configuration guides to create a valid sequence of commands.

The point of Smart CLI and FlexConfig is to allow you to configure features that are not directly supported through Firewall Device Manager policies and settings.

Caution Cisco strongly recommends using Smart CLI and FlexConfig 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 Smart CLI and FlexConfig may cause unintended results with other configured features.

You may contact the Cisco Technical Assistance Center for support concerning Smart CLI and FlexConfig objects 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 Firewall Threat Defense features. Smart CLI and FlexConfig features may become deprecated at any time. For fully guaranteed feature support, you must wait for the Firewall Device Manager support. When in doubt, do not use Smart CLI or FlexConfig.

The following topics explain these features in more detail.

Recommended Usage for Smart CLI and FlexConfig

There are two main recommended uses for FlexConfig:

- You are migrating from ASA to Firewall Threat Defense, and there are compatible features you are using (and need to continue using) that the Firewall Device Manager 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. Verify by comparing **show running-config** output on the two devices.
- You are using the Firewall 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.

Before trying to recreate an ASA configuration, 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 Firewall Threat Defense. Do not attempt to completely recreate an ASA configuration on the Firewall Threat Defense device. You must carefully test any feature that you configure using FlexConfig.

CLI Commands in Smart CLI and FlexConfig Objects

The Firewall Threat Defense uses ASA configuration commands to configure some features. Although not all ASA features are compatible with Firewall Threat Defense, there are some features that can work on the Firewall Threat Defense but that you cannot configure in the Firewall Device Manager policies. You can use Smart CLI and FlexConfig objects to specify the CLI required to configure these features.

If you decide to use Smart CLI or FlexConfig to manually configure a feature, you are responsible for knowing and implementing the commands according to the proper syntax. FlexConfig does 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.

How Software Upgrades Affect the FlexConfig Policy

Each new version of the Firewall Threat Defense software adds support for configuring features in the Firewall Threat Defense. Sometimes, these new features can overlap with features you have previously configured using FlexConfig.

After upgrade, you need to examine the FlexConfig policy and objects. If any contain commands that have become prohibited due to added support within Firewall Device Manager or Smart CLI, icons in the objects list and messages indicate the problem. Please take the time to redo your configuration. Use the list of prohibited commands for help in determining where the commands should now be configured.

The system will not prevent you from deploying changes while the FlexConfig objects that are attached to the FlexConfig policy contain newly-prohibited commands. However, you will be prevented from creating new Smart CLI objects until you resolve all issues noted in the FlexConfig policy.

You can simply remove the problematic objects from the FlexConfig policy, as the restriction applies only to those objects you are actively deploying to the device configuration. Thus, you can remove the objects, then use them as a reference as you create the corresponding Smart CLI or integrated the Firewall Device Manager configuration. Once you are satisfied with the new configuration, you can simply delete the objects. If the removed objects contain some non-prohibited elements, you can edit them to remove the unsupported commands, and then reattach the objects to the FlexConfig policy.

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 Firewall 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 the Firewall Threat Defense configuration. Many Firewall 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 Firewall Threat Defense configuration.

To view this information, either open the CLI Console in the Firewall Device Manager or 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.
- **show running-config** to view the current CLI configuration.
- **show running-config all** to include all the default commands in the current CLI configuration.

Prohibited CLI Commands

The purpose of Smart CLI and FlexConfig is to configure features that are available on ASA devices that you cannot configure on the Firewall Threat Defense devices using the Firewall Device Manager.

Thus, you are prevented from configuring ASA features that have equivalents in the Firewall Device Manager. The following table lists some of these prohibited command areas. This list contains many parent commands that enter configuration modes. The prohibition of the parent includes the prohibition of the children commands. It also includes the **no** version of the commands and their associated **clear** commands.

The FlexConfig object editor prevents you from including these commands in the object. This list does not apply to Smart CLI templates, as they include only those commands you can validly configure.

Prohibited CLI Command	Comments
aaa	Use **Objects** > **Identity Sources**.
aaa-server	Use **Objects** > **Identity Sources**.
access-list	Partially blocked. • You can create **ethertype** access lists. • You cannot create **extended** and **standard** access lists. Create these ACLs using the Smart CLI Extended Access List or Standard Access List objects. You can then use them on FlexConfig-supported commands that refer to the ACL by object name, such as **match access-list** with an extended ACL for service policy traffic classes. • You cannot create **advanced** access lists, which the system uses with the **access-group** command. Instead, use **Policies** > **Access Control** to configure access rules. • You cannot create **webtype** access lists.
anyconnect-custom-data	Use **Device** > **Remote Access VPN** to configure AnyConnect Client.
asdm	This feature does not apply to a Firewall Threat Defense system.
as-path	Create Smart CLI AS Path objects and use them in a Smart CLI BGP object to configure an autonomous system path filter.
attribute	—
auth-prompt	This feature does not apply to a Firewall Threat Defense system.
boot	—
call-home	—
captive-portal	Use **Policies** > **Identity** to configure the captive portal used for active authentication.
clear	—

Prohibited CLI Command	Comments
client-update	—
clock	Use **Device** > **System Settings** > **NTP** to configure system time.
cluster	—
command-alias	—
community-list	Create Smart CLI Expanded Community List or Standard Community List objects and use them in a Smart CLI BGP object to configure a community list filter.
compression	—
configure	—
crypto	On the **Objects** page, use **Certificates**, **IKE Policies**, and **IPSec Proposals**.
ddns	Use **Device** > **System Settings** > **DDNS Service** to configure dynamic DNS.
dhcp-client	—
dhcpd	Use **Device** > **System Settings** > **DHCP Server**. However, the **dhcpd option** command is allowed.
dhcprelay	Use the dhcprelayservices resource in the threat defense API instead.
dns	Configure DNS groups using **Objects** > **DNS Groups**, and assign the groups using **Device** > **System Settings** > **DNS Server**.
dns-group	Configure DNS groups using **Objects** > **DNS Groups**, and assign the groups using **Device** > **System Settings** > **DNS Server**.
domain-name	Configure DNS groups using **Objects** > **DNS Groups**, and assign the groups using **Device** > **System Settings** > **DNS Server**.
dynamic-access-policy-config **dynamic-access-policy-record**	—
enable	—
event	—
failover	—
fips	—
firewall	Firewall Device Manager supports routed firewall mode only.
hostname	Use **Device** > **System Settings** > **Hostname**.

Prohibited CLI Command	Comments
hpm	This feature does not apply to a Firewall Threat Defense system.
http	Use the **Data Interfaces** tab on **Device** > **System Settings** > **Management Access**.
inline-set	—
interface for BVI, Management, Ethernet, GigabitEthernet, and subinterfaces.	Partially blocked. Configure physical interfaces, subinterfaces, and Bridge Virtual Interfaces on the **Device** > **Interfaces** page. You can then configure additional options using FlexConfig. However, the following **interface** mode commands are prohibited for these types of interface. **cts** **ip address** **ip address dhcp** **ipv6 address** **ipv6 enable** **ipv6 nd dad** **ipv6 nd suppress-ra** **mode** **nameif** **security-level** **shutdown** **zone-member**
interface for **vni**, **redundant**, **tunnel**	Configure interfaces on the **Device** > **Interfaces** page. Firewall Device Manager does not support these types of interface.
ip audit	This feature does not apply to a Firewall Threat Defense system. Instead, apply intrusion policies using access control rules.
ip-client	To configure the system to use data interfaces as the management gateway, use **Device** > **System Settings** > **Management Interface**.
ip local pool	Use **Device** > **Remote Access VPN** to configure address pools.
ipsec	—
ipv6	Create Smart CLI IPv6 Prefix List objects and use them in a Smart CLI BGP object to configure prefix list filtering for IPv6.
ipv6-vpn-addr-assign	Use **Device** > **Remote Access VPN** to configure address pools.
isakmp	Use **Device** > **Site-to-Site VPN**.
jumbo-frame	The system automatically enables jumbo frame support if you increase the MTU of any interface over the default 1500.

