

Cisco Secure Firewall Management Center Administration Guide, 7.2

First Published: 2022-06-06

Last Modified: 2026-09-10

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)

© 2022–2025 Cisco Systems, Inc. All rights reserved.



CONTENTS

PART I **Getting Started** 39

CHAPTER 1 **Firewall Management Center overview** 1

- Quick start basic setup 2
  - Installing and performing initial setup on physical appliances 2
  - Deploy virtual appliances 2
  - First-time login 3
  - Configure basic policies and settings 5
- Unsupported screens for the latest device version 7
- Firewall Threat Defense devices 7
- Features 7
  - Understand appliance and system management features 7
  - Use features to detect, prevent, and process potential threats 9
  - Integrate with external tools 11
- Global search features in Firewall Management Center 12
  - Search for web interface menu options 15
  - Search for policies 15
  - Search for Objects 17
  - Search for how to walkthroughs 21
- Switch domains on Secure Firewall Management Center 22
- Context menus 23
- Data sharing features with Cisco 24
- Online help, How To, and documentation 25
  - User guides on Cisco.com 26
  - License statements in the documentation 26
  - Supported devices statements in the documentation 26

Access statements in the documentation 26
IP address conventions 27
Additional Resources 27

CHAPTER 2 Logging into the Management Center 29
User accounts 29
System user interfaces 31
Web interface considerations 32
Session timeout 32
Log into the Secure Firewall Management Center web interface 33
Log into the Firewall Management Center web interface using SSO 34
Log into the Secure Firewall Management Center with CAC credentials 36
Log into the Firewall Management Center Command Line Interface 37
View your last login 37
Logging Out of the Firewall Management Center Web Interface 38
History for logging into the management center 38

PART II System Settings 41

CHAPTER 3 System Configuration 43
Integrate Firewall Management Center with Cisco Security Cloud 44
Enable SecureX integration 44
Configure Firewall Management Center to share usage metrics and statistics with Cisco 48
Configure Firewall Management Center to share device health data with Cisco 49
Meet requirements and prerequisites for the system configuration 50
Manage the Secure Firewall Management Center system configuration 50
Access lists 51
Configure an access list 51
Configure access control preferences 52
Audit logs 53
Stream audit logs to syslog 54
Stream audit logs to an HTTP server 55
Filter syslogs from audit logs 56
Audit log certificates 57

Secure audit log streaming 58
Obtain a signed audit log client certificate for the Firewall Management Center 59
Import an audit log client certificate into the Firewall Management Center 60
Require valid audit log server certificates 61
View the audit log client certificate on the Firewall Management Center 63
Change reconciliation 64
Configure change reconciliation 64
Change reconciliation options 65
DNS caches 66
Configure DNS cache properties 66
Dashboards 67
Enable custom analysis widgets for dashboards 67
Database 67
Configure database event limits 68
Database event limits 68
Email notifications 71
Configure a mail relay host and notification address 71
External database access 72
Enable external access to the database 73
HTTPS certificates 74
Default HTTPS server certificates 74
Custom HTTPS server certificates 75
Meet HTTPS server certificate requirements 75
HTTPS client certificates 77
View the current HTTPS server certificate 78
Generate an HTTPS server certificate signing request 78
Import HTTPS server certificates 80
Require valid HTTPS client certificates 81
Renew the default HTTPS server certificate 82
View system information fields 83
Intrusion policy preferences 84
Set intrusion policy preferences 84
Languages 85
Set the language for the web interface 85

Login banners 86
Customize the login banner 86
Management interfaces 87
Firewall Management Center management interfaces 87
Device management 87
Management connections 88
Management interfaces on the Firewall Management Center 89
View management interface support per Firewall Management Center model 90
Network routes on Firewall Management Center management interfaces 90
NAT environments 92
Recommendation: management and event traffic channel configuration 94
Modify Firewall Management Center management interfaces 96
Change both Firewall Management Center and Threat Defense IP addresses 99
Network analysis policy preferences 103
Processes 104
Shut down or restart the Firewall Management Center 105
REST API preferences 105
Enable REST API access 106
Remote console access management 107
Configure remote console settings 107
Enable Lights-Out Management user access 108
Use IPMI with serial over LAN 109
Configure serial over LAN with IPMItool 111
Configure serial over LAN with IPMIutil 111
Recommendation: using IPMI with Lights-Out management 112
Configure lights-out management with IPMItool 113
Configure lights-out management with IPMIutil 113
Remote storage devices 114
Configure local storage 114
Configure NFS for remote storage 115
Configure SMB for remote storage 116
Configure SSH for remote storage 118
SNMP 119
Configure SNMP polling 119

Session timeouts 120
Configure session timeouts 121
Configure time settings 121
View NTP server status 122
Time synchronization 123
Synchronize time on the Firewall Management Center with an NTP server 125
Synchronize time without access to a network NTP server 126
Restrictions for changing time synchronization settings 127
UCAPL/CC compliance 128
Configure user settings 128
Set a password reuse limit 129
Track successful logins 130
Enable temporary lockouts 130
Set the maximum number of concurrent sessions 131
VMware tools 132
Enable VMware Tools on the Secure Firewall Management Center for VMware 132
Vulnerability mapping 133
Mapping vulnerabilities for servers 133
Configure web analytics 134
History for system configuration 135

CHAPTER 4 Users 139
CLI users 139
Internal and external users 139
Web interface and CLI access 140
User roles 140
User passwords 142
Users and domains 144
Guidelines and limitations for user accounts for Firewall Management Center 146
Requirements and prerequisites for user accounts for Firewall Management Center 147
Add or edit an internal user 148
External authentication configuration for Firewall Management Center 150
External authentication objects for Firewall Management Center 150
LDAP 151

RADIUS 152
Add an LDAP external authentication object for the Firewall Management Center 152
Add a RADIUS external authentication object for Firewall Management Center 160
Enable external authentication for users on the Firewall Management Center 166
Configure common access card authentication with LDAP 167
SAML Single Sign-On 169
SAML SSO features 169
Guidelines for SSO configuration on Firewall Management Center 170
SSO user accounts 171
User role mappings for SSO users 172
Enable SSO at the Firewall Management Center 173
Configure SSO with Okta 174
Review the Okta Org 175
Configure the Firewall Management Center service provider application for Okta 175
Configure the Firewall Management Center for Okta SSO 177
Configure User Role Mapping for Okta in the Firewall Management Center 178
Okta IdP user role mappings 179
Okta user role mapping examples 182
Configure SSO with OneLogin 187
Review the OneLogin subdomain 188
Configure a Firewall Management Center service provider application for OneLogin 188
Configure the Firewall Management Center for OneLogin SSO 190
Configure User Role Mapping for OneLogin in the Firewall Management Center 191
Configure User Role Mapping at the OneLogin IdP 192
OneLogin User Role Mapping Examples 195
Configure SSO with Azure AD 200
Azure tenants 201
Configure the Firewall Management Center service provider application for Azure 202
Configure the Firewall Management Center for Azure SSO 203
Configure user role mapping for Azure in the Firewall Management Center 205
Azure IdP user role mappings 206
Azure User Role Mapping Examples 209
Configure Single Sign-On with PingID 214
Review the PingID PingOne for Customers Environment 215

Configure a service provider application for PingID PingOne for Customers 215
Configure the Firewall Management Center for SSO with the PingID PingOne for Customers 217
Configure Single Sign-On with Any SAML 2.0-Compliant SSO Provider 219
Familiarize Yourself with the SSO Identity Provider and the SSO Federation 220
Configure Firewall Management Center service provider application for any SAML 2.0-compliant SSO provider 220
Configure the Firewall Management Center for SSO using any SAML 2.0-compliant SSO provider 222
Configure User Role Mapping in the Firewall Management Center for SAML 2.0-Compliant SSO Providers 224
Configure Firewall Management Center User Role Mapping at the IdP for SAML 2.0-Compliant SSO Providers 225
Custom user roles for the web interface 226
Create Custom User Roles 226
Deactivate User Roles 228
Enable user role escalation 229
Set the escalation target role 229
Configure a custom user role for escalation 230
Escalate your user role 231
Troubleshoot LDAP authentication connections 232
Configure User Preferences 232
Change your password 232
Change an expired password 233
Change the web interface appearance 233
Specify your home page 234
Configure event view settings 234
Configure event view preferences 235
Configure file download preferences 236
Default time windows 237
Default workflows 239
Set your default time zone 239
Specify your default dashboard 240
Configure How-To settings 241
History for Firewall Management Center user accounts 241

CHAPTER 5 **Domains** 243

Multitenancy using domains 243

Domains terminology 245

Domain properties 245

Requirements and prerequisites for domains 246

Manage domains 247

Create new domains 247

Move data between domains 248

Move devices between domains 249

History for domain management 253

CHAPTER 6 **Updates** 255

Content updates 255

Product upgrades 256

Guidelines for content updates 256

Vulnerability database (VDB) updates 256

Manually update the VDB 257

Geolocation database (GeoDB) updates 258

Schedule GeoDB updates 259

Manually update the GeoDB 259

Intrusion rule updates 260

Schedule intrusion rule updates 262

Manually update intrusion rules 262

Guidelines for importing custom Snort 2 rules 263

Import custom Snort 2 rules 264

Intrusion rule update logs 265

Maintain your air-gapped deployment 267

History for content updates 267

CHAPTER 7 **Licenses** 271

Licenses 271

Smart Software Manager and accounts 271

Licensing options for air-gapped deployments 272

Licensing for Management Center and devices 272
Periodic communications with the Smart Software Manager 272
Evaluation mode 273
Out-of-compliance state 273
Unregistered state 274
End-user license agreement 274
License types and restrictions 274
Firewall Management Center Virtual licenses 276
Base licenses 277
Malware Defense Licenses 277
Threat licenses 278
URL Filtering licenses 279
AnyConnect Client licenses 279
License for export-controlled functionality 280
Firewall Threat Defense Virtual Licenses 281
Prerequisites for licensing 283
Prerequisites for licensing in high availability, clustering, and multi-instance 283
Licensing for Firewall Management Center high availability 283
License requirements for Firewall Threat Defense device high availability 284
License for device clusters 285
Licensing for multi-instance deployments 285
Create a Cisco account 286
Create a smart account and add licenses 287
Configure Smart Licensing 288
Register the Firewall Management Center for Smart Licensing 288
Register the Firewall Management Center with the Smart Software Manager 288
Register the Firewall Management Center with the Smart Software Manager On-Prem 291
Enable the export control feature for accounts without global permission 293
Assign licenses to devices 294
Assign licenses to a single device 294
Assign licenses to multiple managed devices 295
Manage Smart Licensing 296
Deregister the Firewall Management Center 296
Synchronize or reauthorize the Firewall Management Center 296

Monitor Smart License status 297
Monitor Smart Licensing 297
Troubleshoot Smart Licensing 298
Specific License Reservation 299
Prerequisites for Specific License Reservation 299
Verify that your smart account is ready to deploy Specific License Reservation 300
Enable the Specific Licensing menu option 301
Enter the Specific License Reservation authorization code into the Firewall Management Center 301
Assign specific licenses to managed devices 303
Manage specific license reservation 304
Maintain your Specific License Reservation deployment 304
Update a Specific License Reservation 304
Deactivate and return the Specific License Reservation 306
Monitor Specific License Reservation status 308
Troubleshoot Specific License Reservation 309
Configure Legacy Firewall Management Center PAK-Based Licenses 310
Additional information about licensing 311
History for licenses 311

CHAPTER 8 High availability 313
High availability in Firewall Management Center 313
Roles and statuses in Firewall Management Center high availability 315
Event processing on Firewall Management Center high availability pairs 315
AMP cloud connections and malware information 315
URL filtering and security intelligence 316
Processing user data during Firewall Management Center failovers 316
Configuration management on Firewall Management Center high availability pairs 316
Firewall Management Center high availability disaster recovery 317
Configure single sign-on for high availability pairs 317
Firewall Management Center high availability behaviors during a backup 318
Firewall Management Center high availability split-brain 318
Troubleshoot Firewall Management Center high availability 319
Requirements for Firewall Management Center high availability 321
Requirements for hardware high availability 322

Understand virtual platform requirements 322
Requirements for software version consistency in high availability 323
License requirements for Firewall Management Center high availability configurations 324
Prerequisites for Firewall Management Center high availability 325
Establish high availability for Firewall Management Center 325
Configure high availability for Firewall Management Centers hosted on public cloud 327
View Firewall Management Center high availability status 331
Configurations Synced on Management Center High Availability Pairs 332
Recommendation: Configure external access to the Firewall Management Center database only on the active peer 332
Resolve device registration using CLI in Firewall Management Center high availability 333
Switch peer roles in the Firewall Management Center high availability pair 334
Pause communication between paired Firewall Management Centers 334
Restart communication between paired Firewall Management Centers 335
Change the IP address of the Firewall Management Center in a high availability pair 335
Disable Firewall Management Center high availability 336
Replace Firewall Management Centers in a high availability pair 337
Replace a failed primary Firewall Management Center (successful backup) 337
Replace a failed primary Firewall Management Center (unsuccessful backup) 338
Replace a failed secondary Firewall Management Center (successful backup) 339
Replace a failed secondary Firewall Management Center (unsuccessful backup) 340
Firewall Management Center high availability disaster recovery 341
Restore management center in a high-availability pair (no hardware failure) 341
Restore a backup on the primary Firewall Management Center 342
Restore a backup on the secondary Firewall Management Center 342
Unified backups of management centers in high availability 343
Restore Management Center from unified backup 344
History for Firewall Management Center high availability 344

CHAPTER 9 **Security certifications compliance** 347
Security certifications compliance modes 347
Security certifications compliance characteristics 348
Security certifications compliance recommendations 349
Appliance hardening 350

Protect your network 351
Enable security certifications compliance 352

PART III **Health and Monitoring** 355

CHAPTER 10 **Dashboards** 357

Dashboards 357
Best practices for dashboard usage 358
Dashboard widgets 358
Widget availability by appliance, user role, and domain 359
Dashboard widget availability by user role 359
Predefined dashboard widgets 360
The Allow list events widget 361
The appliance information widget 361
Appliance status widget 362
The correlation events widget 362
Current interface status widget 362
Current sessions widget 363
The custom analysis widget 363
Custom analysis widget preferences 365
View associated events from the custom analysis widget 366
Disk usage widget categories and features 367
The interface traffic widget 368
The intrusion events widget 368
Network compliance widget 369
The product licensing widget 369
Product updates widget 370
RSS feed widget 370
The system load widget 371
System time widget 371
Manage dashboards 371
Add a dashboard 372
Adding widgets to a dashboard 372
Customize the widget display 373

Configure widget preferences 374
Create custom dashboards 374
Custom dashboard options 374
Edit dashboard options 375
Modify dashboard time settings 376
Time constraints in dashboard tables 377
Rename a dashboard 377
View dashboards 378

CHAPTER 11 Health 379
Requirements and prerequisites for health monitoring 379
Health monitoring 379
Health modules 381
Configure health monitoring 392
Health policies 392
Default health policy 393
Create a health policy 393
Apply a health policy 394
Edit a health policy 394
Delete a health policy 396
Device exclusion in health monitoring 396
Exclude devices from health monitoring 397
Exclude health policy modules on a device 398
Exclude interface health monitoring 399
Manage expired health monitor exclusions 399
Health monitor alerts 400
Health monitor alert information 400
Create a health monitor alert 401
Edit a health monitor alert 402
Delete health monitor alerts 402
About the health monitor 402
View the Firewall Management Center health monitor dashboard 404
Run all health modules for an appliance 405
Run a health module test 406

Generate health module alert graphs 406
Device health monitors 407
View system details and troubleshoot devices 408
View the device health monitor 408
Health monitor status categories 411
Health event views 412
View health events 412
View health events by module and appliance 413
View the health events table 413
Health event fields 414
History for health monitoring 415

CHAPTER 12 Audit and Syslog 423
System log 423
View the system log 423
Syntax for system log filters 424
About system auditing 425
Audit records 425
View audit records 425
Suppress audit records 429
About sending audit logs to an external location 432