Prohibited CLI Command	Comments
ldap	—
license-server	Use **Device** > **Smart License**.
logging	Use **Objects** > **Syslog Servers** and **Device** > **System Settings** > **Logging Settings**. However, you can configure the **logging history** command in FlexConfig.
management-access	—
migrate	Use **Device** > **Remote Access VPN** and **Device** > **Site-to-Site VPN** to enable IKEv2 support.
mode	Firewall Device Manager supports single context mode only.
mount	—
mtu	Configure MTU per interface on **Device** > **Interfaces**.
nat	Use **Policies** > **NAT**.
ngips	—
ntp	Use **Device** > **System Settings** > **NTP**
object-group network **object network**	Use **Objects** > **Network**. You cannot create network objects or groups in FlexConfig, but you can use network objects and groups defined in the object manager inside the template as variables.
object service |natorigsvc **object service |natmappedsvc**	The **object service** command is allowed in general, but you cannot edit the internal objects named |natorigsvc or |natmappedsvc. In these names, the vertical bar is intentional and it is the first character of the restricted object names.
passwd **password**	—
password-policy	—
policy-list	Create Smart CLI Policy List objects and use them in a Smart CLI BGP object to configure a policy list.
policy-map sub-commands	You cannot configure the following commands in a policy map. **priority** **police** **match tunnel-group**

Prohibited CLI Command	Comments
prefix-list	Create Smart CLI IPv4 Prefix List objects and use them in a Smart CLI OSPF or BGP object to configure prefix list filtering for IPv4.
priority-queue	—
privilege	—
reload	You cannot schedule reloads. The system does not use the **reload** command to restart the system, it uses the **reboot** command.
rest-api	This feature does not apply to a Firewall Threat Defense system. The REST API is always installed and enabled.
route	Use **Device** > **Routing** to configure static routes.
route-map	Create Smart CLI Route Map objects and use them in a Smart CLI OSPF or BGP object to configure route maps.
router bgp	Use the Smart CLI templates for BGP.
router eigrp	Use the Smart CLI templates for EIGRP.
router ospf	Use the Smart CLI templates for OSPF.
scansafe	This feature does not apply to a Firewall Threat Defense system. Instead, configure URL filtering in access control rules.
setup	This feature does not apply to a Firewall Threat Defense system.
sla	—
snmp-server	Use the FTP API SNMP resources to configure SNMP.
ssh	Use the **Data Interfaces** tab on **Device** > **System Settings** > **Management Access**.
ssl	Use **Device** > **System Settings** > **SSL Settings**.
telnet	Firewall Threat Defense does not support Telnet connections. Use SSH instead of Telnet to access the device CLI.
time-range	—
tunnel-group	Use **Device** > **Remote Access VPN** and **Device** > **Site-to-Site VPN**.
tunnel-group-map	Use **Device** > **Remote Access VPN** and **Device** > **Site-to-Site VPN**.
user-identity	Use **Policies** > **Identity**.
username	To create CLI users, open an SSH or console session to the device and use the **configure user** commands.
vpdn	—

Prohibited CLI Command	Comments
vpn	—
vpn-addr-assign	—
vpnclient	—
vpn-sessiondb	—
vpnsetup	—
webvpn	—
zone	—
zonelabs-integrity	This feature does not apply to a Firewall Threat Defense system.

Smart CLI Templates

The following table explains the Smart CLI templates based on the feature.



Note You also configure OSPF and BGP using Smart CLI templates. However, those templates are available through the **Device** > **Routing** page rather than the Advanced Configuration page.

Feature	Templates	Description
Objects: AS Path	ASPath	Create ASPath objects for use with routing protocol objects.
Objects: Access List	Extended Access List Standard Access List	Create extended or standard ACLs for use with routing objects. You can also refer to these objects by name from FlexConfig objects that configure permitted commands that use ACLs.
Objects: Community List	Expanded Community List Standard Community List	Create expanded or standard community lists for use with routing objects.
Objects: Prefix List	IPV4 Prefix List IPV6 Prefix List	Create IPv4 or IPv6 prefix lists for use with routing objects.
Objects: Policy List	Policy List	Create policy lists for use with routing objects.
Objects: Route Map	Route Map	Create route maps for use with routing objects.

Guidelines and Limitations for Smart CLI and FlexConfig

Please keep the following in mind when configuring features through Smart CLI or FlexConfig.

- The commands defined in FlexConfig objects are deployed after all commands for features defined through Firewall Device Manager, including Smart CLI. Thus, you can depend on objects, interfaces, and so forth being configured before these commands are issued to the device. If you need to use a FlexConfig-deployed item in a Smart CLI template, create and deploy the FlexConfig before creating and deploying the Smart CLI template. For example, if you want to use the OSPF Smart CLI template to redistribute EIGRP routes, first use FlexConfig to configure EIGRP, then create the OSPF Smart CLI template.
- If you want to remove a feature or part of a feature that you configured through FlexConfig, but a Smart CLI template refers to that feature, you must first remove the commands in the Smart CLI template that use the feature. Then, deploy the configuration so that the Smart-CLI configured feature no longer refers to it. You can then remove the feature from FlexConfig and redeploy the configuration to finally eliminate it altogether.

Configuring Smart CLI Objects

Smart CLI objects define features that cannot be configured elsewhere in the Firewall Device Manager. Smart CLI objects provide a level of guidance in configuring a feature. For a given feature (template), all possible commands are pre-loaded, and the variables you enter are validated. Thus, although you still use CLI commands to configure a feature, Smart CLI objects are not as free-form as FlexConfig objects.

Although Smart CLI templates do provide a level of guidance, you must still read the ASA configuration guides and command references to understand the command usage so that you pick values that work correctly for your network. Ideally, you already have an ASA configuration to work from, and you merely need to build the same sequence of commands in the Smart CLI object.

Smart CLI objects are grouped according to feature area.



Note All Smart CLI objects that you define are deployed. Unlike FlexConfig, you cannot create several Smart CLI objects and then select which of them to deploy. Create Smart CLI objects only for those features you want to configure.

Procedure

Step 1 Click **View Configuration** in **Device** > **Advanced Configuration**.

Step 2 Click the appropriate feature area under **Smart CLI** in the Advanced Configuration table of contents.

Step 3 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon () for the object.

To delete an object, click the trash can icon ( ) for the object.

Step 4 Enter a Name for the object and optionally, a description.

Step 5 Select the **CLI Template** for the feature you are configuring.

The system loads the command template into the **Template** window. Initially, only the required commands are shown. These represent the minimum configuration required for the template.

Step 6 Fill in the variables and add commands as needed in the template.

Ideally, you are working with an existing configuration from an ASA or Firewall Threat Defense device (one that is managed by the Firewall Management Center). With a configuration in hand, you simply need to make the template conform to it, changing variables such as IP addresses and interface names as appropriate for the location of this specific device in your network.

Following are some tips for filling in the template:

- To select a value for a variable, click the variable and either type in the appropriate value, or select it from a list (in the case of enumerated values). Mousing over variables that require typing shows the valid values for the option, such as a range of numbers. In some cases, the recommended value is mentioned.

  For example, in the OSPF template, the required command **router ospf** *process-id* shows "Process ID (1-65535)" on mouse-over, and when you click *process-id*, the field is highlighted. Simply type in the number you want.

- When you select an option for a variable, if there are additional possible commands to configure the option, these are automatically exposed and disabled or enabled as appropriate. Watch for these additional commands.

- Use the **Show**/**Hide Disabled** link above the template to control your view. Disabled commands will not be configured, but you must display them to configure them. To see the full template, click the **Show Disabled** link above the template. To see only those commands that will be configured, click the **Hide Disabled** link above the table.

- To clear all of your edits since you last saved the object, click the **Reset** link above the template.

- To enable an optional command, click the + button to the left of the line number.

- To disable an optional command, click the - button to the left of the line number. If you edited the line, your edits are not deleted.

- To duplicate a command, click the Options **...** button and select **Duplicate**. You are allowed to duplicate commands only if it is valid to enter the command more than once.

- To delete a duplicated command, click the Options **...** button and select **Delete**. You cannot delete the commands that are a part of the base template.

Step 7 Click **OK**.

Configuring the FlexConfig Policy

The FlexConfig Policy is simply a list of the FlexConfig objects that you want to deploy to the device configuration. Only those objects included in the policy are deployed, all others are simply defined and unused.

The commands defined in FlexConfig objects are deployed after all commands for features defined through Firewall Device Manager, including Smart CLI. Thus, you can depend on objects, interfaces, and so forth being configured before these commands are issued to the device. If you need to use a FlexConfig-deployed item in a Smart CLI template, create and deploy the FlexConfig before creating and deploying the Smart CLI template. For example, if you want to use the OSPF Smart CLI template to redistribute EIGRP routes, first use FlexConfig to configure EIGRP, then create the OSPF Smart CLI template.



Note If there is a Smart CLI template for a feature, you cannot configure it using FlexConfig. You must use the Smart CLI object.

Before you begin

Create the FlexConfig objects. See the following topics:

- Configuring FlexConfig Objects, on page 843
- Creating Variables in a FlexConfig Object, on page 845
- Configuring Secret Key Objects, on page 852

Procedure

Step 1 Click **View Configuration** in **Device** > **Advanced Configuration**.

Step 2 Click **FlexConfig** > **FlexConfig Policy** in the Advanced Configuration table of contents.

Step 3 Manage the list of objects in the **Group List**.

- To add an object, click the **+** button. If the object does not yet exist, click **Create New FlexConfig Object** to define it.
- To delete an object, click the **X** button at the right of the object entry.

Note
We recommend that each object be completely self-contained and not depend on the configuration defined in any other FlexConfig object. This ensures that you can add or remove objects without affecting other objects.

Step 4 Evaluate the proposed commands in the **Preview** pane.

You can click the **Expand** button (and subsequently, **Collapse**) to widen the screen so you can see long commands more clearly.

The preview evaluates variables and produces the exact commands that will be issued. Ensure that these commands are correct and valid. You are responsible for ensuring the commands will not result in errors or poor configurations that make the device unusable.

Caution
The system does not validate the commands. It is possible for you to deploy invalid and even potentially destructive commands. Examine the preview very carefully before deploying changes.

Step 5 Click **Save**.

What to do next

After editing the FlexConfig policy, carefully examine the results of the next deployment. If there are errors, correct the CLI in the object. See Troubleshooting the FlexConfig Policy, on page 853.

Configuring FlexConfig Objects

A FlexConfig object contains the ASA commands required to configure a particular feature that you cannot otherwise configure using the Firewall Device Manager. You are responsible for ensuring that you enter the right sequence of commands, without typos. The system does not validate the content of FlexConfig objects.

We recommend that you create separate objects for each general feature you intend to configure. For example, if you want to define banners and also configure the RIP routing protocol, use 2 separate objects. Isolating features in separate objects makes it easier for you to pick and choose which to objects to deploy, and also makes troubleshooting more straight-forward.

Note Do not include the **enable** and **configure terminal** commands. The system enters the right mode for configuration commands automatically.

Procedure

Step 1 Click **View Configuration** in **Device** > **Advanced Configuration**.

Step 2 Click **FlexConfig** > **FlexConfig Objects** in the Advanced Configuration table of contents.

Step 3 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 4 Enter a Name for the object and optionally, a description.

Step 5 In the **Variables** section, create any variables that you want to use in the object body.

The only variables that you must create are those that point to objects defined within the Firewall Device Manager, specifically, the Network, Port, and Secret Key variable types, or the Interface variable, which points to a named interface. For other variable types, you can simply type in the values into the object body.

For detailed information on creating and using variables, see Creating Variables in a FlexConfig Object, on page 845.

Step 6 In the **Template** section, type in the ASA commands required to configure the feature.

You must enter commands in the right order for configuring the feature. Use the ASA CLI configuration guides to learn how to enter the commands. Ideally, you should have a pre-tested configuration file from an ASA or another Firewall Threat Defense device that you can use as a reference.

You can also use Mustache notation to refer to and process variables. For detailed information, see Referring to FlexConfig Variables and Retrieving Values, on page 846.

Following are some tips for creating the object body:

- To add lines, put the cursor at the end of a line and press Enter.
- To use a variable, type the variable name between double-braces: {{*variable_name*}}. For variables that refer to objects, you must include the attribute whose value you are retrieving: {{*variable_name.attribute*}}. The available attributes differ based on object type. For complete information, see Variable References: {{variable}} or {{{variable}}}, on page 846.
- To use a Smart CLI object, type the name of the object. If you need to refer to a routing process configured in Smart CLI, enter the process identifier. See Referring to Smart CLI Objects in a FlexConfig Object, on page 851.
- Click the **Expand/Collapse** link above the template body to make the body larger or smaller.
- Click the **Reset** link to erase any changes you made since you last saved the object.

Step 7 In the **Negate Template** section, enter the commands needed to remove or reverse the commands configured in the object body.

The Negate section is very important and serves two purposes:

- It simplifies deployment. Before re-deploying the commands in the body, the system uses these commands to first erase or undo the configuration. This ensures a clean deployment.
- If you decide to remove the feature by removing the object from the FlexConfig Policy, the system uses these commands to remove the commands from the device.

If you do not supply the commands needed to negate or reverse the CLI in the object body, the deployment might need to clear the entire device configuration and redeploy all policies, not just the commands within the object. This will make deployment take longer and also disrupt traffic. Ensure that you have all, and only, those commands needed to undo the configuration defined in the object body. Although negate commands are typically the **no** or **clear** form of the commands in the template, if you are actually turning off a feature that was already enabled, the “negate” command is actually the positive form of the command, the one that enables the feature.

Use the ASA configuration guides and command reference to determine the appropriate commands. Sometimes, you can undo a configuration with a single command. For example, in an object that configures RIP, a simple **no router rip** command removes the entire **router rip** configuration, including subcommands.

Likewise, if you entered several **banner login** commands to create a multi-line banner, a single **no banner login** command negates the entire login banner.

If your template creates several nested objects, the negate template needs to remove the objects in reverse order, to first remove references to the objects before deleting the objects. For example, if you first create an ACL, then refer to it in a traffic class, then refer to the traffic class in a policy map, and finally enable the policy map using a service policy, the negate template must undo the configuration by first removing the service policy, then the policy map, then the traffic class, and finally the ACL.

Step 8 Click **OK**.

What to do next

Simply creating a FlexConfig object is not enough to get it deployed. You must add the object to the FlexConfig Policy. Only those objects in the FlexConfig policy get deployed. This makes it possible for you to refine your FlexConfig objects, and have some ready for special uses, without having all of them automatically deployed. See Configuring the FlexConfig Policy, on page 841.

Creating Variables in a FlexConfig Object

The variables you use inside a FlexConfig object are defined within the object itself. There is no separate list of variables. Thus, you cannot define a variable and then use it within separate FlexConfig objects.

Variables provide these main benefits:

- They make it possible to point to objects defined using the Firewall Device Manager. This includes network, port, and secret key objects.
- They isolate values that might change from the object body. Thus, if you need to change a value, you simply edit the variable and you do not need to edit the object body. This can be especially helpful if you need to refer to the object in several command lines.

This procedure explains the process of adding variables to a FlexConfig object.

Procedure

Step 1 Edit or create a FlexConfig object from the **Device** > **Advanced Configuration** page.

See Configuring FlexConfig Objects, on page 843.