CHAPTER 13 Statistics 433
System statistics 433
Host statistics section 433
Disk usage section 434
The processes section 434
Process status fields 434
System daemons 436
Executables and system utilities 437
SFDataCorrelator process statistics section 439
Intrusion event information section 440
View system statistics 441

CHAPTER 14 **Troubleshooting** 443

Guidelines and best practices for troubleshooting 443

System messages 444

Message types 444

Message management 446

View basic system information 446

View appliance information 447

Manage system messages 447

View deployment messages 448

View upgrade messages 449

View health messages 449

View task messages 450

Manage task messages 450

Memory usage thresholds for health monitor alerts 451

Disk usage and drain of events health monitor alerts 452

Disk usage for device configuration history files 456

Health monitor reports for troubleshooting 456

Generate troubleshooting files for specific system functions 457

Download advanced troubleshooting files 458

General troubleshooting 458

Connection-based troubleshooting 458

Troubleshoot a connection 459

Advanced troubleshooting for the Secure Firewall Threat Defense device 459

Packet capture 460

Use a packet capture trace 462

Packet tracer tool overview 463

Run a packet trace 464

Use the Firewall Threat Defense diagnostic CLI from the web interface 466

Feature-specific troubleshooting 467

PART IV **Tools** 469

CHAPTER 15 **Backup/Restore** 471

Backup and restore 471
Requirements: backup and restore configuration 473
Guidelines and limitations for backup and restore 474
Recommendation: Configuration Import/Export Guidelines for Firepower 4100/9300 475
Best practices for backup and restore 476
Backing up Firewall Management Centers or managed devices 478
Back up the Firewall Management Center 478
Back up a Device from the Firewall Management Center 480
Export an FXOS configuration file 481
Create a backup profile 482
Restoring Firewall Management Centers and managed devices 484
Restore Firewall Management Center from Backup 484
Restore Firewall Threat Defense from Backup 485
Zero-Touch Restore Firewall Threat Defense from Backup: ISA 3000 489
Restore Firewall Threat Defense from Backup: Firepower 4100/9300 Chassis 492
Importing a Configuration File 495
Restore Firewall Threat Defense Virtual from Backup 496
Manage backups and remote storage 499
Backup storage locations 501
History for backup and restore 502

CHAPTER 16 **Scheduling** 505
Task scheduling 505
Guidelines for scheduled tasks 505
Scheduled task types 506
Schedule a task 511
History for scheduling 511

CHAPTER 17 **Import/Export** 513
Import/Export feature overview 513
Exceptions to export behavior 513
Import objects and object groups 514
Conflict resolution for duplicate configurations 514
Requirements and prerequisites for import/export 516

Guidelines for Import/Export 516
Export configurations 517
Import configurations 518

CHAPTER 18 Data Purge and Storage 521
Data stored on the Firewall Management Center 521
Purge data from the Firewall Management Center Database 522
External data storage 523
Comparison of Security Analytics and Logging remote event storage options 524
Remote data storage in Cisco Secure Cloud Analytics 524
Remote data storage on a Secure Network Analytics appliance 525
History for data storage 525

PART V Reporting and Alerting 527

CHAPTER 19 Reports 529
Requirements and prerequisites for reports 529
Reports 529
Risk reports 530
Risk report templates 530
Generate, view, and print risk reports 530
Standard reports 531
About designing reports 532
Report templates 532
Report template section fields 532
Report templates 534
Create a custom report template 535
Create a report template from an existing template 536
Create a report template from an event view 536
Create a report template by importing a dashboard or workflow 537
Data source options in the import report sections window 538
Report templates 538
Configure table and data format in a report template section 539
Specify search or filter criteria for a report template 540

Modify fields in a report table 540
Add a text section to a report template 541
Add a page break to a report template 541
Global time windows in report template sections 542
Set the global time window for a report template and its sections 542
Set the local time window for report template sections 542
Rename a report template section 543
Preview a report template section 543
Report template searches 543
Search in report template sections 544
Input parameters 544
Report template predefined input parameters 545
User-defined input parameters 546
Create user-defined input parameters 546
Edit user-defined input parameters 547
Constrain a search with user-defined input parameters 547
Document attributes in report templates 548
Edit document attributes in a report template 548
Customize a cover page 549
Manage report template logos 549
Add a new logo 550
Change the logo for a report template 550
Delete a logo 550
Manage report templates 551
Editing report templates 551
Export report templates 552
About generating reports 553
Generate reports 553
Report generation options 554
Distribute reports by email when generated 555
Scheduled future reports 555
Generated reports 556
View reports 556
Download reports 557

Store reports remotely 557
Moving reports to remote storage 558
Delete reports 558
Report template feature history 559

CHAPTER 20 External Alerting with Alert Responses 561
Configuring external alerts with alert responses 561
External alert types 562
Prerequisites for external alerting with alert responses 562
Guidelines for notifying external systems with alert responses 562
Alert responses 563
Create an SNMP alert response 563
Syslog alert responses 564
SYSLOG alert facilities 565
Syslog severity levels 566
Create a syslog alert response 566
Create an email alert response 567
External alerts 568
Configure impact flag alerts 568
Configure discovery event alerts 569
Configure Malware defense alerts 569
Troubleshoot external alerts and alert response configuration 570
History for external alerting with alert responses 571

CHAPTER 21 External Alerting for Intrusion Events 573
External alerts for intrusion events 573
License requirements for external alerting for intrusion events 574
Requirements for configuring external alerting for intrusion events 574
Configure SNMP alerting for intrusion events 574
Intrusion SNMP alert configuration options 575
Configure syslog alerting for intrusion events 576
Facility and severity values for intrusion syslog alerts 577
Configure email alerting for intrusion events 579
Intrusion email alert options 579

PART VI Event and Asset Analysis Tools 581

CHAPTER 22 Context Explorer 583

Context explorer 583

Differences between the dashboard and the context explorer 584

Traffic and intrusion event counts time graph 585

Indications of compromise section 585

Hosts by indication graph 585

Indications by host graph 586

Network information section 586

Operating systems graph 586

Traffic by source IP graph 587

Traffic by source user graph 587

Connections by access control action graph 587

Traffic by destination IP graph 588

Traffic by ingress or egress security zone graphs 588

Application information section 589

Change the focus of Application Protocol Information section 589

Traffic by Risk or Business relevance and application graph 589

Intrusion events by risk or business relevance and application graph 590

Hosts by risk or business relevance and application graph 590

Application details list 591

Security intelligence section 591

Security intelligence traffic by category graph 591

Security intelligence traffic by source IP graph 592

Security intelligence traffic by destination IP graph 592

Intrusion information section 592

Intrusion events by impact graph 593

The top attackers graph 593

The Top Users Graph 593

Intrusion events by priority graph 593

Top targets graph 594

Top ingress or egress security zones graph 594

Intrusion event details list 594
Files information section 595
The Top File Types Graph 595
Top file names graph 595
Files by disposition graph 596
Top hosts sending files graph 596
Top hosts receiving files graphs 597
Top malware detections graph 597
Geolocation information section 597
Connections by Initiator or Responder country graph 598
Intrusion events by Source or Destination country graph 598
File events by sending or receiving country graph 598
URL information section 599
Traffic by URL graph 599
Traffic by URL category graphs 599
Traffic by URL reputation graph 600
Requirements for the Context Explorer 600
Reload the context explorer 600
Set the context explorer time range 601
Minimize and maximize context explorer sections 601
Drill down on context explorer data 602
Filters in the context explorer 603
Data type field options 604
Create a filter from the Add Filter window 606
Create a quick filter from the context menu 606
Save filtered context explorer views 607
View filter data 607
Delete a filter 607

CHAPTER 23 **Unified Events** 609
Unified events 609
Requirements for using unified events 609
Work with unified events 610
Set a time range in unified events 613

Enable live event monitoring in unified events 613
Filters in unified events 614
Unified events column details 615
History for unified events 616

CHAPTER 24 **Network Map** 619
Requirements for using network map 619
Network maps 619
Hosts network map 620
Network devices network map 621
Mobile devices network map 622
Indications of compromise network map 622
Application protocols network map 622
Vulnerabilities network map 623
Host attributes network map 624
View network maps 624
Custom network topologies 625
Create custom topologies 626
Import networks from the network discovery policy 627
Manually add networks to your custom topology 627
Activate or deactivate custom topologies 628
Edit custom topologies 628

CHAPTER 25 **Lookups** 631
Lookups in Firewall Management Center 631
Requirements for using lookups 631
Lookup IP address details using Whois 631
Check the category and reputation of the URLs 632
Find geolocation information for an IP address 633

CHAPTER 26 **Event Analysis Using External Tools** 635
Cisco cloud event settings 635
Enable sending events to the Cisco Security Cloud 636
Analyze events using Cisco XDR 637

Event investigation using web-based resources 638
Manage contextual cross-launch resources 639
Requirements for custom contextual cross-launch resources 639
Add contextual cross-launch resources 640
Investigate events using contextual Cross-Launch 641
Configure cross-launch links for Secure Network Analytics 641
Sending syslog messages for security events 643
Security event syslog platform settings 643
Best practices for configuring security event syslog messaging 644
Send security event syslog messages from your Firewall Threat Defense devices 645
Configure security event syslog messages for Classic devices 647
Configuration locations for security event syslogs 648
Components of security event syslog messages 652
Facility values and locations for security event syslog messages 655
Syslog message types supported by Secure Firewall devices 655
Limitations of syslog for security events 656
eStreamer server streaming 656
Comparison of syslog and eStreamer for security eventing 657
Data sent only via eStreamer, not via syslog 658
Choose eStreamer event types 659
Configure eStreamer client communications 659
Event analysis in splunk 660
IBM QRadar events 661
History for analyzing event data using external tools 661

PART VII **Workflows and Tables** 665

CHAPTER 27 **Workflows** 667
Workflows 667
Predefined workflows 667
Predefined intrusion event workflows 668
Predefined malware workflows 669
Predefined file workflows 669
Predefined captured file workflows 670

Predefined connection data workflows 670
Predefined security intelligence workflows 671
Predefined host workflows 672
Predefined indications of compromise workflows 672
Predefined applications workflows 672
Predefined application details workflows 673
Predefined servers workflows 673
Predefined host attributes workflows 674
The predefined discovery events workflow 674
Predefined user workflows 674
Predefined vulnerabilities workflows 674
Predefined third-party vulnerabilities workflows 675
Predefined correlation and Allow list workflows 675
Predefined system workflows 675
Custom table workflows 676
Use workflows 676
Workflow access by user role 678
Workflow selection 678
Workflow pages 680
Workflow page navigation tools 681
Workflow page traversal tools 681
File trajectory icons 682
Host profile icons 682
Threat score icons 683
User icons 683
Workflow toolbar 683
Use drill-down pages 684
Use table view pages 684
Work in Firewall Management Center with connection events stored on a Secure Network Analytics appliance 685
Geolocation 686
Connection event graphs 686
Use connection event graphs 687
Event time constraints 692

Per-session time window customization for events 693
Default time window for events 696
Event view constraints 698
Constrain events 699
Compound event view constraints 700
Use compound event view constraints 701
Navigate between workflows 701
Work with the unified event viewer 702
Bookmarks 703
Create bookmarks 703
View bookmarks 704
History for workflows 704

CHAPTER 28 **Event Search** 707
Event searches 707
Search constraints 708
Best practice for general search constraints 708
Wildcards and symbols in searches 709
Objects and application filters in searches 709
Time constraints in searches 709
IP addresses in searches 710
URLs in searches 711
Managed devices in searches 711
Ports in searches 711
Event fields in searches 712
Perform a search 713
Save a search 713
Load a saved search 714
Query overrides via the shell 715
Shell-Based query management syntax 715
Stop long-running queries 716
History for search for events 716

CHAPTER 29 **Custom Workflows** 717

Custom workflows 717
Saved custom workflows 717
How custom workflows work 718
Create custom workflows based on non-connection data 719
Create custom connection data workflows 720
How to access custom workflows 721
View custom workflows based on predefined tables 722
View custom workflows based on custom tables 722
Edit custom workflows 723

CHAPTER 30 **Custom Tables** 725
Custom tables 725
Predefined custom tables 725
Possible table combinations 726
User-defined custom tables 729
Create a custom table 729
Modify a custom table 730
Delete a custom table 730
View a workflow based on a custom table 731
Search custom tables 731
History for custom tables 732

PART VIII **Events and Assets** 735

CHAPTER 31 **Connection Logging** 737
Connection logging 737
Logged connections 738
Other Connections You Can Log 739
Rules and policy actions effect on logging 740
Logging for fastpathed connections 740
Logging for monitored connections 740
Logging for trusted connections 741
Logging for blocked connections 741
Logging for allowed connections 743

Connection logging modes 744
Secure Firewall Management Center vs external logging 745
Connection logging limitations 746
Event viewer event appearances 747
Best practices for connection logging 747
Requirements and prerequisites for connection logging 750
Connection logging 750
Logging connections with tunnel and prefilter rules 750
Configure logging for decryptable connections with rules 751
Logging connections with security intelligence 752
Logging connections with access control rules 752
Logging connections with a policy default action 753
Limiting logging of long URLs 754

CHAPTER 32 **Connection and Security-Related Connection Events** 755
Connection events 755
Connection vs. Security Intelligence events 756
NetFlow connections 756
Connection summaries (aggregated data for graphs) 756
Long-running connections 757
Combined connection summaries from external responders 757
Connection and Security Intelligence Event Fields 758
Connection and Security Intelligence Event data overview 773
Initiator and responder, source and destination, and sender and receiver fields 773
Connection event reasons 774
Requirements for populating connection event fields 776
Information available in connection event fields 778
View Connection and Security Intelligence Event Tables 779
View files and malware detected in a connection 781
View intrusion events for a connection 782
Encrypted connection certificate details 783
View the connection summary page 783
History for Connection and Security Intelligence Events 784

CHAPTER 33 **Intrusion Events** 787

Intrusion events 787
Tools for reviewing and evaluating intrusion events 788
License requirements for intrusion events 788
Requirements and prerequisites for intrusion events 789
View intrusion events 789
Intrusion event data overview 790
Intrusion event fields 790
Intrusion event impact levels 802
View connection data associated with intrusion events 803
Mark intrusion events reviewed 803
View previously reviewed intrusion events 804
Mark reviewed intrusion events unreviewed 804
Preprocessor events 805
Preprocessor generator IDs 805
Intrusion event workflow pages 807
Use intrusion event workflows 808
Intrusion event drill-down page constraints 810
Intrusion event table view constraints 811
Using the intrusion event packet view 812
Event information fields 813
Configure intrusion rules within the packet view 816
Set threshold options within the packet view 817
Set suppression options within the packet view 818
Frame information fields 819
Data link layer information fields 819
View network layer information 820
View transport layer information 823
View packet byte information 825
Internally sourced intrusion events 826
View intrusion event statistics 826
Host statistics 827
Event overview 827

Event statistics 827
View intrusion event performance graphs 828
Intrusion event performance statistics graph types 828
View intrusion event graphs 832
History for intrusion events 832

CHAPTER 34 **File/Malware Events and Network File Trajectory 835**
File and malware events and network file trajectory 835
File and malware events 836
File and malware event types 837
File events 837
Malware events 837
File rule actions and file disposition in malware events 837
Retrospective malware events 838
Malware events generated by Secure Endpoint 839
Malware event analysis with Secure Endpoint 839
Event data from Secure Endpoint 840
Using file and malware event workflows 841
File and malware event fields 841
Malware event sub-types 852
Information available in file and malware event fields 853
View details about analyzed files 855
File composition reports 855
View file details in AMP private cloud 856
Threat scores and dynamic analysis summary reports 856
View Dynamic Analysis Results in the Cisco Secure Malware Analytics Cloud 857
Using captured file workflows 858
Captured file fields 859
Stored files download 861
Submit files for analysis manually 862
Network file trajectory 863
Recently detected malware and analyzed trajectories 863
Network file trajectory 863
Network file trajectory summary information 864