Step 2 Do one of the following in the **Variables** section:

- To add a variable, click the + button (or click **Add Variable** if there are none yet defined).
- To edit a variable, click the edit icon ( ) for the variable.

To delete a variable, click the trash can ( ) icon for the variable. Make sure you remove any references to it from the template body.

Step 3 Enter a Name for the variable and optionally, a description.

Step 4 Select a data **Type** for the variable, then enter or select the value.

You can create the following types of variable. Choose a type that fits the data requirements of the commands in which you will use the variable.

- **String**—A text string. For example, hostnames, usernames, and so forth.
- **Numeric**—An integer number. Do not include commas, decimals, signs (such as negative), or hexadecimal notation. For non-integer numbers, use a string variable.
- **Boolean**—A logical true/false. Select either True or False.

- **Network**—A network object or group defined on the Objects page. Select the network object or group.
- **Port**—A TCP or UDP port object defined on the Objects page. Select the port object. You cannot select groups or objects for other protocols.
- **Interface**—A named interface defined on the Device > Interfaces page. Select the interface. You cannot select interfaces that have no names.
- **IP**—A single IPv4 or IPv6 IP address without netmask or prefix length.
- **Secret**—A secret key object defined for FlexConfig. Select the object. For information on creating secret key objects, see Configuring Secret Key Objects, on page 852.

Step 5 Click **Add** or **Save** in the Variable dialog box.

You can now use the variable within the body of the FlexConfig object. The way you refer to the variable differs based on variable type. For details on how to use these variables, see the following topics:

- Variable References: {{variable}} or {{{variable}}}, on page 846
- Sections {{#key}} {{/key}} and Inverse Sections {{^key}} {{/key}}, on page 849

Step 6 Click **OK** in the FlexConfig Object dialog box.

Referring to FlexConfig Variables and Retrieving Values

FlexConfig uses Mustache as the template language, but support is limited to the features explained in the following sections. Use these features to refer to variables, retrieve their values, and process them.

Variable References: {{variable}} or {{{variable}}}

To refer to a variable, which you define within a FlexConfig object, you use the following notation:

{{*variable_name*}}

Or:

{{{*variable_name*}}}

This is sufficient for simple variables that are single values, which includes variables of the following types: **Numeric**, **String**, **Boolean**, and **IP**. Use triple braces if the variable contains special characters such as &. Alternatively, you can always use triple braces for all variables.

However, for variables that point to elements that are modeled as objects in the configuration database, you must use dot notation and include the name of the object attribute you want to retrieve. You can find these attribute names by examining the models in the API Explorer for the related object type. You must use the following notation to use variables of the following types: **Secret**, **Network**, **Port**, and **Interface**.

{{*variable_name.attribute*}}

For example, to retrieve the address from a network variable named net-object1 (which points to a network object, not a network group), you would use:

{{net-object1.value}}

If you are trying to retrieve an attribute value from an object within an object, you need to use a series of dotted attributes to drill down to the desired value. For example, the IP addresses for an interface are modeled as sub-objects, named ipv4 and ipv6, to the interface object. Thus, to retrieve the IPv4 address and subnet mask for an interface variable named int-inside (which points to the inside interface), you would use:

{{int-inside.ipv4.ipAddress.ipAddress}} {{int-inside.ipv4.ipAddress.netmask}}



Note To open API Explorer, click the more options button (⋮) and choose **API Explorer**.

The following table lists the variable types, how to refer to them, and for objects, the name of the API model and the most likely references that you might use.

Variable Type	Reference Models	Description
Boolean (simple variable)	**Variable:** {{*variable_name*}} **Section:** `{{#variable_name}}` `commands` `{{/variable_name}}` **Inverse Section:** `{{^variable_name}}` `commands` `{{/variable_name}}`	A logical true/false. The main purpose for Boolean variables is for sections or inverse sections. You can edit the value of a Boolean variable to turn a section of commands on or off, for example, if you need to enable a feature periodically or under special circumstances only. Some objects also have Boolean attributes in their models, which you can use to provide optional processing of a section.

Variable Type	Reference Models	Description
Interface (object variable: API model is Interface)	**Variable:** {{*variable_name*.*attribute*}} **Section:** `{{#variable_name.attribute}}` `commands` `{{/variable_name.attribute}}` **Inverse Section:** `{{^variable_name.attribute}}` `commands` `{{/variable_name.attribute}}`	A named interface defined on the Device > Interfaces page. You cannot point to unnamed interfaces. There is a wide variety of attributes available in the interface model. Also, the interface model includes sub-objects, for example, for IP addresses. Following are some of the main attributes that you might find useful: • *variable_name*.**name** returns the logical name of the interface. • *variable_name*.**hardwareName** returns the interface port name, such as GigabitEthernet1/8. • *variable_name*.**managementOnly** is a Boolean value. TRUE means that the interface is defined as management only. FALSE means the interface is for through-the-device traffic. You could use this option as a section key. • *variable_name*.**ipv4.ipAddress.ipAddress** returns the IPv4 address for the interface. • *variable_name*.**ipv4.ipAddress.netmask** returns the subnet mask for the IPv4 address for the interface.
IP (simple variable)	**Variable:** {{*variable_name*}}	A single IPv4 or IPv6 IP address without netmask or prefix length.
Network (object variable: API model is NetworkObject)	**Variable (Network Objects):** {{*variable_name*.*attribute*}} **Section (Group Objects):** `{{#variable_name.networkObjects}}` `commands referring to one of` `  {{value}}` `  {{name}}` `{{/variable_name.networkObjects}}`	A network object or group defined on the Objects page. You can use sections to process network groups. Following are the main attributes that you might find useful: • {{*variable_name*.**name**}} returns the name of the network object or group. • {{*variable_name*.**value**}} returns the IP address contents of a network object (but not a network group). Ensure that you use a network object that has the right type of contents for a given command, for example, a host address rather than a subnet address. • {{*variable_name*.**groups**}} returns the list of network objects contained within a network group. Use this only with variables that point to network groups, and use it on a section tag to iteratively process the contents of the group. Use either {{**value**}} or {{**name**}} to retrieve the contents of each network object in turn.

Variable Type	Reference Models	Description
Numeric (simple variable)	**Variable:** {{*variable_name*}}	An integer number. Do not include commas, decimals, signs (such as negative), or hexadecimal notation. For non-integer numbers, use a string variable.
Port (object variable: API model is PortObject, tcpports or udpports)	**Variable:** {{*variable_name.attribute*}}	A TCP or UDP port object defined on the Objects page. This must be a port object, not a port group. Following are the main attributes that you might find useful: • {{*variable_name*.**port**}} returns the port number. The protocol is not included. • {{*variable_name*.**name**}} returns the name of the port object.
Secret (object variable: API model is Secret)	**Variable:** {{*variable_name*.**password**}} Or: {{{*variable_name*.**password**}}}	A secret key object defined for FlexConfig. The only reference you should make is to the **password** attribute, which returns the encrypted string. If the password includes special characters such as &, use triple braces.
String (simple variable)	**Variable:** {{*variable_name*}}	A text string. For example, hostnames, usernames, and so forth.

Sections {{#key}} {{/key}} and Inverse Sections {{^key}} {{/key}}

A section or an inverse section is a block of commands between the section start and end tags, which use a key as the processing criteria. How the section is processed depends on whether it is a regular or an inverse section:

- A regular section (or simply, a section) is processed if the key is TRUE or has non-empty contents. If the key is FALSE or the object has no content, the commands in the section are not configured. The section is bypassed.

  The following is the syntax for a regular section.

```
{{#key}}
one or more commands
{{/key}}
```

- An inverse section is the opposite of a section. It is processed if the key is FALSE or the object has no contents. If the key is TRUE or the object has contents, the inverse section is bypassed.