Network file trajectory maps and related events lists 866
Using a network file trajectory 867
Work with Event Data in the Secure Endpoint Console 868
History for file and malware events and network file trajectory 869

CHAPTER 35 **Host Profiles** 871
Requirements and prerequisites for host profiles 871
Host profiles 872
Host profile limitations 873
View host profiles 873
Basic host information in the host profile 874
Operating systems in the host profile 876
View operating system identities 877
Set the current operating system identity 878
Operating system identity conflicts 879
Make a conflicting operating system identity current 879
Resolving an operating system identity conflict 879
Servers in the host profile 880
Server details in the host profile 881
View server details 883
Editing server identities 883
Resolving server identity conflicts 884
Web applications in the host profile 884
Delete web applications from the host profile 885
Host protocols in the host profile 886
Delete a protocol from the host profile 886
Indications of compromise in the host profile 886
VLAN tags in the host profile 887
User history in host profiles 887
Host attributes in the host profile 888
Predefined Host Attributes 888
Allow list host attributes 889
User-Defined Host Attributes 889
Create text- or URL-based host attributes 890

Create integer-based host attributes 890
Create list-based host attributes 891
Set host attribute values 891
Allow list violations in the host profile 892
Create Shared Allow List Host Profiles 893
Malware detections in the host profile 893
Vulnerabilities in the host profile 894
Download patches for vulnerabilities 894
Deactivate vulnerabilities for individual hosts 895
Deactivate individual vulnerabilities 895
Scan results in the host profile 896
Scan a host from the host profile 897
History for host profiles 897

CHAPTER 36 **Discovery Events** 899
Requirements and prerequisites for discovery events 899
Discovery and identity data in discovery events 899
View discovery event statistics 900
Statistics summary section 901
The event breakdown section 902
Protocol breakdown section 903
The application protocol breakdown section 903
OS breakdown section 903
View discovery performance graphs 903
Discovery performance graph types 904
Access discovery and identity workflows 905
Discovery and host input events 906
Discovery event types 907
Host input event types 911
View discovery and host input events 912
Discovery event fields 913
Host data 914
View host data 915
Host data fields 915

Create a traffic profile for selected hosts 919
Creating a compliance allow list based on selected hosts 920
Host attribute data 920
View host attributes 920
Host attribute data fields 921
Set host attributes for selected hosts 922
Indications of compromise data 922
View and work with indications of compromise data 923
Indications of compromise data fields 925
Edit indication of compromise rule states for a single host or user 925
View source events for indication of compromise tags 925
Resolving indication of compromise tags 926
Server data 926
View server data 927
Server data fields 927
Application and application details data 930
View application data 930
Application data fields 931
View application detail data 932
Application detail data fields 933
Vulnerability data 934
Vulnerability data fields 935
Vulnerability deactivation 936
View vulnerability data 937
View vulnerability details 938
Deactivate multiple vulnerabilities 938
Third-party vulnerability data 939
View third-party vulnerability data 939
Third-party vulnerability data fields 940
Active sessions, users, and user activity data 941
User-related fields 942
Active sessions data 948
User data 949
User activity data 952

User profile and host history 955
History for working with discovery events 956

CHAPTER 37 Correlation and Compliance Events 959
View correlation events 959
Correlation event fields 960
Analyze compliance events using compliance allow list workflows 962
View allow list events 963
Allow list event fields 964
View allow list violations 965
Allow list violation fields 966
Remediation status events 967
View remediation status events 967
Remediation status table fields 968
Configure the remediation status events table 969

PART IX Correlation and Compliance 971

CHAPTER 38 Compliance Lists 973
Introduction to compliance allow lists 973
Compliance allow list target networks 974
Compliance allow list host profiles 975
Operating system-specific host profiles 976
Shared host profiles 976
Allow violation triggers 977
Requirements and prerequisites for compliance 978
Create a compliance allow list 979
Set target networks for a compliance allow list 980
Configure allow list host profiles 981
Add an application protocol to a compliance allow list 982
Add a client to a compliance allow list 983
Add a web application to a compliance allow list 983
Add a protocol to a compliance allow list 983
Manage compliance allow lists 984

Edit a compliance allow list 985
Manage shared host profiles 986

CHAPTER 39 Correlation Policies 989
Correlation policies and rules 989
Requirements and prerequisites for compliance 990
Configure correlation policies 991
Add responses to rules and allow lists 991
Manage correlation policies 992
Configure correlation rules 993
Syntax for VPN troubleshoot event trigger criteria 994
Syntax for intrusion event trigger criteria 995
Syntax for malware event trigger criteria 997
Discovery event trigger criteria and event types 999
Syntax for user activity event trigger criteria 1000
Syntax for host input event trigger criteria 1001
Syntax for Connection Event Trigger Criteria 1002
Syntax for traffic profile changes 1005
Syntax for correlation host profile qualifications 1007
Syntax for user qualifications 1009
Connection trackers 1010
Add a connection tracker 1011
Syntax for connection trackers 1011
Syntax for connection tracker events 1014
How excessive connection detection works 1014
Sample configuration for excessive BitTorrent data transfers 1016
Snooze and inactive periods 1018
Correlation rule building mechanics 1019
Add AND link conditions in correlation rules 1021
Use multiple values in correlation rule conditions 1022
Manage correlation rules 1023
Configure correlation response groups 1024
Manage correlation response groups 1024

CHAPTER 40 **Traffic profiling** 1027

Introduction to traffic profiles 1027

Traffic profile conditions 1029

Requirements and prerequisites for traffic profiles 1031

Manage traffic profiles 1031

Configure traffic profiles 1032

Add traffic profile conditions 1033

Add host profile qualifications to a traffic profile 1034

Syntax for traffic profile conditions 1034

Syntax for host profile qualifications in a traffic profile 1035

Use multiple values in a traffic profile condition 1037

CHAPTER 41 **Remediations** 1039

Requirements and prerequisites for remediations 1039

Introduction to remediations 1039

Cisco ISE EPS remediations 1040

Configure ISE EPS remediations 1041

Cisco IOS null route remediations 1043

Configure remediations for Cisco IOS routers 1043

Nmap scan remediations 1048

Set attribute value remediations 1048

Configure set attribute remediations 1048

Manage remediation modules 1050

Manage remediation instances 1050

Manage instances for a single remediation module 1051

PART X **Reference** 1053

CHAPTER 42 **Secure Firewall Management Center command line reference** 1055

About the Secure Firewall Management Center CLI 1055

Firewall Management Center CLI modes 1056

Secure Firewall Management Center CLI management commands 1056

Secure Firewall Management Center CLI show commands 1057

Secure Firewall Management Center CLI configuration commands 1058
Secure Firewall Management Center CLI system commands 1058
History for the Secure Firewall Management Center CLI 1060

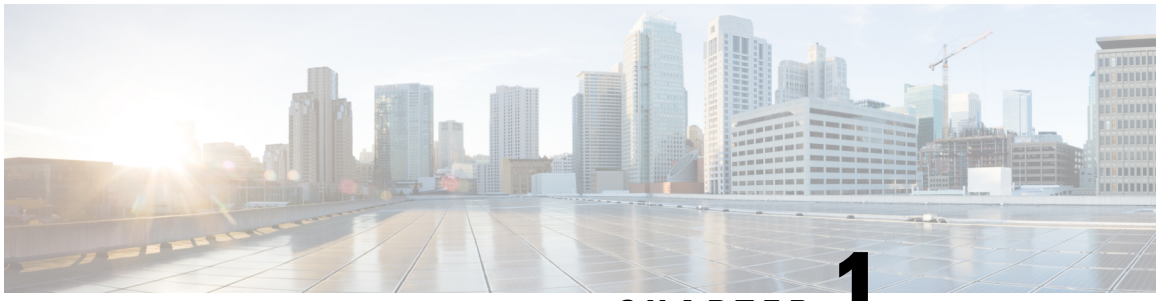
CHAPTER 43 **Communication Ports and Internet Access** 1063
Guidelines for secure management communications 1063
Communication ports for Firewall Management Center 1063
Communication ports for managed devices 1066
Internet access requirements for Firewall Management Center 1067
Internet access requirements for managed devices 1071
Internet resources accessed by your browser 1072



PART I

Getting Started

- Firewall Management Center overview, on page 1
- Logging into the Management Center, on page 29



CHAPTER 1

Firewall Management Center overview

This guide applies to an *on-premises* Firewall Management Center, either as your primary manager or as an analytics-only manager. If you use the Security Cloud Control Cloud-Delivered Firewall Management Center as your primary manager, you can use an on-prem Firewall Management Center for analytics. Use the Cloud-Delivered Firewall Management Center documentation when managing devices with Cloud-Delivered Firewall Management Center. For guidance, refer to *Cisco Security Cloud Control: Cloud-Delivered Firewall Management Center for Firewall Threat Defense*.

The Firewall Management Center is a powerful web-based multi-device manager that runs on its own server hardware or as a virtual device on a hypervisor. Use the Firewall Management Center if you want a multi-device manager and require all features on the Firewall Threat Defense. The Firewall Management Center also provides powerful analysis and monitoring of traffic and events.

Note If you have a Security Cloud Control-managed device and are using the on-prem Firewall Management Center for analytics only, then the on-prem Firewall Management Center does not support policy configuration or upgrading. This guide applies to devices managed by an on-premises Firewall Management Center. For devices managed by Security Cloud Control, refer to the relevant documentation.

For the Firewall Management Center used as the primary manager: The Firewall Management Center is not compatible with other managers because the Firewall Management Center owns the Firewall Threat Defense configuration. Configure the Firewall Threat Defense using the Firewall Management Center. Direct configuration is not supported.

- Quick start basic setup, on page 2
- Unsupported screens for the latest device version, on page 7
- Firewall Threat Defense devices, on page 7
- Features, on page 7
- Global search features in Firewall Management Center, on page 12
- Switch domains on Secure Firewall Management Center, on page 22
- Context menus, on page 23
- Data sharing features with Cisco, on page 24
- Online help, How To, and documentation, on page 25
- IP address conventions, on page 27
- Additional Resources, on page 27

Quick start basic setup

Quick start basic setup is a configuration approach that enables rapid deployment of Secure Firewall feature sets and managed devices to begin controlling and analyzing network traffic with minimal initial configuration effort.

Setup capabilities

The Secure Firewall feature set provides powerful and flexible capabilities that support both basic and advanced configurations for Secure Firewall Management Center and its managed devices.

Installing and performing initial setup on physical appliances

Procedure

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

- **Firewall Management Center**
  - Cisco Secure Management Center Getting Started Guide for your hardware model, available from Cisco Secure Firewall Management Center Getting Started Guides
- **Firewall Threat Defense managed devices**
  - Firepower 1010 getting started guides
  - Firepower 1100 getting started guides
  - *Cisco Firepower 2100 Getting Started Guide*
  - Cisco Secure Firewall 3100 Getting Started Guide
  - *Cisco Firepower 4100 Getting Started Guide*
  - *Cisco Firepower 9300 Getting Started Guide*
  - Cisco Secure Firewall Threat Defense for the ISA 3000 Using Secure Firewall Management Center Quick Start Guide

Deploy virtual appliances

This task enables you to deploy virtual appliances for Cisco Secure Firewall by identifying supported platforms and using the correct deployment guides for Management Centers and devices.

Deploy virtual appliances when your network architecture requires virtualized Cisco Secure Firewall Management Centers or devices, such as in cloud or virtualized environments.

Use the documentation roadmap to locate the relevant deployment guides: Navigating the Cisco Secure Firewall Threat Defense Documentation. Follow these steps if your deployment includes virtual appliances.

Before you begin

Review your deployment requirements. Ensure you have access to the compatibility guide and the appropriate deployment documentation for your environment. Determine whether you are deploying Management Centers, devices, or both as virtual appliances. Follow these steps to deploy virtual appliances:

Procedure

Step 1 Determine the supported virtual platforms you will use for the Management Center and devices (these may not be the same). Refer to the *Cisco Secure Firewall Compatibility Guide*.

Step 2 Deploy virtual Secure Firewall Management Centers using the documentation for your environment:

- Firewall Management Center Virtual running on VMware: *Cisco Secure Firewall Management Center Virtual Getting Started Guide*
- Firewall Management Center Virtual running on AWS: *Cisco Secure Firewall Management Center Virtual Getting Started Guide*
- Firewall Management Center Virtual running on KVM: *Cisco Secure Firewall Management Center Virtual Getting Started Guide*

Step 3 Deploy virtual devices using the documentation for your appliance:

- Firewall Threat Defense Virtual running on VMware: *Cisco Secure Firewall Threat Defense Virtual for VMware Getting Started Guide*
- Firewall Threat Defense Virtual running on AWS: *Cisco Secure Firewall Threat Defense Virtual for AWS Getting Started Guide*
- Firewall Threat Defense Virtual running on KVM: *Cisco Secure Firewall Threat Defense Virtual for KVM Getting Started Guide*
- Firewall Threat Defense Virtual running on Azure: *Cisco Secure Firewall Threat Defense Virtual for Azure Getting Started Guide*

After you complete these steps, your virtual appliances are deployed. Configure and integrate them into your network environment.

What to do next

- Refer to the configuration guides for further setup and integration of your deployed virtual appliances.
- Verify connectivity and functionality of the Management Centers and devices.

First-time login

A first-time login is a system access event that

- requires preparation of the appliance as described in Installing and performing initial setup on physical appliances, on page 2 or Deploy virtual appliances, on page 2.
- uses the **admin** account for CLI access or web interface access and refers to the instructions in the *Cisco Secure Firewall Management Center Getting Started Guide* for the specific model, and
- completes initial configuration, setting system aspects such as passwords, network settings, and enabling telemetry features.

Initial configuration settings

These aspects of your system are configured during the first login:

- During initial configuration, the system sets the passwords for the two **admin** accounts (web interface and CLI) to the same value. This process enforces strong password requirements as described in Guidelines and limitations for user accounts for Firewall Management Center, on page 146. The system synchronizes the passwords only during initial configuration. If you later change the password for either account, the passwords will no longer match, and the web interface **admin** account will no longer require a strong password. (Refer to Add or edit an internal user, on page 148.)
- During the first login, you must configure network settings for the Firewall Management Center. These settings determine how the system communicates through its management interface (eth0). By default, values are supplied, but you can set custom values as needed:
  - Fully qualified domain name (`<hostname>.<domain>`)
  - Boot protocol for IPv4 configuration (DHCP or Static Manual)
  - IPv4 address
  - Subnet mask
  - Gateway
  - DNS servers
  - NTP servers

  You can view and change these settings through the Firewall Management Center web interface. For more information, refer to Modify Firewall Management Center management interfaces, on page 96 and Time synchronization, on page 123.
- As part of the initial configuration, the system schedules multiple update-related tasks. We recommend you review these tasks to make sure they run at the right time for your environment:
  - Daily intrusion rule updates. Refer to Schedule intrusion rule updates, on page 262.
  - Weekly geolocation database (GeoDB) updates. Refer to Schedule GeoDB updates, on page 259.
  - Weekly maintenance release, patch, and vulnerability database (VDB) downloads. Refer to Scheduling, on page 505.
  - Weekly configuration-only Firewall Management Center backups (locally stored). Refer to Scheduling, on page 505.
- As part of the initial configuration, the system downloads and installs the latest VDB. To keep the system up to date, schedule recurring updates. Refer to Scheduling, on page 505.

After you complete the Firewall Management Center initial configuration, the web interface displays the device management page described in *Cisco Secure Firewall Management Center Device Configuration Guide*.

(This is the default login page only for the first time the **admin** user logs in. On subsequent logins by the **admin** or any user, the default login page is determined as described in Specify your home page, on page 234.)

After you complete the initial configuration, you can begin controlling and analyzing traffic by configuring basic policies. For details, refer to Configure basic policies and settings, on page 5.

Configure basic policies and settings

This task guides you to configure and deploy basic policies and system settings, so you can begin to monitor and manage your deployment effectively.

- Establishes initial access and configuration for your system.
- Ensures that essential policies are in place for data visibility and system operation.