  The following is the syntax for an inverse section. The only difference is that a caret replaces the hash tag.

```
{{^key}}
one or more commands
{{/key}}
```

The following topics explain the main uses for sections and inverse sections.

How to Process Multiple-Value Variables

The primary example of processing a multiple-value variable is a network variable that points to a network group. Because the group contains multiple objects (under the **objects** attribute), you can iteratively go through the values in the network group to configure the same command multiple times with different values.

Although an object group defines the contained network objects within the objects attribute, those objects do not include the contents of the contained objects. Instead, you use the **networkObjects** attribute to get at the contents of the contained objects.

For example, if you have a network group named net-group with the hosts 192.168.10.0, 192.168.20.0, and 192.168.30.0, you can use the following technique to configure a network command for each address for RIP routing. Note that you use the network object's **value** attribute alone, because the use of **net-group.networkObjects** in the section start implies that the value attribute will be taken from the member objects. (You do not create a separate variable for the “value” attribute within the FlexConfig object.)

```
router rip
{{#net-group.networkObjects}}
 network {{value}}
{{/net-group.networkObjects}}
```

The system translates the section structure as:

```
router rip
   network 192.168.10.0
   network 192.168.20.0
   network 192.168.30.0
```

How to Perform Optional Processing Based on a Boolean Value or an Empty Object



Note The examples in this topic are for illustration purposes only. For example, you cannot use FlexConfig to configure SNMP starting with version 6.7; you must use the Firewall Threat Defense API SNMP resource instead.

If the content of the variable in the section start tag is TRUE, or an object is not empty, the section is processed. If a Boolean value is FALSE or empty (such as an empty object), the section is bypassed.

The main use here is for Boolean values. For example, you could create a Boolean variable, and put commands within a section covered by the variable. Then, if you need to enable or disable a section of the commands in the FlexConfig object, you merely need to change the value of the Boolean variable, you do not need to delete those lines from the code. This makes it easy to turn features on or off.

For example, you might want to be able to turn off SNMP traps if you use FlexConfig to enable SNMP. You could create a Boolean variable named enable-traps, and initially set it to TRUE. Then, if you need to turn off traps, you simply edit the variable, change it to FALSE, save the object and then redeploy the configuration. The command sequence could look like the following:

```
snmp-server enable
snmp-server host inside 192.168.1.5
snmp-server community clearTextString
```

```
{{#enable-traps}}
snmp-server enable traps all
{{/enable-traps}}
```

You can also do this type of processing based on Boolean values within an object. For example, you could check whether an interface is management-only before configuring some characteristic on it. In the following example, int-inside is an interface variable that points to the interface named inside. The FlexConfig configures the EIGRP-related interface options on the interface only if the interface is not set to management only. You would use an inverse section so that the commands are configured only if the Boolean value is FALSE.

```
router eigrp 2
 network 192.168.1.0 255.255.255.0
{{^int-inside.managementOnly}}
interface {{int-inside.hardwareName}}
 hello interval eigrp 2 60
 delay 200
{{/int-inside.managementOnly}}
```

Referring to Smart CLI Objects in a FlexConfig Object

When you create a FlexConfig object, you can use variables to point to objects that you can configure within the Firewall Device Manager. For example, you can create variables that point to interface elements or network objects.

However, you cannot point to Smart CLI objects in the same way.

Instead, if you create a Smart CLI object that you need to use in a FlexConfig policy, you simply type in the name of the Smart CLI object at the appropriate location.

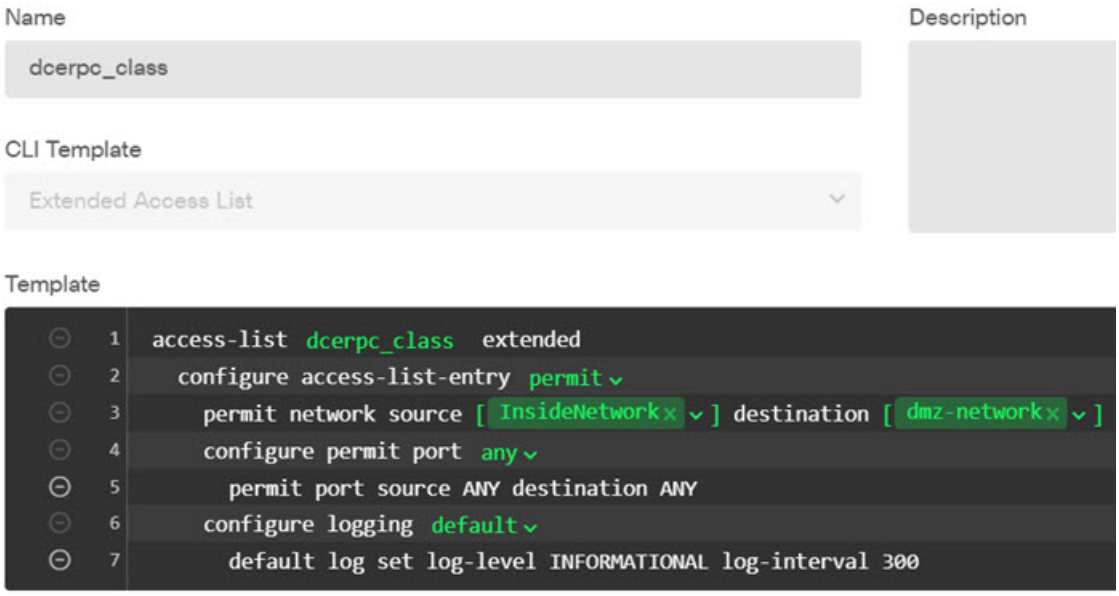
For example, you might want to use an extended access list as the traffic class when you configure protocol inspection. Because there is a Smart CLI object for extended access lists, you need to use the Smart CLI object to create the ACL: you cannot use the **access-list** command in the FlexConfig object.

As an example, if you wanted to enable DCERPC inspection between networks 192.168.1.0/24 and 192.168.2.0/24 globally, you would do the following.

Procedure

Step 1 Create separate network objects for the two networks. For example, InsideNetwork and dmz-network.

Step 2 Use these objects in a Smart CLI extended access list object.



Step 3 Create a FlexConfig object that points to the Smart CLI object by name.

For example, if the object is named "dcerpc_class," your FlexConfig object might look like the following. Note that in the negate template, you do not negate the access list created through the Smart CLI object, as that object is not actually created through FlexConfig.



Step 4 Add the object to the FlexConfig policy.

Configuring Secret Key Objects

The point of a secret key object is to obscure passwords or sensitive strings. If you do not want to risk someone seeing a string used in a FlexConfig object or Smart CLI template, create a secret key object for the string.

Procedure

Step 1 Select **Objects**, then select **Secret Keys** from the table of contents.

Step 2 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon ( ) for the object.

To delete an unreferenced object, click the trash can icon ( ) for the object.

Step 3 Enter a Name for the object and optionally, a description.

Step 4 Enter the password or other secret string in both the **Password** and **Confirm Password** fields.

The system obscures the text as you type.

Step 5 Click **OK**.

What to do next

- If this is a new object, to use it in FlexConfig, edit a FlexConfig object, create a variable of the secret key type, and select the object. Then, refer to the variable within the object body. For more information, see Creating Variables in a FlexConfig Object, on page 845.
- If you are editing an existing object that is used in a FlexConfig object that is part of the FlexConfig policy, you need to deploy the configuration to update the device with the new string.
- In Smart CLI templates, if a command requires a secret key, you will see a list of these objects when editing the relevant property. Select the right key for the purpose.

Troubleshooting the FlexConfig Policy

After editing the FlexConfig policy, carefully examine the results of the next deployment. If you get a “Last Deployment Failed” message in the Pending Changes dialog box, click the **See Details** link. The link takes you to the audit log, where you can find the failed deployment job. Open the job to find the specific error messages.