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

This guide does not provide a comprehensive discussion of policy or feature capabilities. For information about additional features and advanced configurations, refer to other sections in this guide.

Before you begin

Log in to the web interface using the **admin** account for web interface or CLI. Perform the initial configuration as described in the *Cisco Secure Firewall Management Center Getting Started Guide*, which is available for your hardware model from the Install and Upgrade Guides.

Follow these steps to configure basic policies and settings:

Procedure

Step 1 Set a time zone for this account as described in Set your default time zone, on page 239.

Step 2 If needed, add licenses as described in Licenses, on page 271.

Step 3 Add managed devices to your deployment. For detailed instructions, refer to *Add a Device to the Firewall Management Center* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Step 4 Configure your managed devices as described in:

- *Interface Overview* in the *Cisco Secure Firewall Management Center Device Configuration Guide*, to configure transparent or routed mode on Firewall Threat Defense devices.
- *Interface Overview* in the *Cisco Secure Firewall Management Center Device Configuration Guide*, to configure interfaces on the Firewall Threat Defense devices.

Step 5 Configure an access control policy as described in *Creating a Basic Access Control Policy* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

- Set the **Balanced Security and Connectivity** intrusion policy as your default action in most cases. For more information, refer to *Access Control Policy Default Action* and *System-Provided Network Analysis and Intrusion Policies* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.
- In most cases, Cisco suggests enabling connection logging to meet the security and compliance needs of your organization. Consider the traffic on your network when deciding which connections to log so that you do not clutter your displays or overwhelm your system. For more information, refer to Connection logging, on page 737.

Step 6 Apply the system-provided default health policy as described in Apply a health policy, on page 394.

Step 7 Customize a few of your system configuration settings:

- To allow inbound connections for a service such as SNMP or syslog, modify the relevant ports in the access list. For detailed instructions, refer to Configure an access list, on page 51.
- Edit your database event limits if needed. For additional information, refer to Configure database event limits, on page 68.
- To change the display language, edit the language setting as described in Set the language for the web interface, on page 85.
- To use a proxy server for restricted network access, edit your proxy settings as described in Modify Firewall Management Center management interfaces, on page 96.

Step 8 Customize your network discovery policy as described in *Configuring the Network Discovery Policy* in the *Cisco Secure Firewall Management Center Device Configuration Guide*. By default, the network discovery policy analyzes all network traffic. Limit discovery to RFC 1918 addresses in most cases.

Step 9 Consider customizing these other common settings:

- If you want to customize the default values for system variables, understand their use as described in *Variable Sets* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.
- If you want to create additional locally authenticated user accounts to access the Firewall Management Center, refer to Add or edit an internal user, on page 148.
- To allow access to the Firewall Management Center through LDAP or RADIUS external authentication, refer to External authentication configuration for Firewall Management Center, on page 150.

Deploy configuration changes; see the *Cisco Secure Firewall Management Center Device Configuration Guide*.

After you complete these steps, your system allows you to view data in dashboards and event tables and proceed with more feature customization if needed.

What to do next

Review and consider configuring other features described in Features, on page 7 and the rest of this guide.

Unsupported screens for the latest device version

You have reached this help page because the Firewall Management Center screen is a deprecated feature and is not supported on the latest version of device software.

- This guide includes features supported on the latest version of device software. Although the Firewall Management Center can manage devices running previous versions (as specified in the compatibility matrix available at *Cisco Secure Firewall Threat Defense Compatibility Guide*), refer to the appropriate guide for features supported only on older versions.
- Refer to the guide that matches your device version for features that are supported only on older device versions. Use the on-prem guide for other versions, as

Firewall Threat Defense devices

A Firewall Threat Defense device is a next-generation firewall (NGFW) that

- provides NGIPS capabilities
- supports site-to-site and remote access VPN, robust routing, NAT, clustering, and application inspection, and
- is available on a wide range of physical and virtual platforms.

Manager-device compatibility

This section explains compatibility between managers and devices, covering software, device models, virtual hosting environments, and operating systems.

Refer to the Release Notes for Secure Firewall Threat Defense, *Cisco Secure Firewall Management Center Compatibility Guide*, and *Cisco Secure Firewall Threat Defense Compatibility Guide* for more information.

Typical deployment scenario

In a typical deployment, multiple traffic-handling devices report to one Secure Firewall Management Center, which is used for administrative, management, analysis, and reporting tasks.

Features

The tables present commonly used features.

Understand appliance and system management features

To locate documents, refer to: Navigating the Cisco Secure Firewall Threat Defense Documentation.

This reference lists the main features for managing appliances and systems, such as user accounts, device health, backup and restore, upgrades, licensing, high availability, routing, VPN, multitenancy, REST API access, and troubleshooting.

Table 1: Appliance and system management features

If you want to...	Configure...	As described in...
Manage user accounts for logging in to your Secure Firewall devices	Device authentication	Users, on page 139 and *Users for Devices* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Monitor the health of system hardware and software	Health monitoring policy	Health monitoring, on page 379
Back up data on your appliance	Backup and restore	Backup/Restore, on page 471
Upgrade to a new version	System updates	Secure Firewall Threat Defense upgrade guides for Firewall Management Center Release Notes for Secure Firewall Threat Defense
Baseline your physical appliance	Restore to factory defaults (reimage)	Cisco FXOS Troubleshooting Guide for the Firepower 1000/2100 and Secure Firewall 3100/4200 with Threat Defense
Update the VDB, intrusion rule updates, or GeoDB on your appliance	Vulnerability Database (VDB) updates, intrusion rule updates, or Geolocation Database (GeoDB) updates	Updates, on page 255
Apply licenses in order to take advantage of license-controlled functionality	Smart licensing	Licenses, on page 271
Ensure continuity of appliance operations	Managed device high availability and/or Firewall Management Center high availability	*About Secure Firewall Threat Defense "High Availability chapter"* in the *Cisco Secure Firewall Management Center Device Configuration Guide* High availability in Firewall Management Center, on page 313
Configure a device to route traffic between two or more interfaces	Routing	*Reference for Routing* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Configure packet switching between two or more networks	Device switching	*Configure Bridge Group Interfaces* in the *Cisco Secure Firewall Management Center Device Configuration Guide*

If you want to...	Configure...	As described in...
Translate private addresses into public addresses for internet connections	Network Address Translation (NAT)	*Network Address Translation* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Establish a secure tunnel between managed Firewall Threat Defense devices	Site-to-Site virtual private network (VPN)	*VPN Overview* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Establish secure tunnels between remote users and managed Firewall Threat Defense devices	Remote Access VPN	*VPN Overview* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Segment user access to managed devices, configurations, and events	Multitenancy using domains	Multitenancy using domains, on page 243
View and manage appliance configuration using a REST API client	REST API and REST API Explorer	REST API preferences, on page 105 *Secure Firewall Mangement Center REST API Quick Start Guide*
Troubleshoot issues	N/A	Troubleshooting, on page 443

Use features to detect, prevent, and process potential threats

To locate documents, refer to: Navigating the Cisco Secure Firewall Threat Defense Documentation.

This reference lists the features and policies you can use to detect, prevent, and process potential threats on your network, along with links to configuration topics and guides.

Table 2: Features for detecting, preventing, and processing potential threats

If you want to...	Configure...	As described in...
Inspect, log, and take action on network traffic	Access control policy, the parent of several other policies	*Access Control Policy* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Block or monitor connections to or from IP addresses, URLs, or domain names, or both	Security Intelligence within your access control policy	*About Security Intelligence* in the *Cisco Secure Firewall Management Center Device Configuration Guide*

If you want to...	Configure...	As described in...
Control the websites that users on your network can access	URL filtering within your policy rules	*URL Filtering* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Monitor malicious traffic and intrusions on your network	Intrusion policy	*Intrusion Policy Basics* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Block encrypted traffic without inspection Inspect encrypted or decrypted traffic	SSL policy	*SSL Policies Overview* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Tailor deep inspection to encapsulated traffic and improve performance with fastpathing	Prefilter policy	*About Prefiltering* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Rate limit network traffic that is allowed or trusted by access control	Quality of Service (QoS) policy	*About QoS Policies* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Allow or block files (including malware) on your network	File/malware policy	*Network Malware Protection and File Policies* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Operationalize data from threat intelligence sources	Cisco Threat Intelligence Director (TID)	*Secure Firewall Threat Intelligence Director Overview* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Configure passive or active user authentication to perform user awareness and user control	User awareness, user identity, identity policies	*About User Identity Sources* in the *Cisco Secure Firewall Management Center Device Configuration Guide* *About Identity Policies* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Collect host, application, and user data from traffic on your network to perform user awareness	Network Discovery policies	*Network Discovery Policies* in the *Cisco Secure Firewall Management Center Device Configuration Guide*

If you want to...	Configure...	As described in...
Use tools beyond your device to collect and analyze data about network traffic and potential threats	Integration with external tools	Event Analysis Using External Tools, on page 635
Perform application detection and control	Application detectors	*Application Detection* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Troubleshoot issues	N/A	Troubleshooting, on page 443

Integrate with external tools

This reference describes how to integrate with external tools for automation, event streaming, database access, host input, event analysis, and troubleshooting in your network environment.

Table 3: Integration options with external tools

If you want to...	Configure...	As described in...
Automatically launch remediations when conditions on your network violate an associated policy	Remediations	Introduction to remediations, on page 1039 *Firepower System Remediation API Guide*
Stream event data from a Firewall Management Center to a custom-developed client application	eStreamer integration	eStreamer server streaming, on page 656 *Secure Firewall Mangement Center Event Streamer Integration Guide*
Query database tables on a Firewall Management Center using a third-party client	External database access	External database access, on page 72 *Secure Firewall Mangement Center Database Access Guide*
Augment discovery data by importing data from third-party sources	Host input	*Host Input Data* in the *Cisco Secure Firewall Management Center Device Configuration Guide* *Firepower System Host Input API Guide*
Investigate events with external event data storage tools and additional resources.	Integration with external event analysis tools	Event Analysis Using External Tools, on page 635
Troubleshoot issues	N/A	Troubleshooting, on page 443

Global search features in Firewall Management Center

A global search feature is a configuration navigation tool that

- enables users to quickly locate and access elements within the Firewall Management Center configuration
- supports searching for web interface pages, policy types, object types, and walkthroughs,
- provides search history, wildcard support, and results based on relevance.

Global search configuration entities

You can search the Firewall Management Center configuration for these entities:

- Names of web interface pages in top-level menus. (Refer to Search for web interface menu options, on page 15)
- For certain policy types:
  - Policy names
  - Policy descriptions
  - Rule names
  - Rule comments

  (Refer to Search for policies, on page 15)
- For certain object types:
  - Object names
  - Object descriptions
  - Configured values

  (Refer to Search for Objects, on page 17)
- How To walkthroughs.

  The search returns a list of walkthroughs that contain the search term, with links to each. (Refer to Search for how to walkthroughs, on page 21)

Global search usage details:

- When you open the global search tool, the ten most recent searches appear in a history list below the search text box. You can select an item from this list to repeat a search.
- When you type a search expression, the system updates the search results and replaces the search history. You do not need to press Enter to execute the search.
- Navigate the history list or search results using the mouse, keyboard arrow keys, or the Enter key. Press Enter to select the highlighted item in the search results. For web interface pages, the Firewall Management Center interface displays the highlighted page. For objects and policies, the interface displays details about the found entity.

- Search is not case-sensitive.
- You can use these wildcard characters in your search:
  - ? matches any single character.
  - * matches any 0 or more characters.
  - ^ anchors the search term it precedes to the beginning of matched entities.
  - $ anchors the search term it follows to the end of matched entities.

  Wildcards cannot be escaped.
- For greater efficiency, global search does not return indirect search results. It does not return policies or objects referencing those with the search term. To see which policies or objects reference found items, view the **Usages** tab in the search detail pane.
- Global search returns the top results for your search expression determined by its relevance to the most commonly used configuration entities in the Firewall Management Center. If global search does not return a result you expect, try refining your search. You can also use the search or filter tool at the top of many GUI pages, or explore configuration-specific search features that the web interface offers:
  - *Searching for Rules* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
  - *Searching and Filtering the NAT Rule Table* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
  - Searching for Events
  - Searching Custom Tables

Note This feature is supported in Light and Dusk themes only. To change the theme, refer to Change the web interface appearance, on page 233.

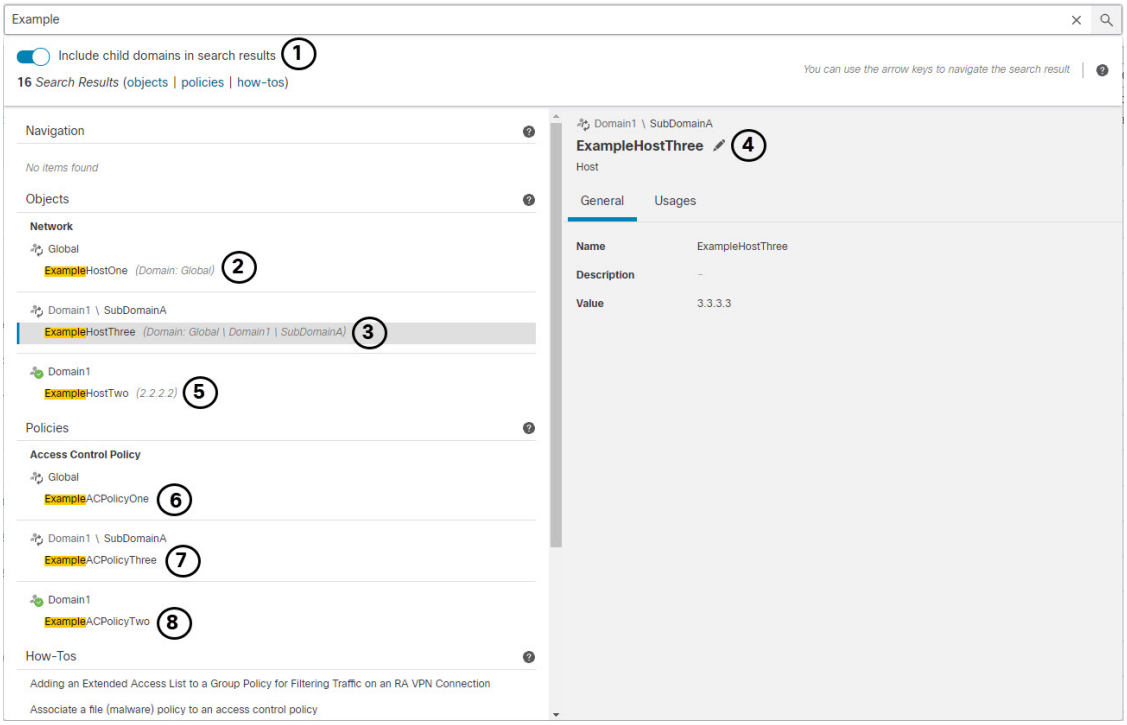
Global search in a multidomain deployment

In a multidomain deployment, by default search returns only objects and policies defined within the current domain and its ancestor domains. You can view objects and policies in child domains by toggling an option in the search results dialog.

For an object search, if your search expression is found in objects defined in domains other than your current domain, the search results display the names of the domains within which those objects reside. If your search expression is found in objects defined within your current domain, the search results display the object values.

The screenshot demonstrates a deployment with three domains at three levels: Global, Domain1, and SubDomainA. The user, with Domain1 as the current domain, has entered a search for the string “example” in both ancestor and child domains.

Figure 1: Example of global search in a multidomain environment



1	The user has chosen to search child domains (SubDomainA) as well as the current domain (Domain1) and its ancestor (Global).	2	A matching network object ExampleHostOne defined in the parent domain Global is displayed with the domain name, and the **External Domain** ( ) icon indicating the user must switch domains to edit details.
3	The matching network object ExampleHostThree defined in the child domain SubDomainA is displayed with the domain name, and the **External Domain** ( ) icon indicating the user must switch domains to edit details. This object is currently selected.	4	The matching network object ExampleHostThree is currently selected, and information is provided in the right pane. The **External Domain** ( ) icon indicates that when the user clicks **Edit** ( ), the system will prompt the user to confirm a domain change before allowing edit access to the object.
5	The matching network object ExampleHostTwo, defined in the current domain, is displayed with the object value, and with the **Current Domain** ( ) icon indicating the user may edit this object without switching domains.	6	The matching access control policy ExampleACPolicyOne defined in the parent domain Global is displayed with the domain name, and the **External Domain** ( ) icon indicating the user must switch domains to edit details.