If the deployment fails because of a FlexConfig problem, the details will mention the FlexConfig object with the bad command, and show the command that failed. Use this information to correct the object and try deployment again. The object name is a link, click it to open the edit dialog for the object.

For example, you might want to configure the maximum TCP segment size (TCP MSS). You can control this setting with the **sysopt connection tcpmss** command. When configured by Firewall Device Manager, the Firewall Threat Defense default for this option is 0, compared to the ASA default of 1380.

The ASA default is designed for optimal processing when running an IPv4 VPN on interfaces that use the default MTU of 1500. The system needs 120 bytes for the VPN headers. For IPv6, the system needs 140 bytes. The Firewall Threat Defense default of 0 simply allows the endpoints to negotiate the MSS, which is the ideal setting for normal traffic, especially if you use different MTUs across the interfaces on the device,

including MTUs over 1500. Because TCP MSS is a global setting and not per-interface, you would change it only if a significant percentage of your traffic is over VPN and you are getting excessive fragmentation. In that case, you might set TCP MSS to MTU minus 120 (for IPv4) or 140 (for IPv6), and use the same MTU for all interfaces. Note that 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.

For purposes of illustration, suppose you want to set TCP MSS to 3 bytes. The command takes 48 bytes as the minimum value, so you will get a deployment error similar to the following:



The error is composed of these elements:

1. The deployment error message, which includes the name of the FlexConfig object that caused the error. The object name is linked to the edit dialog box so you can quickly open the object and correct the error. This is the first sentence of the message.
2. The text starting with “ERROR:” is the message returned from the device. This is exactly how an ASA would respond if you typed in the errant command, without the formatting of an SSH client. In this example, the error message is “ERROR: [3] is smaller than the minimum allowed MSS of 48 by RFC 791.” The text that starts with "Config Error" mentions the specific line that generated the error message.
3. The text in black is the actual line from the FlexConfig object that caused the error. You must fix this line. In this example, if you are trying to accommodate IPv4 VPN traffic on MTU 1500 interfaces (the common situation), you would change 3 to 1380.

   In fixing this example, you can keep the CLI Console open and use **show running-config all sysopt** to see all of the **sysopt** command settings. Most of the **sysopt** commands have default settings appropriate for most uses, so they do not appear in the running configuration. The **all** keyword includes these default settings in the output.

Examples for FlexConfig

The following topics provide some examples for using FlexConfig to configure features.

How to Enable and Disable Global Default Inspections

Some protocols embed IP addressing information in the user data packet or open secondary channels on dynamically assigned ports. These protocols require the system to do a deep packet inspection so that NAT can be applied and secondary channels can be allowed. Several common inspection engines are enabled on the system by default, but you might need to enable others, or disable default inspections, depending on your network.

To see the list of currently enabled inspections, use the **show running-config policy-map** command, either in CLI Console or an SSH session. Following is what you would see on a system where no changes have been made to the inspection configuration. In this output, the list of **inspect** commands at the end of the output shows which protocol inspections are enabled. The preceding commands enable these inspections on the

inspection_default traffic class (which is the normal protocols and, if applicable, port numbers, for the inspected protocol). This class is part of the global_policy policy map, which applies these inspections on all interfaces using a service-policy command that is not shown in the output. For example, ICMP inspection is done on all ICMP traffic that passes through the device.

```
> show running-config policy-map
!
policy-map type inspect dns preset_dns_map
 parameters
  message-length maximum client auto
  message-length maximum 512
  no tcp-inspection
policy-map global_policy
 class inspection_default
  inspect dns preset_dns_map
  inspect ftp
  inspect h323 h225
  inspect h323 ras
  inspect rsh
  inspect rtsp
  inspect sqlnet
  inspect skinny
  inspect sunrpc
  inspect xdmcp
  inspect sip
  inspect netbios
  inspect tftp
  inspect ip-options
  inspect icmp
  inspect icmp error
!
```

Note For a detailed discussion of each inspection, see the *Cisco ASA Series Firewall Configuration Guide* available from https://www.cisco.com/c/en/us/support/security/asa-firepower-services/products-installation-and-configuration-guides-list.html.

The following procedure shows you how to enable, or disable, inspections in this globally-applied default inspection class. For purposes of illustration, the example:

- Enables PPTP (Point-to-Point Tunneling Protocol). This protocol is used for tunneling a point-to-point connection between to endpoints.
- Disables SIP (Session Initiation Protocol). You would typically disable SIP only if the inspection is causing problems in the network. However, if you disable SIP, you must ensure that your access control policies allow the SIP traffic (UDP/TCP 5060) and any dynamically allocated ports, and that you do not need NAT support for SIP connections. Adjust the access control and NAT policies accordingly through the standard pages, not through FlexConfig.

Before you begin

Good planning will help you use FlexConfig efficiently. In this example, we are changing two different and unrelated inspections, although we are making the changes in the same traffic class. But it is highly likely that if you need to alter these policies, you will do so independently.

Therefore, we recommend creating separate FlexConfig objects for each inspection in this example. That way, you can easily change your setting for one inspection without changing the other, and without needing to edit the FlexConfig object.

Procedure

Step 1 Click **View Configuration** in **Device** > **Advanced Configuration**.

Step 2 Click **FlexConfig** > **FlexConfig Objects** in the Advanced Configuration table of contents.

Step 3 Create the object to enable PPTP inspection.

a) Click the + button to create a new object.

b) Enter a name for the object. For example, **Enable_PPTP_Global_Inspection**.

c) In the **Template** editor, enter the following lines, including indentations.

```
policy-map global_policy
 class inspection_default
  inspect pptp
```

d) In the **Negate Template** editor, enter the lines required to undo this configuration.

Just as you need to include the parent commands to enter the correct sub-mode for a command to enable it, you also need to include those commands in the negate template.

The negate template will be applied if you remove this object from the FlexConfig policy (after having deployed it successfully), and also during an unsuccessful deployment (to reset the configuration to its previous condition).

Thus, for this example, the negate template would be the following:

```
policy-map global_policy
 class inspection_default
  no inspect pptp
```

The object should look like the following:



Note
Because the inspection_default class has other inspection commands enabled, you do not want to negate the entire class. Similarly, the global_policy policy map includes these other inspections, and you do not want to negate the policy map either.

e) Click **OK** to save the object.

Step 4 Create the object to disable SIP inspection.

a) Click the + button to create a new object.
b) Enter a name for the object. For example, **Disable_SIP_Global_Inspection**.
c) In the **Template** editor, enter the following lines, including indentations.

```
policy-map global_policy
 class inspection_default
  no inspect sip
```

d) In the **Negate Template** editor, enter the lines required to undo this configuration.

The "negate" command for a disabling "no" command is the command that enables the feature. Thus, the "negate" template is not just the commands to disable a feature, it is the commands to reverse whatever you do in the "positive" template. The point of the negate template is to undo your changes.

Thus, for this example, the negate template would be the following:

```
policy-map global_policy
```

```
class inspection_default
 inspect sip
```

The object should look like the following:



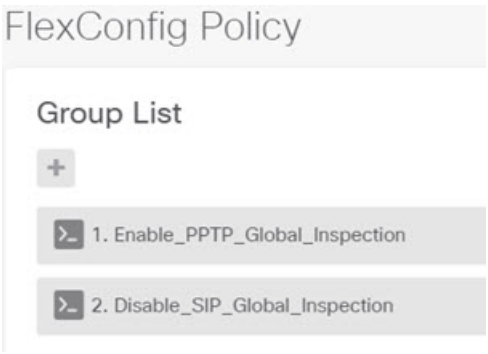
e) Click **OK** to save the object.

Step 5 Add the objects to the FlexConfig policy.