7	The matching access control policy ExampleACPolicyThree defined in the child domain SubDomainA is displayed with the domain name, and the **External Domain** ( ) icon indicating the user must switch domains to edit details.	8	The matching access control policy ExampleACPolicyTwo defined in the current domain is displayed with the **Current Domain** ( ) icon indicating the user may edit details without switching domains.

Search for web interface menu options

Search in the Secure Firewall Management Center web interface to quickly find menu options and pages. Use the search function to navigate efficiently and access configuration pages.

Search the top-level menus in the web interface to locate the pages you need.

To view or configure Quality of Service settings, search for `QoS`. The search results update as you type and are grouped by category.

Before you begin

This feature is unavailable in the Classic theme. To change the theme, refer to Change the web interface appearance, on page 233. Ensure you use a supported theme for search functionality.

Follow these steps to search for web interface menu options:

Procedure

Step 1 Use one of the two methods to initiate a search:

- In the menu bar at the top of the Firewall Management Center web interface, click **Search** ( ).
- With focus outside of a text box, type the forward slash (/) .

Step 2 Enter one or more letters of the menu option's name. The search results appear and update as you type. You do not need to press Enter.

Step 3 Search results appear grouped by category. To go to a page listed under **Navigation**, click the menu path in the search results list.

After completing these steps, you can access desired pages or menu options in the Secure Firewall Management Center web interface. The search results display in groups by category, which allows efficient navigation.

Search for policies

Configure global search to efficiently locate policies by name or rule comments in the web interface. This task helps you find policy types and their details so you can streamline policy management and review.

You can search for specific policy types by name using global search. This table shows which policy types are in scope for search and which are not:

In Scope	Out of Scope
Access Control Policy	Threat Defense Platform Settings
Prefilter Policy	Firepower Settings Policy
Threat Defense NAT Policy	Firepower NAT Policy
Intrusion category • Intrusion Policy • Network Analysis Policy	QoS Policy FlexConfig Policy DNS Policy Malware and File Policy SSL Policy Identity Policy Network Discovery Application Detector Correlation Policy VPN category • Dynamic Access Policy • Site-to-site • Remote Access

Global search returns policies whose names match the search term, as well as access control policies using rules whose name or comments match the search term. If you view an access control policy in the search result list whose name does not match your search, the match occurs because the name or comments for a rule configured within the policy matched your search term.

- Global search returns the top results for your search expression determined by its relevance to the most commonly used configuration entities in the Firewall Management Center.
- Your search term may exist in policy types that are not in scope for this search feature.
- For a full description of the global search feature and alternative search methods, refer to Search the Firewall Management Center.

Before you begin

This feature is not available in the Classic theme. To change the theme, refer to Change the web interface appearance, on page 233.

Follow these steps to search for policies using global search:

Procedure

Step 1 Use one of the two methods to initiate a search:

- In the menu bar at the top of the Firewall Management Center web interface, click **Search** ( ).
- With focus outside of a text box, type forward slash (/).

Step 2 Enter a search expression in the search text box. Search results are displayed under the text box and update as you type. You do not need to press Enter to execute the search.

Step 3 (Optional) In a multidomain deployment, if your current domain has descendant domains, you can toggle **Include child domains in search results** to view policies in those descendant domains.

Step 4 Search results appear grouped by category. In a multidomain deployment, within the **Policies** category the search results are grouped by the domains within which found policies are defined. Under the **Policies** category, you can:

To:	Do this:
View search results for a single policy type.	Click the policy type in the search results, such as Access Control Policy.
View details about a policy.	Click the policy name in the search results list to view the details pane and display the **General** tab.
View the Access Control policies that reference Intrusion and Network Analysis policies.	Click the name of the Intrusion or Network Analysis policy in the search results to view the details pane and display the **Usages** tab.
Open the policy configuration page for a policy in a separate browser window.	Click the policy name in the search results. In the details pane, click the **Edit** ( ). In a multidomain deployment, if you choose to edit a policy not defined within your current domain the system will prompt you to change your current domain.

After completing these steps, you will be able to locate and access policies and their details using global search. Search results are grouped by category and domain, allowing for efficient policy management.

Search for Objects

This table indicates which object types listed on the Object Management page (**Objects** > **Object Management**) are in scope for the global search feature:

In Scope	Out of Scope
AAA Server category • RADIUS Server Group • Single Sign-On Server Access List category • Extended Access List • Standard Access List Address Pools category • IPv4 Pools • IPv6 Pools AS Path Community List category • Extended Community DNS Server Group External Attributes Category • Dynamic Object • Security Group Tag Geolocation Interface category • Security Zone • Interface Group Key Chain Network (includes Network, Host, Range, FQDN, Network Group) PKI category Cert Enrollment Policy List Port (objects and groups, TCP, UDP, ICMP, ICMP6, other)	Application Filters Cipher Suite List Community List Category • Community Distinguished Name category • Individual Distinguished Name Objects • Distinguished Name Object Groups File List FlexConfig category • FlexConfig Object • Text Object PKI category • External Cert Groups • External Certs • Internal CA Groups • Internal CAs • Internal Cert Groups • Internal Certs • Trusted CA Groups • Trusted CAs Security Intelligence category • DNS Lists and Feeds • Network Lists and Feeds • URL Lists and Feeds Sinkhole Variable Set

In Scope	Out of Scope
Prefix List category • IPV4 Prefix List • IPV6 Prefix List Route Map SLA Monitor Time Range Time Zone Tunnel Zone URL (Objects, groups) VLAN Tag (Objects, groups) VPN category • Certificate Map • Group Policy • IKEv1 IPsec Proposal • IKEv1 Policy • IKEv2 IPSec Proposal • IKEv2 Policy	VPN category • AnyConnect File • Custom Attribute

Global search shows objects whose names, descriptions, or configured values match your search term. If you find an object in the results and its name does not match your search, the match comes from the description or another configured value.



Important Global search returns the top results for your search expression, based on relevance to the most commonly used configuration entities in the Firewall Management Center. Your search term may exist in object types that are not in scope for this search feature. For a full description of the global search feature and alternative search methods, refer to Search the Firewall Management Center .

Object searches can be particularly useful when you need to locate network information within your deployment. You can search for these in object names, descriptions, or configured values:

- IPv4 and IPv6 address information, including these formats:
  - Full addresses (For example, `194.164.0.23`, `2001:0db8:85a3:0000:0000:8a2e:0370:7334`.)
  - Partial addresses (For example, `194.164, 2001:db8`.)

  - Ranges (For example, `198.51.1.1-198.51.1.5` or `2001:db8::0202-2001:db8::8329`. Do not add a space before or after the hyphen.) Global search returns objects using network addresses that match any within the specified range.
  - CIDR notation. (For example `198.51.1.0/24`, `2002::1234:abcd:ffff:101/64`.) Global search returns objects using network addresses that match any within the specified CIDR block.
- Port information:
  - Port numbers (For example, `22` or `80`.)
  - Protocols. (For example, `https` or `ssh`.)
- Fully qualified domain names. (For example, `www.cisco.com`.)
- URLs. (For example, `http://www.cisco.com`.)
- Encryption standards or hash types. (For example, `AES-128` or `SHA`.)
- VLAN tag numbers. (For example, `568`.)

Before you begin

This feature is not available in the Classic theme. To change the theme, refer to Change the web interface appearance, on page 233.

Procedure

Step 1 Use one of two methods to initiate a search:

- In the menu bar at the top of the Firewall Management Center web interface, click **Search** ( ) .
- With focus outside of a text box, type `/` (forward slash).

Step 2 Enter a search expression in the search text box. Search results are shown in the text box and update as you type. There is no need to press Enter to execute the search.

If your search expression is found in objects defined in domains other than your current default domain, the search results display the names of the domains within which those objects reside. If your search expression is found in objects defined within your current domain, the search results display the object values.

Step 3 (Optional) In a multidomain deployment, if your current domain has descendant domains, you can toggle **Include child domains in search results** to view objects in those descendant domains.

Step 4 Search results appear grouped by category. In a multidomain deployment, within the **Objects** category, the results are organized by the domains in which the objects are defined. Under the **Objects** category you can:

To:	Do this:
View search results for a single object type.	Click the object type in the search results, such as **Network**.
View details about an object in the search results.	Click the object name in the search results to view the details pane and display the **General** tab.

To:	Do this:
View a list of policies or objects that use an object in the search results.	Click the object name in the search results to view the details pane and display the **Usages** tab. **Note** Global search does not provide usage information for all object types.
Open the object configuration page for an object in a separate browser window.	Click the object name in the search results, and in the details pane click **Edit** ( ). In a multidomain deployment, if you choose to edit an object not defined within your current domain the system will prompt you to change your current domain.

Search for how to walkthroughs

Search for How To walkthroughs to access step-by-step instructions for your tasks. This feature helps you find guides for device setup and other procedures quickly.

This task is relevant when you need guidance on specific procedures or want to find walkthroughs for tasks such as device setup.

Search for How To walkthroughs about tasks you are interested in. For example, to find walkthroughs for device setup, search for the word "device." Use this task when you need clear instructions for a specific action.

Procedure

Step 1 Use one of two methods to initiate a search:

- In the menu bar located at the top of the Firewall Management Center web interface, click **Search** ( ).
- With focus outside of a text box, type the forward slash (/).

Step 2 Enter a search term associated with a task for which you would like to view a walkthrough. Search results appear near the text box and update as you type. You do not need to press Enter to execute the search.

Step 3 Search results appear grouped by category. To view a walkthrough listed under **How-Tos**, click the walkthrough title in the search results list. For more information on How To walkthroughs, refer to Online help, How To, and documentation, on page 25.

After completing these steps, you will view a list of How To walkthroughs relevant to your search term, grouped by category. You can access detailed instructions for the selected task.

Switch domains on Secure Firewall Management Center

Switching domains enables users to access data and menu options specific to each domain in a multidomain deployment.

- Configure access to multiple domains for a single user account.
- Assign different privileges for each domain to the same user.

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

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

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

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

When you switch to a domain, the system displays:

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

Before you begin

You do not need explicit prerequisites to switch domains. However, your user account must be associated with multiple domains and have the necessary privileges.

- Ensure your user account is assigned to more than one domain.
- Verify you have privileges in each domain you want to access.

Follow these steps to switch domains on Secure Firewall Management Center:

Procedure

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

The system displays data and menu options relevant to the selected domain, based on your assigned privileges.

What to do next

No additional steps are required after switching domains. You can repeat the process to switch to another domain as needed.

Context menus

A context menu is a web interface feature that

- provides shortcuts for accessing other features through right-click or left-click actions
- contains contents that depend on the specific page and data where you access it, and
- appears as the normal browser context menu on pages or locations that do not support it.

Context menu capabilities by location

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

For example:

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

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

Specific context menu capabilities by location:

- Policy Editors: Many policy editors contain hotspots over each rule. You can insert new rules and categories; cut, copy, and paste rules; set the rule state; and edit the rule.
- Intrusion Rules Editor: The intrusion rules editor contains hotspots over each intrusion rule. You can edit the rule, set the rule state, configure thresholding and suppression options, and view rule documentation. Optionally, after clicking **Rule documentation** in the context menu, you can click **Rule Documentation** in the documentation pop-up window to view more-specific rule details.
- Event Viewer: Event pages (the drill-down pages and table views available under the Analysis menu) contain hotspots over each event, IP address, URL, DNS query, and certain files' SHA-256 hash values. While viewing most event types, you can:
  - View related information in the Context Explorer.
  - Drill down into event information in a new window.
  - View the full text in places where an event field contains text too long to fully display in the event view, such as a file's SHA-256 hash value, a vulnerability description, or a URL.
  - Open a web browser window with detailed information about the element from an external source, using the Contextual Cross-Launch feature. For more information, see Event investigation using web-based resources, on page 638.
- Intrusion Event Packet View: Intrusion event packet views contain IP address hotspots. The packet view uses a left-click context menu.
- Dashboard: Many dashboard widgets contain hotspots to view related information in the Context Explorer. Dashboard widgets can also contain IP address and SHA-256 hash value hotspots.

- Context Explorer: The Context Explorer contains hotspots over its charts, tables, and graphs. If you want to examine data from graphs or lists in more detail than the Context Explorer allows, you can drill down to the table views of the relevant data. You can also view related host, user, application, file, and intrusion rule information. The Context Explorer uses a left-click context menu, which also contains filtering and other options unique to the Context Explorer.

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

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

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

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

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

- Edit the triggering rule.
- Set the rule state, including disabling the rule.
- Configure thresholding and suppression options.
- View rule documentation. Optionally, after clicking **Rule documentation** in the context menu, you can click **Rule Documentation** in the documentation pop-up window to view more-specific rule details.

Data sharing features with Cisco

Data sharing features in Firewall Management Center allow you to share usage metrics and analytics with Cisco. This data helps Cisco analyze product usage and improve the user experience. Data sharing features include Cisco Success Network, Cisco Support Diagnostics, and web analytics.

You can opt to share data with Cisco using these features:

- Cisco Success Network

  Refer to Configure Firewall Management Center to share usage metrics and statistics with Cisco, on page 48

- Web analytics

  Refer to Configure web analytics, on page 134

The Cisco Success Network and Cisco Support Diagnostic capabilities are enabled by default.

The Cisco Success Network capability collects customer usage metrics and statistics. Cisco analyzes product usage data and improves user experience based on the collected metrics. To opt out of sending Cisco Success Network telemetry data to Cisco, refer to Configure Firewall Management Center to share usage metrics and statistics with Cisco, on page 48. For more information about the telemetry data that Cisco collects, refer to Cisco Success Network Telemetry Data Collected from the Management Center Devices.

The Cisco Support Diagnostics capability collects essential information from your devices. Cisco uses this information to enhance your support experience. To opt out of sending Cisco Support Diagnostics metrics to Cisco, refer to Configure Firewall Management Center to share device health data with Cisco, on page 49.

You can opt to share data with Cisco using web analytics. For more information, refer to Configure web analytics, on page 134.

Online help, How To, and documentation

Online help, How Tos, and documentation provide support resources that

- provide context-sensitive assistance directly from the web interface
- include documentation roadmaps for additional reference information.

Access methods for online help and documentation

You can reach online help from the web interface:

- By clicking the context sensitive help link on each page.
- By choosing **Help** ( ) > **Page-level Help**.

You can find additional documentation using the documentation roadmap:

Navigating the Cisco Secure Firewall Threat Defense Documentation

How To features

How To is a widget that provides walkthroughs to help you navigate tasks on the Firewall Management Center. The walkthroughs guide you through each required step. You may need to use different UI screens to complete the task. The **How To** widget is enabled by default. To disable the widget, choose **User Preferences** from the drop-down list under your user name, and uncheck the **Enable How-Tos** check box in **How-To Settings**.

To open the How To widget, choose **Help** ( ) > **On-screen Assistance** > **How-Tos**.

For a list of walkthroughs available in your Firewall Management Center, refer to Walkthroughs in Secure Firewall Management Center.

Note Walkthroughs are generally available for all UI pages and do not depend on user role. However, if your privileges do not allow access, some menu items do not appear on the Firewall Management Center interface. As a result, you cannot run the walkthroughs on these pages.

User guides on Cisco.com

You can find documents for configuring Secure Firewall Threat Defense at Navigating the Cisco Secure Firewall Threat Defense Documentation.

License statements in the documentation

A license statement is a documentation element that

- indicates which Classic or Smart license you must assign to a managed device to enable the described feature
- shows only the highest required license for each feature because licensed capabilities are often additive, and
- uses "or" statements to specify that a particular license is required, but an additional license can add functionality.

License statement usage details

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

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

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

For more information about licenses, refer to Licenses, on page 271.

Related Topics

Licenses, on page 271

Supported devices statements in the documentation

A Supported Devices statement is a documentation element that

- appears at the beginning of a chapter or topic
- indicates that a feature is supported only on specific device series, families, or models,
- helps users identify feature support for specific devices such as Secure Firewall Threat Defense.

Platform support reference information

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

Access statements in the documentation

An access statement is a documentation element that

- indicates the predefined user roles required to perform a procedure
- allows any of the listed roles to perform the procedure, and
- provides guidance regarding access permissions for users with custom roles.

User roles and access permissions

Custom roles may have permission sets that differ from predefined roles. If you have a custom role with similar permissions, you also have access whenever a predefined role is listed for that procedure.

Some users with custom roles may use slightly different menu paths to reach configuration pages. For example, users who have a custom role with only intrusion policy privileges access the network analysis policy through the intrusion policy instead of the standard path through the access control policy.

IP address conventions

The system uses a set of IP address conventions to interpret IP address blocks entered using IPv4 CIDR notation or IPv6 prefix length notation, including the addresses that are not entered on the standard network boundary.

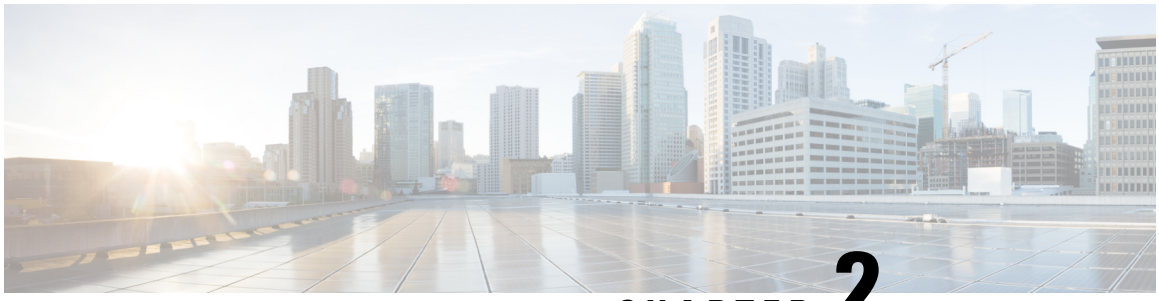
- You can use IPv4 Classless Inter-Domain Routing (CIDR) notation and the similar IPv6 prefix length notation to define address blocks in many places in the system.
- When you use CIDR or prefix length notation to specify a block of IP addresses, the system uses **only** the portion of the network IP address specified by the mask or prefix length. For example, if you type 10.1.2.3/8, the system uses 10.0.0.0/8.
- Cisco recommends using network IP addresses on the bit boundary when employing CIDR or prefix length notation, but the system does not enforce this requirement.

Additional Resources

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

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

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



CHAPTER 2

Logging into the Management Center

The following topics describe how to log into the system:

- User accounts, on page 29
- System user interfaces, on page 31
- Log into the Secure Firewall Management Center web interface, on page 33
- Log into the Firewall Management Center web interface using SSO, on page 34
- Log into the Secure Firewall Management Center with CAC credentials, on page 36
- Log into the Firewall Management Center Command Line Interface, on page 37
- View your last login, on page 37
- Logging Out of the Firewall Management Center Web Interface, on page 38
- History for logging into the management center, on page 38

User accounts

You must provide a username and password to obtain local access to the web interface or CLI on Firewall Management Center or a managed device. On managed devices, CLI users with Config level access can use the **expert** command to access the Linux shell. On the Firewall Management Center, all CLI users can use the **expert** command. The Firewall Threat Defense and Firewall Management Center can be configured to use external authentication, storing user credentials on an external LDAP or RADIUS server; you can withhold or provide CLI access rights to external users. The Firewall Management Center can be configured to support Single Sign-On (SSO) using any SSO provider conforming to the Security Assertion Markup Language (SAML) 2.0 open standard for authentication and authorization.

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

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

Secure Firewall Management Centers

Secure Firewall Management Centers support these user account types:

- A pre-defined **admin** account for web interface access, which has the administrator role and can be managed through the web interface.

- Custom user accounts, which provide web interface access and which **admin** users and users with administrator privileges can create and manage.
- A pre-defined **admin** account for CLI access. Users logging in with this account can use the **expert** command to gain access to the Linux shell.

  During initial configuration, the passwords for the CLI **admin** account and the web interface **admin** account are synchronized but, optionally, thereafter you can configure separate passwords for the two **admin** accounts.

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

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

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

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

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

Secure Firewall Threat Defense and Secure Firewall Threat Defense Virtual Devices

Secure Firewall Threat Defense and Secure Firewall Threat Defense Virtual devices support these user account types:

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

The Secure Firewall Threat Defense supports external authentication for SSH users.

System user interfaces

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

For information on browser requirements, see the Secure Firewall Release Notes.



Note On all appliances, after a user makes three consecutive failed attempts to log into the CLI via SSH, the system terminates the SSH connection.

Appliance	Web-Based GUI	Auxiliary CLI	Linux Shell
Secure Firewall Management Center	• Supported for predefined **admin** user and custom user accounts. • Can be used for administrative, management, and analysis tasks.	• Supported for predefined **admin** user and custom external user accounts. • Accessible using an SSH, serial, or keyboard and monitor connection. • Should be used only for administration and troubleshooting directed by Cisco TAC.	• Supported for predefined **admin** user. • Must be accessed via **expert** command from the Secure Firewall Management Center CLI. • Accessible using an SSH, serial, or keyboard and monitor connection. • Should be used only for administration and troubleshooting directed by Cisco TAC or by explicit instructions in the Firewall Management Center documentation.
Secure Firewall Threat Defense Secure Firewall Threat Defense Virtual	—	• Supported for predefined **admin** user and custom user accounts. • Accessible in physical devices using an SSH, serial, or keyboard and monitor connection. Accessible in virtual devices via SSH or VM console. • Can be used for setup and troubleshooting directed by Cisco TAC.	• Supported for predefined **admin** user and custom user accounts. • Accessible by CLI users with Config access using the **expert** command. • Should be used only for administration and troubleshooting directed by Cisco TAC or by explicit instructions in the Firewall Management Center documentation..

Related Topics

Add or edit an internal user, on page 148

Web interface considerations

Web interface considerations are a set of guidelines that

- address authentication requirements for Common Access Cards (CACs) when external users access the appliance through LDAP
- explain how menu options and privileges affect user access and system logging, and
- describe browser behavior when processing time-intensive operations.

Web interface operational details

The web interface has specific operational characteristics:

- If your organization uses Common Access Cards (CACs) for authentication, external users authenticated with LDAP can use CAC credentials to obtain access to the web interface of an appliance.
- Menus and menu options at the top of the default home page are based on your user account privileges. Links on the home page offer choices from the full range of user privileges. If you select a link requiring privileges you do not have, the system shows a warning and records the event.
- Some processes that take a significant amount of time may cause your web browser to display a message that a script has become unresponsive. If this occurs, make sure you allow the script to continue until it finishes.

Related Topics

Specify your home page, on page 234

Session timeout

A session timeout is a security mechanism that

- automatically logs users out of a session after 1 hour of inactivity by default
- can be configured to exempt certain users from automatic logout, and
- allows Administrator role users to modify the timeout interval for an appliance.

Session timeout behavior

By default, the system automatically logs you out of a session after 1 hour of inactivity unless you are exempt from session timeout by configuration.

Note For SSO users, when the Firewall Management Center session times out, the display briefly redirects to the IdP interface, and then the Firewall Management Center login page. Unless the SSO session has been terminated from elsewhere, anyone can access the Firewall Management Center without providing login credentials simply by clicking on the **Single Sign-On** link on the login page. To ensure Firewall Management Center security and prevent others from accessing the Firewall Management Center using your SSO account, we recommend you not leave a Firewall Management Center login session unattended, and log out of the SSO federation at the IdP when you log out of the Firewall Management Center.

Users with the Administrator role can change the session timeout interval for an appliance using **System** (⚙) > **Configuration** > **Session Timeout**.

Related Topics

Configure session timeouts, on page 121
SAML Single Sign-On, on page 169

Log into the Secure Firewall Management Center web interface

This task allows you to access the Secure Firewall Management Center web interface to manage and configure your firewall system.

Note This task applies to internal users and external users authenticated by LDAP or RADIUS servers. For SSO login, see Log into the Firewall Management Center web interface using SSO, on page 34.

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

In a NAT environment where multiple Firewall Management Centers share the same IP address:

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

Before you begin

- If you do not have access to the web interface, contact your system administrator to modify your account privileges, or log in as a user with Administrator access and modify the privileges for the account.
- Create user accounts as described in Add or edit an internal user, on page 148.

Follow these steps to log into the Secure Firewall Management Center web interface:

Procedure

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

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

- User names are *not* case-sensitive.
- In a multidomain deployment, prepend the user name with the subdomain where your user account was created. You do not need to specify the Global domain. For example, if your user account was created in SubdomainA, enter your username in the following format:

  `SubdomainA\username`

  If your user was added to SubdomainB, whose parent domain is SubdomainA, enter your username in the following format:

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

Step 3 Click **Login**.

You are successfully logged into the Secure Firewall Management Center web interface and can begin managing your firewall configuration.

Related Topics

Session timeout, on page 32

Log into the Firewall Management Center web interface using SSO

This task enables secure access to the Firewall Management Center web interface using existing SSO credentials, eliminating the need for separate authentication.

The Firewall Management Center can be configured to participate in any Single-Sign On (SSO) federation implemented with an SSO provider conforming to the Security Assertion Markup Language (SAML) 2.0 open standard. SSO user accounts must be established at the identity provider (IdP) and must use email addresses for their account names. If your user name is not an email address, or SSO login fails, contact your system administrator.

Note The Firewall Management Center does not support logging in with CAC credentials for SSO accounts.

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

In a NAT environment where multiple Firewall Management Centers share the same IP address:

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

Before you begin

- Configure the Firewall Management Center for SSO access. See SAML Single Sign-On, on page 169.
- If you do not have access to the web interface, contact your system administrator to configure your account at the SSO IdP.

Follow these steps to log into the Firewall Management Center web interface using SSO:

Procedure

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

Note
SSO users must consistently access the Firewall Management Center using the login URL specifically configured for SSO access; ask your administrator for this information.

Step 2 Click the **Single Sign-On** link.

Step 3 The system may display one of two responses:

- If you are already logged into the SSO federation, the Firewall Management Center default home page appears.
- If you are not already logged into the SSO federation, the Firewall Management Center redirects your browser to the login page for your IdP. After you complete the login process at the IdP, the Firewall Management Center default home page appears.

You are successfully logged into the Firewall Management Center web interface and can access the default home page to begin managing your firewall configurations.

Related Topics

Session timeout, on page 32
SAML Single Sign-On, on page 169

Log into the Secure Firewall Management Center with CAC credentials

This task enables secure access to the Secure Firewall Management Center web interface using CAC authentication credentials.

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

In a NAT environment where multiple Firewall Management Centers share the same IP address:

- Each Firewall Management Center can support only one login session at a time.
- To access different Firewall Management Centers, use a different browser for each login, such as Firefox or Chrome, or set the browser to incognito or private mode.



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

Before you begin

- If you do not have access to the web interface, ask your system administrator to change your account privileges. Alternatively, log in as a user with administrator access to modify account privileges.
- Create user accounts. For more information, refer to Add or edit an internal user, on page 148.
- Configure CAC authentication and authorization. For more information, refer to Configure common access card authentication with LDAP, on page 167.

Follow these steps to log into the Secure Firewall Management Center with CAC credentials:

Procedure

Step 1 Insert a CAC as instructed by your organization.

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

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

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

Step 5 Click **Continue**.

You are successfully logged into the Secure Firewall Management Center web interface using CAC authentication.

Related Topics

Configure common access card authentication with LDAP, on page 167

Session timeout, on page 32
Guidelines for SSO configuration on Firewall Management Center, on page 170

Log into the Firewall Management Center Command Line Interface

This task enables CLI access to the Firewall Management Center for administrative and external users.

The **admin** CLI user and certain custom external users can log into the Firewall Management Center CLI.



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



Note For all appliances, after a user makes three consecutive failed attempts to log into the CLI via SSH, the system terminates the SSH connection.

Before you begin

Complete the initial configuration process as the **admin** user. See First-time login, on page 3.

Procedure

Step 1 Use the **admin** user name and password to connect to the Firewall Management Center via SSH or the console port.

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

Step 2 Use any of the available CLI commands.

You are now logged into the Firewall Management Center CLI and can execute administrative commands.

View your last login

View your last login information to verify that no unauthorized users have accessed your account using your credentials.

If you suspect that an unauthorized user has used your credentials to sign in to the Secure Firewall Management Center, you can see the date, time, and IP address from which your credentials were last used to log in.

Before you begin

This feature is not available if you are using the Classic theme. You can select a UI theme in User Preferences.

Procedure

Step 1 Sign in to the Secure Firewall Management Center.

Step 2 At the top right corner of your browser window, look for the User ID that you used to sign in.

Step 3 Click your user name.

Step 4 Information about your previous login is shown at the bottom of the menu that appears.

You can view the date, time, and IP address of your last login to verify account security.

Logging Out of the Firewall Management Center Web Interface

Logging out ends your web session and ensures that no one can use the interface with your credentials.

When you are no longer actively using the Firewall Management Center web interface, Cisco recommends that you log out, even if you are only stepping away from your web browser for a short period of time.



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

Procedure

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

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

Your web session is terminated and you are logged out of the Firewall Management Center web interface.

Related Topics

Session timeout, on page 32

History for logging into the management center

This reference provides a comprehensive history of login-related features and enhancements for the Management Center across different software versions. The information includes feature descriptions, supported versions, platform compatibility, and implementation details for Single Sign-On capabilities, CLI access controls, login security measures, and user interface modifications.

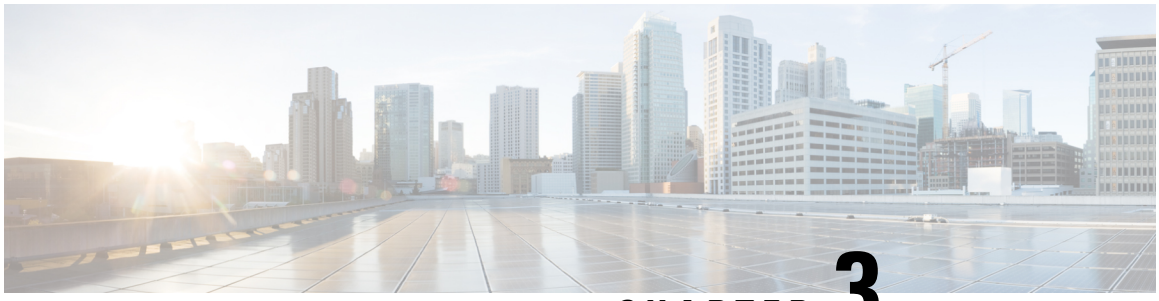
Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Added support for Single Sign-On (SSO) using any SAML 2.0-compliant SSO provider.	6.7	Any	Added the ability for users configured at any third-party SAML 2.0-compliant identity provider (IdP) to log into the Firewall Management Center using a new **Single Sign-On** link on the login page. New/Modified screen: Login screen
View information about the last time you signed in to the Secure Firewall Management Center	6.5	Any	View the date, time, and IP address from which you last logged in. New/Modified menus: The menu at the top right of the window that shows the username that you used to log in. Supported platforms: Firewall Management Center
Automatic CLI access for the Firewall Management Center	6.5	Any	When you use SSH to log into the Firewall Management Center, you automatically access the CLI. You can then use the CLI `expert` command to access the Linux shell, however, doing so is strongly discouraged.". **Note** This feature deprecates the Version 6.3 ability to enable and disable CLI access for the Firewall Management Center. As a consequence of deprecating this option, the virtual Firewall Management Center no longer displays the **System** > **Configuration** > **Console Configuration** page, which still appears on physical Firewall Management Centers.
Limit number of SSH login failures	6.3	Any	When a user accesses any device via SSH and fails three successive login attempts, the device terminates the SSH session.
Ability to enable and disable CLI access for the Firewall Management Center	6.3	Any	New/Modified screens: New check box available to administrators in Firewall Management Center web interface: **Enable CLI Access** on the **System** > **Configuration** > **Console Configuration** page. • Checked: Logging into the Firewall Management Center using SSH accesses the CLI. • Unchecked: Logging into Firewall Management Center using SSH accesses the Linux shell. This is the default state for fresh Version 6.3 installations as well as upgrades to Version 6.3 from a previous release. Supported platforms: Firewall Management Center



PART II

System Settings

- System Configuration, on page 43
- Users, on page 139
- Domains, on page 243
- Updates, on page 255
- Licenses, on page 271
- High availability, on page 313
- Security certifications compliance, on page 347



CHAPTER 3

System Configuration

This chapter explains how to configure system configuration settings on the Secure Firewall Management Center.