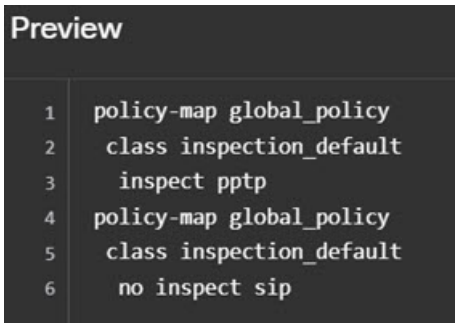
Creating an object isn't enough. Objects are deployed only if you add them to the FlexConfig policy (and save your changes). This allows you to experiment with objects (and leave them partially complete) without risking deployment failures on unfinished work. You can then easily turn features on or off simply by adding and removing objects: there is no need to recreate the object each time.

a) Click **FlexConfig Policy** in the table of contents.
b) Click + in the Group List.
c) Select the Enable_PPTP_Global_Inspection and Disable_SIP_Global_Inspection objects and click **OK**.

The group list should look like the following:



The preview should update with the commands in the template. Verify you are seeing the expected commands.



d) Click **Save**.

You can now deploy the policy.

Step 6 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.



b) Click the **Deploy Now** button.

You can wait until deployment completes, or click **OK** and check the task list or deployment history later.

Step 7 In CLI Console or an SSH session, use the **show running-config policy-map** command and verify that the running configuration has the correct changes.

In the following output, note that **inspect pptp** is added to the bottom of the inspection_default class, and that **inspect sip** is no longer in the class. This confirms that the changes defined in the FlexConfig object were successfully deployed.

```
> show running-config policy-map
!
policy-map type inspect dns preset_dns_map
 parameters
  message-length maximum client auto
  message-length maximum 512
  no tcp-inspection
policy-map global_policy
```

```
 class inspection_default
  inspect dns preset_dns_map
  inspect ftp
  inspect h323 h225
  inspect h323 ras
  inspect rsh
  inspect rtsp
  inspect sqlnet
  inspect skinny
  inspect sunrpc
  inspect xdmcp
  inspect netbios
  inspect tftp
  inspect ip-options
  inspect icmp
  inspect icmp error
  inspect pptp
!
```

How to Undo Your FlexConfig Changes

If you enter a correct negate template in a FlexConfig object, removing the changes made using that object is trivial. You simply delete the object from the FlexConfig policy, and upon the next deployment, the system uses your negate template to undo your changes.

You do not need to create a new object to undo your changes.

The following example shows how to re-enable global SIP inspection. The example reverts the change explained in How to Enable and Disable Global Default Inspections, on page 854, which disabled SIP inspection.

Before you begin

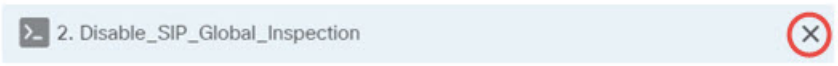
Verify that the FlexConfig object has the correct negate template. If it does not, edit the object to correct the negate template.

Procedure

Step 1 Click **View Configuration** in **Device** > **Advanced Configuration**.

Step 2 Click **FlexConfig** > **FlexConfig Policy** in the Advanced Configuration table of contents.

Step 3 Click the **X** on the right side of the **Disable_SIP_Global_Inspection** object's entry in the FlexConfig policy to delete it from the policy.



The commands from the object are removed from the preview. The negate commands are not added to the preview, these will be executed behind the scenes.

Step 4 Click **Save**.

Step 5 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.



b) Click the **Deploy Now** button.

You can wait until deployment completes, or click **OK** and check the task list or deployment history later.

Step 6 In CLI Console or an SSH session, use the **show running-config policy-map** command and verify that the running configuration has the correct changes.

In the following output, note that **inspect sip** is added to the bottom of the inspection_default class. This confirms that the changes defined in the FlexConfig object were successfully deployed. (Order is not important in this class, so it does not matter that **inspect sip** is at the end and not in its original location.)

```
> show running-config policy-map
!
policy-map type inspect dns preset_dns_map
 parameters
  message-length maximum client auto
  message-length maximum 512
  no tcp-inspection
policy-map global_policy
 class inspection_default
  inspect dns preset_dns_map
  inspect ftp
  inspect h323 h225
  inspect h323 ras
  inspect rsh
  inspect rtsp
  inspect sqlnet
  inspect skinny
  inspect sunrpc
  inspect xdmcp
  inspect netbios
  inspect tftp
  inspect ip-options
  inspect icmp
  inspect icmp error
  inspect pptp
  inspect sip
!
```

How to Enable Inspections for Unique Traffic Classes

In this example, we will enable PPTP inspection for traffic between two endpoints on a specific interface. This targets the inspection to just those endpoints that have a point-to-point tunnel configured between them.

The CLI required to enable PPTP inspection between 2 endpoints involves the following:

1. An ACL with the source and destination set to the IP addresses of the endpoint hosts.
2. A traffic class that refers to this ACL.
3. A policy map that includes the traffic class, and that enables PPTP inspection on the traffic class.

4. A service policy that applies the policy map to the desired interface. This is the step that actually activates the policy and enables the inspection.

Note For a detailed discussion of service policies related to inspections, see the *Cisco ASA Series Firewall Configuration Guide* available from https://www.cisco.com/c/en/us/support/security/asa-firepower-services/products-installation-and-configuration-guides-list.html.

Procedure

Step 1 Click **View Configuration** in **Device** > **Advanced Configuration**.

Step 2 Click **FlexConfig** > **FlexConfig Objects** in the Advanced Configuration table of contents.

Step 3 Click the + button to create a new object.

Step 4 Enter a name for the object. For example, **Enable_PPTP_Inspection_on_Interface**.

Step 5 Add a variable for the inside interface.

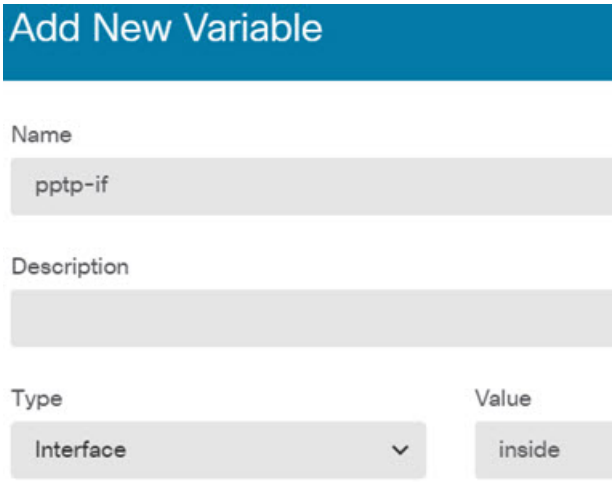
a) Click + above the Variables list.

b) Enter a name for the variable, for example, **pptp-if**.

c) For **Type**, select **Interface**.

d) For **Value**, select the **inside** interface.

The dialog box should look like the following:



e) Click **Add**.

Step 6 In the **Template** editor, enter the following lines, including indentations.

```
access-list MATCH_ACL permit ip host 192.168.1.55 host 198.51.100.1
class-map MATCH_CMAP
 match access-list MATCH_ACL
policy-map PPTP_POLICY
 class MATCH_CMAP
  inspect pptp
```

```
service-policy PPTP_POLICY interface {{pptp-if.name}}
```

Note that to use the variable, you type the variable name between double braces. You also need to use dot notation to pick out the attribute you want to retrieve, because the object that defines an interface has many attributes. Because the interface name is held in the "name" attribute, entering **{{pptp-if.name}}** retrieves the value of the name attribute for the interface assigned to the variable. If you need to change the interface for PPTP inspection, you simply need to select a different interface in the variable definition.

Step 7 In the **Negate Template** editor, enter the lines required to undo this configuration.