- Integrate Firewall Management Center with Cisco Security Cloud, on page 44
- Meet requirements and prerequisites for the system configuration, on page 50
- Manage the Secure Firewall Management Center system configuration, on page 50
- Access lists, on page 51
- Configure access control preferences, on page 52
- Audit logs, on page 53
- Audit log certificates, on page 57
- Change reconciliation, on page 64
- DNS caches, on page 66
- Dashboards, on page 67
- Database, on page 67
- Email notifications, on page 71
- External database access, on page 72
- HTTPS certificates, on page 74
- View system information fields, on page 83
- Intrusion policy preferences, on page 84
- Languages, on page 85
- Login banners, on page 86
- Management interfaces, on page 87
- Network analysis policy preferences, on page 103
- Processes, on page 104
- REST API preferences, on page 105
- Remote console access management, on page 107
- Remote storage devices, on page 114
- SNMP, on page 119
- Session timeouts, on page 120
- Configure time settings, on page 121
- Time synchronization, on page 123
- UCAPL/CC compliance, on page 128
- Configure user settings, on page 128
- VMware tools, on page 132

• Vulnerability mapping, on page 133
• Configure web analytics, on page 134
• History for system configuration, on page 135

Integrate Firewall Management Center with Cisco Security Cloud

An integration of Firewall Management Center with Cisco Security Cloud is a security management solution that

- connects firewall deployments to Cisco's integrated security cloud services for a consistent experience,
- unifies visibility, enables automation, and strengthens security across network, endpoints, and applications, and
- offers a platform approach with simpler, more integrated cloud services that reduce the complexity of managing multiple products.

Integration features

The integration provides these capabilities:

- A centralized view of inventory across Firewall Management Centers.
- Zero-Touch Provisioning for Firewall Threat Defense.
- Cross-launching to the Firewall Management Center to manage devices and objects.
- Help with establishing consistent policies across Firewall Management Centers.
- Cloud event storage and services to enrich threat hunts and investigations.

Use your Security Cloud Control account to authorize and register (onboard) your Firewall Management Center; Onboard an On-Prem Management Center.

To integrate the Secure Firewall Management Center with Cisco XDR, refer to the Cisco Secure Firewall Management Center and Cisco XDR Integration Guide.

Enable SecureX integration

Enable SecureX integration to connect your management center and managed devices to a Cisco Defense Orchestrator (CDO) tenant, allowing for centralized cloud-based management and enhanced security capabilities.

- Centralizes device management and event forwarding through the cloud.
- Provides access to Cisco Security Cloud features such as AI Assistant, Policy Analyzer, XDR Automation, and Zero-Touch Provisioning.

Use this task when you need to onboard your management center and its managed devices to a CDO tenant for integration with Cisco Security Cloud services.

This integration is relevant for organizations seeking to leverage Cisco cloud-based security features, streamline device management, and enable advanced automation and analytics. Perform this task during initial setup or when adding new devices to your Cisco Security Cloud environment.

Before you begin

- Security Cloud Control uses Cisco Security Cloud Sign On as its identity provider and Duo for multifactor authentication. Ensure that you have your Cisco Security Cloud Sign On credentials and can sign in to the Cisco regional cloud where your account was created.
- You need a Security Cloud Control tenant to integrate the Secure Firewall Management Center with Cisco Security Cloud. If you do not already have aSecurity Cloud Control tenant, request one or create a tenant during this workflow. For more information, refer to Request a Security Cloud Control Tenant.
- Link your Security Cloud Control tenant, the one you want to use for onboarding the management center, to your Security Services Exchange (SSE) account. For more information, refer to Link Your Security Cloud Control and Cisco XDR Tenant Accounts.

Follow these steps to enable SecureX integration:

Procedure

Step 1 In the Secure Firewall Management Center, choose **Integration** > **SecureX**.

Step 2 Choose a Cisco regional cloud from the **Current Region** drop-down list.

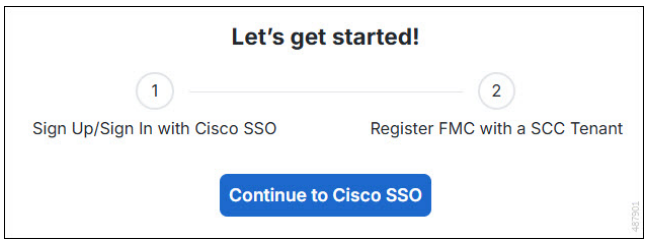
This cloud is also used for Cisco Success Network, Cisco Support Diagnostics, and the Secure Network Analytics cloud using Security Analytics and Logging (SaaS). If you are registered to the Smart Software Manager, your region is preselected.

Step 3 Click Enable **SecureX**.

A separate browser tab opens to log you in to your Security Cloud Control account. Make sure this page is not blocked by a pop-up blocker.

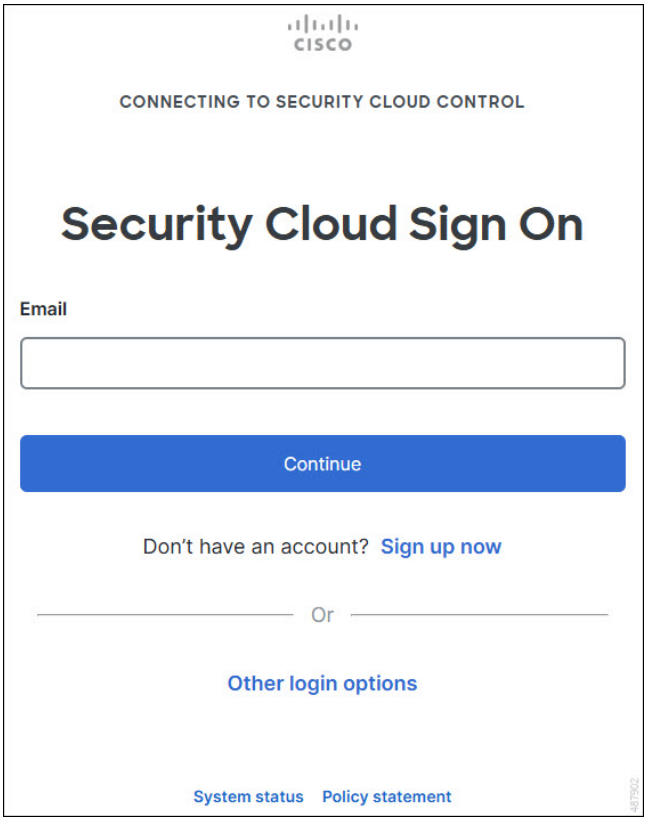
Step 4 Click **Continue to Cisco SSO**.

Figure 2: Welcome Page



Step 5 Log in to your Security Cloud Control account.

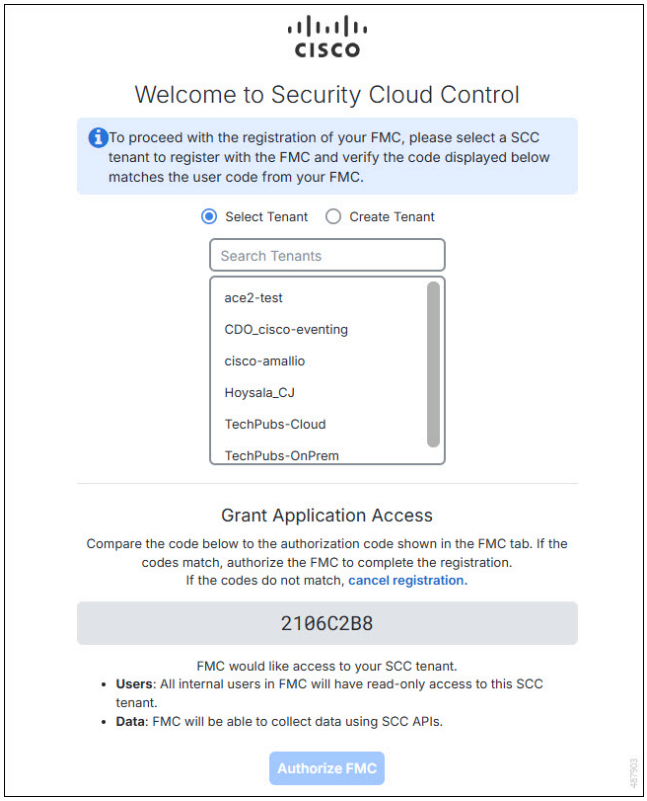
Figure 3: Security Cloud Control Sign On



If you do not have a Security Cloud Sign On account to log in to Security Cloud Control and you want to create one, click **Sign up now** in the **Security Cloud Sign On** page. Refer to Create a New Cisco Security Cloud Sign On Account.

Step 6 Choose a Security Cloud Control tenant that you want to use for this integration. The Secure Firewall Management Center and the managed devices get onboarded to the Security Cloud Control tenant that you choose here.

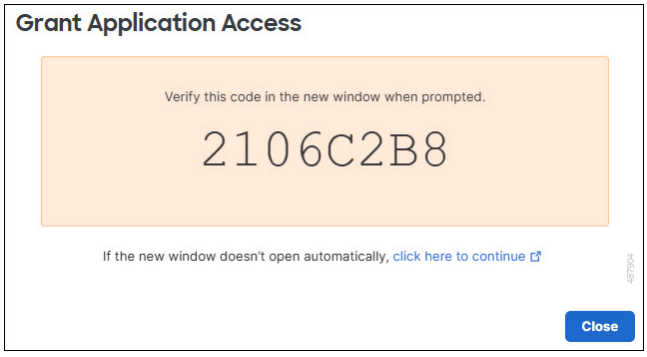
Figure 4: Choose the Security Cloud Control Tenant



If you do not already have a Security Cloud Control tenant or if you want to use a new tenant for this integration, create a new tenant. Refer to Request a Security Cloud Control Tenant for more information.

Step 7 Verify that the code displayed in the Security Cloud Control login page matches the code provided by the Secure Firewall Management Center.

Figure 5: Verification Code in Firewall Management Center



Step 8 Click **Authorize FMC**.

Step 9 In the Secure Firewall Management Center, configure these settings:

- **Event Configuration**: Enable this setting to allow Firewall Threat Defense devices to send events directly to the cloud. The event types that you configure on this page can be used for multiple integrations when

applicable and can be enabled as needed. For more information, refer to Enable sending events to the Cisco Security Cloud, on page 636.

- **Cisco Security Cloud Support**: Enable the Cisco Success Network and Cisco Support Diagnostics capabilities to participate in the customer success initiatives and for an enhanced support experience. For more information, refer to Configure Firewall Management Center to share usage metrics and statistics with Cisco, on page 48 and Configure Firewall Management Center to share device health data with Cisco, on page 49.

After you complete this task, your management center and its managed devices are onboarded to the CDO tenant and integrated with Cisco Security Cloud. You can now manage devices, forward events, and use advanced Cisco security features from the cloud.

Configure Firewall Management Center to share usage metrics and statistics with Cisco

Configure Cisco Success Network to allow your Firewall Management Center to establish a secure connection to Cisco cloud and stream usage information and statistics. Cisco can then inform you about available features that you have not used, provide additional support services, and improve its products.

Cisco Success Network is a cloud service that enables the Firewall Management Center to establish a secure connection to Cisco cloud and stream usage information and statistics. Streaming telemetry enables you to select data of interest from the Firewall Threat Defense device. The device sends this data in a structured format to remote management stations. This transfer supports several purposes:

- To inform you of available, but unused features that can improve the effectiveness of the product in your network.
- To inform you of additional technical support services and monitoring that are available for your product.
- To help Cisco improve its products.

To know more about the telemetry data that Cisco collects, refer to Cisco Success Network Telemetry Data Collected from Cisco Secure Firewall Management Center Devices.

The Firewall Management Center establishes and maintains a secure connection with Cisco cloud at all times when either Cisco Support Diagnostics or Cisco Success Network is enabled. However, the Firewall Management Center and the Firewall Threat Defense devices establish and maintain secure connections with the Cisco cloud when Cisco Support Diagnostics is enabled. You can turn off this connection at any time by disabling both Cisco Success Network and Cisco Support Diagnostics, which disconnects the Firewall Management Center from the Cisco cloud.

You can enable Cisco Success Network when you register the Firewall Management Center with the Smart Software Manager.

Note

- Cisco Success Network is not supported in evaluation mode.
-
- Cisco Success Network is not supported if the Firewall Management Center uses a valid Smart Software Manager On-Prem (formerly known as Smart Software Satellite Server) configuration or the Specific License Reservation.

Before you begin

Enable SecureX integration or register your Firewall Management Center with the Smart License to perform this task.

Follow these steps to configure your Firewall Management Center to share usage metrics and statistics with Cisco.

Procedure

Step 1 Click **Integration** > **SecureX**.

Step 2 Under **Cisco Cloud Support**, check the **Enable Cisco Success Network** check box to enable this service.

Note
Read the information provided next to the **Enable Cisco Success Network** check box before you proceed.

Step 3 Click **Save**.

After completing these steps, your Firewall Management Center will securely share usage metrics and statistics with Cisco, enabling enhanced support and product improvement.

Configure Firewall Management Center to share device health data with Cisco

Configure Firewall Management Center to enable sharing of configuration and operational health data with Cisco. This capability allows automated problem detection and supports Cisco TAC assistance. You can automate sending device health data to Cisco for analysis and proactive issue notification.

Cisco Support Diagnostics sends configuration and operational health data to Cisco, and processes that data through our automated problem detection system. This feature also allows Cisco TAC to collect essential information from your devices during the course of a case. For users with specific service contracts, we can proactively notify you of issues.

Both the Firewall Management Center and its devices communicate with Cisco. For more information, refer to Communication Ports and Internet Access, on page 1063.

Before you begin

- Enable SecureX integration or register with the Smart Software Manager.
- Verify that your system meets the integration or registration requirements before proceeding.

Follow these steps to configure Firewall Management Center to share device health data with Cisco.

Procedure

Step 1 Choose **Integration** > **SecureX**.

Step 2 Under **Cisco Cloud Support**, check **Enable Cisco Support Diagnostics**.

Read the information provided next to the **Enable Cisco Support Diagnostics** check box before you proceed.

Step 3 Click **Save**.

After you complete this task, Firewall Management Center will send configuration and operational health data to Cisco, enabling automated diagnostics and proactive support notifications.

Meet requirements and prerequisites for the system configuration

This reference lists the requirements and prerequisites for the system configuration, including supported models, domains, and user roles.

Model support

Firewall Management Center

Supported domains

Global

User roles

Admin

Manage the Secure Firewall Management Center system configuration

Configure the system to establish and maintain the essential settings for the Secure Firewall Management Center. Ensure that these settings align with organizational requirements.

The system configuration identifies basic settings for the Firewall Management Center. Use this task when you need to review or update the fundamental configuration options for the management center.

Before you begin

Access to the Secure Firewall Management Center interface may be required. Follow these steps to manage the Secure Firewall Management Center system configuration.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Use the navigation panel to choose configurations to change.

When you complete the steps, the basic settings for the Secure Firewall Management Center will update to maintain proper system operation.

What to do next

Verify that the configuration changes are applied and the system is functioning as expected.