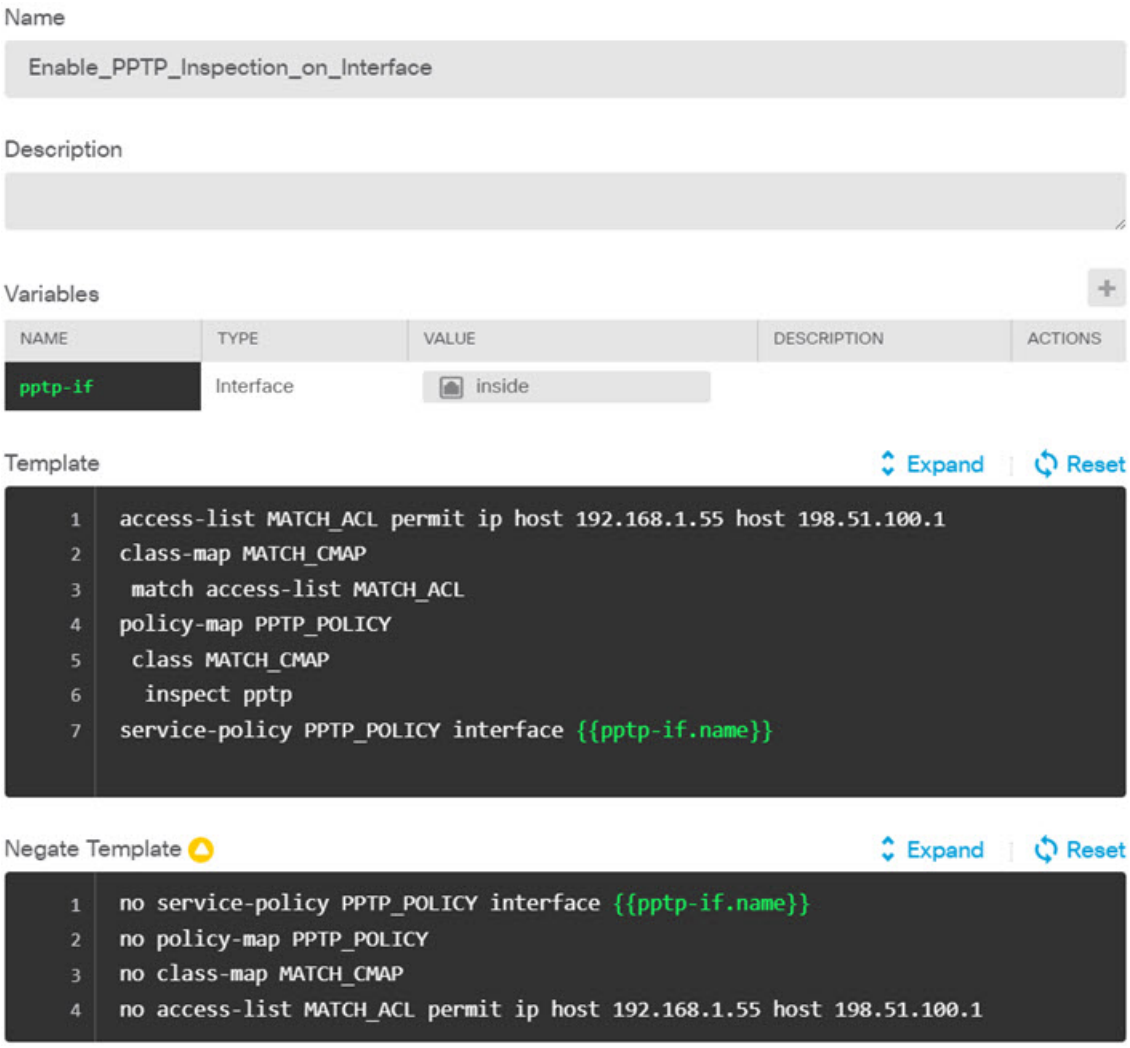
For this example, we will assume that the class map, policy map, and service policy exist for the sole purpose of applying PPTP inspection. Thus, in the negate template, we want to remove all of these.

If, however, you are actually adding PPTP inspection to an existing service policy on an interface, you would not negate the policy map or service policy. You would either negate the class from the policy map, or simply turn off inspection within the class within the policy map. You need to have a clear understanding of what you are implementing in other FlexConfig objects to ensure that your negate template does not have unintended consequences.

When deleting nested items, you need to do it in the reverse order in which you created them. Thus, you start by deleting the service policy, and end by deleting the access list. Otherwise, you would be trying to delete objects that are in use, and the system will return errors and not let you do that.

```
no service-policy PPTP_POLICY interface {{pptp-if.name}}
no policy-map PPTP_POLICY
no class-map MATCH_CMAP
no access-list MATCH_ACL permit ip host 192.168.1.55 host 198.51.100.1
```

The object should look like the following:



Step 8 Click **OK** to save the object.

Step 9 Add the objects to the FlexConfig policy.

a) Click **FlexConfig Policy** in the table of contents.

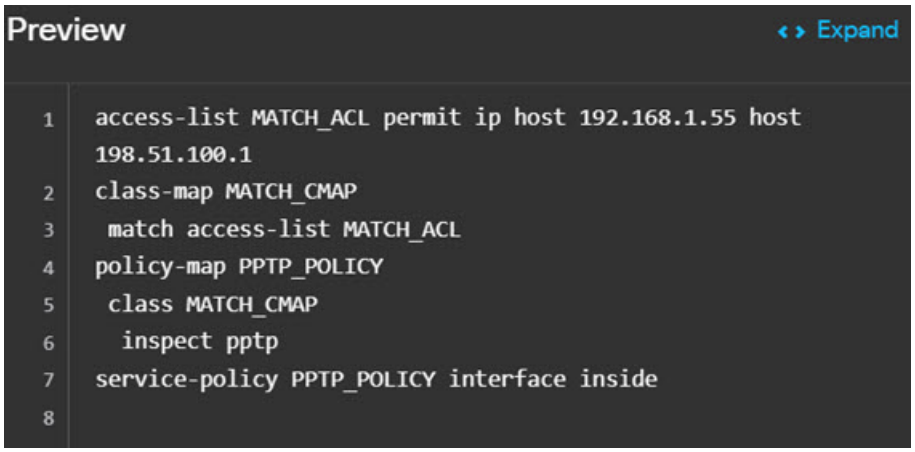
b) Click + in the Group List.

c) Select the **Enable_PPTP_Inspection_on_Interface** object and click **OK**.

The group list should look like the following:



The preview should update with the commands in the template. Verify you are seeing the expected commands, as shown in the following graphic. Notice that the interface variable resolves to the name "inside" in the preview. Pay special attention to variables: if they do not resolve correctly in the preview, they will not deploy correctly. Edit the FlexConfig object until you get the correct variable translation in the preview.



d) Click **Save**.

You can now deploy the policy.

Step 10 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.



b) Click the **Deploy Now** button.

You can wait until deployment completes, or click **OK** and check the task list or deployment history later.

Step 11 In CLI Console or an SSH session, use variations of the **show running-config** command and verify that the running configuration has the correct changes.

You can enter **show running-config** and examine the entire CLI configuration, or you can use the following commands to verify each part of this configuration:

- **show running-config access-list MATCH_ACL** to verify the ACL.
- **show running-config class** to verify the class map. This command will show all of the class maps.
- **show running-config policy-map PPTP_POLICY** to verify the class and policy map configuration.
- **show running-config service-policy** to verify that the policy map was applied to the interface. This will show all service policies.

The following output shows this sequence of commands, and you can see that configuration is correctly applied.

```
> show running-config access-list MATCH_ACL
access-list MATCH_ACL extended permit ip host 192.168.1.55 host 198.51.100.1
```

```
> show running-config class
!
class-map MATCH_CMAP
 match access-list MATCH_ACL
class-map inspection_default
 match default-inspection-traffic
!

> show running-config policy-map PPTP_POLICY
!
policy-map PPTP_POLICY
 class MATCH_CMAP
  inspect pptp
!
> show running-config service-policy
service-policy global_policy global
service-policy PPTP_POLICY interface inside
```