Access lists

An access list is a security configuration that

- limits access to the Firewall Management Center by IP address and port
- enables default ports for any IP address, and
- allows optional access for SNMP polling after SNMP is enabled.

Default and optional port access

You can limit access to the Firewall Management Center by IP address and port. These ports are enabled by default for any IP address:

- Port 443 (HTTPS) for web interface access.
- Port 22 (SSH) for CLI access.

You can also add access to poll for SNMP information over port 161. Because SNMP is disabled by default, enable SNMP before adding SNMP access rules. For more information, refer to Configure SNMP polling, on page 119.



Caution By default, access is not restricted. To operate in a more secure environment, consider adding access for specific IP addresses and then deleting the default **any** option.

Configure an access list

To enable secure management of your device, specify which IP addresses can connect and choose the protocols (SSH, HTTPS, or SNMP) they are permitted to use.

This access list does not control external database access. For more information, refer to Enable external access to the database, on page 73.

If you remove access for the IP address you are using and there is no rule for "`IP=any port=443`", you will lose your connection when you save the configuration.

Before you begin

By default, the access list includes rules for HTTPS and SSH. To add SNMP rules to the access list, you must first enable SNMP. For more information, refer to Configure SNMP polling, on page 119.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 (Optional) Click **SNMP** to configure SNMP if you want to add SNMP rules to the access list. By default, SNMP is disabled; refer to Configure SNMP polling, on page 119.

Step 3 Click **Access List**.

Step 4 Click **Add Rules** to add access for IP addresses.

Step 5 In the **IP Address** field, enter an IP address or address range, or `any`.

Step 6 Choose **SSH**, **HTTPS**, **SNMP**, or a combination of these options to specify which ports you want to enable for these IP addresses.

Step 7 Click **Add**.

Step 8 Click **Save**.

The management center allows connections only from the IP addresses and protocols you specify.

Related Topics

IP address conventions, on page 27

Configure access control preferences

You can track changes to access control rules by allowing or requiring users to comment when they save. This allows you to assess why critical policies in a deployment were modified. By default, this feature is disabled.

Configure object optimization to evaluate and optimize network or host policy objects that are used in rules. The system then creates associated network object groups on the device. For more information, refer to Object-group optimization.

Procedure

Step 1 Choose **System** (⚙) > **Configuration** > **Access Control Preferences**.

Step 2 From the **Comments on rule change** list, choose an option:

- **Disabled**: Disables comments for access control rule changes.
- **Optional**: Allows users to optionally add a comment when they change an access control rule.
- **Required**: Requires users to add a comment when they change an access control rule.

Step 3 From the **Object-group optimization** list, choose an option:

- **Enabled**: Enables object-group optimization. The setting takes effect after deployment.
- **Disabled**: Disables object-group optimization.

Warning
For the first deployment to a threat defense device after you enable object-group optimization, the device can take several minutes to an hour to reevaluate the policy configuration and optimize object groups. CPU utilization on the device can also increase. Schedule the deployment during a low-traffic period or a maintenance window.

Step 4 Click **Save**.

What to do next

Deploy the access control policies for object-group optimization to take effect.

Audit logs

An audit log is a security record that

- records user activity in read-only logs,
- allows review, sorting, filtering, deletion, and reporting of audit information, and
- supports streaming audit log messages to external servers such as syslog and HTTP servers.

Audit log review and streaming options

You can review audit log data in several ways:

- Use the web interface: Audit and Syslog, on page 423.

  Audit logs appear in a standard event view. In this view, you can view, sort, and filter audit log messages based on any item. You can also delete audit information, generate reports, and view detailed reports of user changes.

- Stream audit log messages to the syslog: Stream audit logs to syslog, on page 54.
- Stream audit log messages to an HTTP server: Stream audit logs to an HTTP server, on page 55.

Streaming audit log data to an external server allows you to conserve space on the Firewall Management Center. Note that sending audit information to an external URL may affect system performance.

Optionally, you can secure the channel for audit log streaming, enable TLS and mutual authentication using TLS certificates; refer to Audit log certificates, on page 57.

Streaming to multiple syslog servers

You can stream audit log data to a maximum of five syslog servers. However, if you have enabled TLS for secured audit log streaming, you can stream only to a single syslog server.

Stream audit logs to syslog

Configure audit log streaming to syslog servers to ensure centralized monitoring and compliance tracking for your system. Enable external logging of configuration changes and security events.

When this feature is enabled, audit log records appear in the syslog in the following format:

```
Date Time Host: [Tag] Sender: User_Name@User_IP, Subsystem, Action
```

Where the local date, time, and originating hostname precede the bracketed optional tag, and the sending device name precedes the audit log message.

For example, if you specify a tag of `FMC-AUDIT-LOG` for audit log messages from your management center, a sample audit log message from your Firewall Management Center could appear as follows:

```
Mar 01 14:45:24 localhost: [FMC-AUDIT-LOG] Dev-MC7000: admin@10.1.1.2, Operations >
Monitoring, Page View
```

If you specify a severity and facility, these values do not appear in syslog messages; instead, they tell the system that receives the syslog messages how to categorize them.

This task is relevant when you need to send audit logs from your management center to an external syslog server for compliance or operational monitoring. Use this procedure when integrating with centralized log management solutions.

Before you begin

Make sure the Firewall Management Center can communicate with the syslog server. When you save your configuration, the system uses ICMP/ARP and TCP SYN packets to verify syslog server connectivity. By default, the system uses port 514/UDP to stream audit logs. To secure the channel (optional, refer to Audit log certificates, on page 57), manually configure port 1470 for TCP.

Follow these steps to stream audit logs to a syslog server:

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **Audit Log**.

Step 3 Choose **Enabled** from the **Send Audit Log to Syslog** drop-down menu.

Step 4 The following fields are applicable only for audit logs sent to syslog:

Option	Description
Host	The IP address or the fully qualified name of the syslog server to which you will send audit logs. You can add a maximum of five syslog hosts, separated by commas. **Note** You can specify multiple syslog hosts, only when TLS is disabled for the Audit Server Certificate.
Facility	The subsystem that creates the message. Choose a facility described in SYSLOG alert facilities, on page 565. For example, choose AUDIT.

Option	Description
Severity	The severity of the message. Choose a severity described in Syslog severity levels, on page 566.
Tag	An optional tag to include in audit log syslog messages. Best practice: Enter a value in this field to easily differentiate audit log messages from other, similar syslog messages such as health alerts. For example, if you want all audit log records sent to the syslog to be labeled with `FMC-AUDIT-LOG`, enter `FMC-AUDIT-LOG` in the field.

Step 5 (Optional) To test whether the IP address of the syslog servers is valid, click **Test Syslog Server**.

The system sends the following packets to verify whether the syslog server is reachable:

a. ICMP echo request

b. TCP SYN on 443 and 80 ports

c. ICMP time stamp query

d. TCP SYN on random ports

Note
If the Firewall Management Center and syslog server are in the same subnet, ARP is used instead of ICMP.

The system displays the result for each server.

Step 6 Click **Save**.

After you complete this task, audit logs are sent from the management center to the specified syslog server, allowing you to monitor and review system activity externally.

What to do next

Verify that the syslog server is receiving audit log messages as expected. Check for log entries corresponding to configuration changes or security events.

Stream audit logs to an HTTP server

Stream audit logs from the device to an HTTP server to centralize log management and support compliance requirements.

Use this task to forward audit logs from your device to an external HTTP server for monitoring, analysis, or compliance.

When this feature is enabled, the appliance sends audit log records to an HTTP server in the following format:

```
Date Time Host [Tag] Sender: User_Name@User_IP, Subsystem, Action
```

Where the local date, time, and originating hostname precede the bracketed optional tag, and the sending appliance name precedes the audit log message.

For example, if you specify a tag of `FROMMC`, a sample audit log message could appear as follows:

```
Mar 01 14:45:24 localhost [FROMMC] Dev-MC7000: admin@10.1.1.2, Operations > Monitoring,
Page View
```

Before you begin

Make sure the device can communicate with the HTTP server. Optionally, secure the channel. Refer to Audit log certificates, on page 57.

Follow these steps to stream audit logs to an HTTP server.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **Audit Log**.

Step 3 Optionally, in the **Tag** field, enter the tag name that you want to appear with the message. For example, if you want all audit log records to be preceded with `FROMMC`, enter `FROMMC` in the field.

Step 4 Choose **Enabled** from the **Send Audit Log to HTTP Server** drop-down list.

Step 5 In the **URL to Post Audit** field, designate the URL where you want to send the audit information. Enter a URL that corresponds to a Listener program that expects the HTTP POST variables as listed:

- `subsystem`
- `actor`
- `event_type`
- `message`
- `action_source_ip`
- `action_destination_ip`
- `result`
- `time`
- `tag` (if defined; see Step 3)

Caution
To allow encrypted posts, use an HTTPS URL. Sending audit information to an external URL may affect system performance.

Step 6 Click **Save**.

When you complete this task, the device sends audit logs to the specified HTTP server. Centralized monitoring and record-keeping become possible.

Filter syslogs from audit logs

When you enable **Send Audit Log to Syslog** and provide **Host** information, syslog messages are also sent to the configured host in addition to the audit logs. This behavior occurs because the system creates

`/etc/syslog-ng.d/syslog-tls.conf` when you deploy the Firepower platform settings policy. As a result, syslog messages are forwarded to the configured host rather than only sending the audit logs.

If your auditing policy does not need these syslog records, you can prevent syslogs from being streamed to the configured host. To filter syslogs from audit logs, ensure you have access to the appliance's **admin** user account. You must also be able to access the appliance's console or open a secure terminal.

Caution Make sure that only authorized personnel have access to the appliance and to its **admin** account.

Procedure

Step 1 In the `/etc/syslog-ng.conf` file, comment out the `@include "/etc/syslog-ng.d/*.conf"` line.

Example:

```
#@include "/etc/syslog-ng.d/*.conf"
```

Step 2 Reload the syslog configuration file. Use the `syslog-ng-ctl reload` command to reload the configuration file without having to restart the application.

Example:

```
syslog-ng-ctl reload
```

Audit log certificates

An audit log certificate is a security credential that

- secures communications between the Firewall Management Center and a trusted audit log server,
- enables authentication and encryption using Transport Layer Security (TLS), and
- supports mutual authentication and certificate revocation checks for enhanced security.

Client certificates

Client certificates are required to secure communications between the Firewall Management Center and the audit log server.

- Generate a certificate signing request (CSR). Submit it to a Certificate Authority (CA) for signing. Import the signed certificate onto the Firewall Management Center.
- Use the local system configuration: Obtain a signed audit log client certificate for the Firewall Management Center, on page 59 and Import an audit log client certificate into the Firewall Management Center, on page 60.

Server certificates

Server certificates are optional and provide additional security through mutual authentication.

- Require mutual authentication between the Firewall Management Center and the audit log server by loading one or more certificate revocation lists (CRLs).
- You cannot stream audit logs to servers with revoked certificates listed in those CRLs.
- Secure Firewall supports CRLs encoded in Distinguished Encoding Rules (DER) format. These are the same CRLs used to validate HTTPS client certificates for the Firewall Management Center web interface.
- Use the local system configuration: Require valid audit log server certificates, on page 61.

Audit log certificate example

For example, a signed client certificate imported onto the Firewall Management Center enables secure TLS communication with an audit log server, ensuring log integrity and confidentiality.

Non-compliant audit log certificate

A certificate that is not signed by a trusted Certificate Authority or is listed in a certificate revocation list (CRL) cannot be used to stream audit logs securely.

Audit log certificate analogy

An audit log certificate is like a passport for secure communication, verifying identity and granting access between the management center and the audit log server.

Secure audit log streaming

Configure secure audit log streaming to ensure that audit logs are transmitted safely between the Firewall Management Center and a trusted server.

When you stream the audit log to a trusted HTTP server or syslog server, use Transport Layer Security (TLS) certificates to secure the channel between the Firewall Management Center and the server.

You must generate a unique client certificate for each appliance you want to audit. Mutual authentication requires the client certificate to be signed by the same CA as the server certificate. For more considerations on client and server certificate requirements, refer to Audit log certificates, on page 57.

Before you begin

Refer to ramifications of requiring client and server certificates at Audit log certificates, on page 57. Ensure you have access to a recognized certificate authority (CA) for signing client certificates.

Follow these steps to secure audit log streaming.

Procedure

Step 1 Obtain and install a signed client certificate on the Firewall Management Center.

a) Obtain a signed audit log client certificate for the Firewall Management Center, on page 59:

Generate a Certificate Signing Request (CSR) from the Firewall Management Center based on your system information and the identification information you supply.

- Submit the CSR to a recognized, trusted certificate authority (CA) to request a signed client certificate.
- If you will require mutual authentication between the Firewall Management Center and the audit log server, ensure the client certificate is signed by the same CA that signed the server certificate.

b) After you receive the signed certificate from the certificate authority, import it into the Firewall Management Center. Refer to Import an audit log client certificate into the Firewall Management Center, on page 60.

Step 2 Configure the communication channel with the server to use Transport Layer Security (TLS) and enable mutual authentication.

Refer to Require valid audit log server certificates, on page 61.

Step 3 If you have not yet configured audit log streaming, do so now.

Refer to Stream audit logs to syslog, on page 54 or Stream audit logs to an HTTP server, on page 55.

Audit logs are securely streamed to the trusted server using TLS, with mutual authentication enabled. The Firewall Management Center and the server verify each other's certificates, ensuring secure transmission and integrity of audit log data.

What to do next

Monitor the audit log streaming status and verify certificate validity periodically. Renew client and server certificates before expiration to maintain secure log streaming.

Obtain a signed audit log client certificate for the Firewall Management Center

Obtain a signed audit log client certificate for the Firewall Management Center. This certificate enables secure communication between the appliance and the audit log server. It ensures that the server can authenticate audit log data and that the data is transmitted securely.

The system generates certificate request keys in Base-64 encoded PEM format.



Important The **Audit Log Certificate** page is not available on a standby Firewall Management Center in a high availability setup. You cannot perform this task from a standby Firewall Management Center.

Perform this task on the active appliance to generate a certificate signing request for audit log client authentication. Do not perform this task on a standby appliance in a high availability setup.

Before you begin

Keep this in mind:

- To ensure security, use a globally recognized and trusted certificate authority (CA) to sign your certificate.
- If you will require mutual authentication between the appliance and the audit log server, the same certificate authority must sign both the client certificate and the server certificate.

Follow these steps to obtain a signed audit log client certificate for the Firewall Management Center:

Procedure

Step 1 Choose **System** (⚙) > **Configuration** and click **Audit Log Certificate**.

Step 2 Click **Generate New CSR**.

Step 3 Enter a country code in the **Country Name (two-letter code)** field.

Step 4 Enter a state or province postal abbreviation in the **State or Province** field.

Step 5 Enter a **Locality or City**, an **Organization** name, and an **Organizational Unit (Department)** name.

Step 6 Enter the fully qualified domain name of the server for which you want to request a certificate in the **Common Name** field.

Note
If the common name and the DNS hostname do not match, audit log streaming will fail.

Step 7 Click **Generate** and then open a new blank file with a text editor.

Step 8 Copy the block of text from the certificate request, including the `BEGIN CERTIFICATE REQUEST` and `END CERTIFICATE REQUEST` lines. Paste this information into a blank text file.

Step 9 Save the file as *`clientname`*`.CSR`. Replace *`clientname`* with the appliance name where you plan to use the certificate. Click **Close**.

After completing these steps, you will have a certificate signing request (CSR) file that can be submitted to a trusted certificate authority. After it is signed, you can import the certificate to the appliance for secure audit log streaming.

What to do next

- Submit the certificate signing request to the certificate authority you selected, according to the guidelines in the 'Before You Begin' section.
- When you receive the signed certificate, import it to the appliance; refer to Import an audit log client certificate into the Firewall Management Center, on page 60.

Import an audit log client certificate into the Firewall Management Center

Importing an audit log client certificate enables secure communication between the Secure Firewall Management Center and external systems during audit log transmission. This task verifies that you import the correct signed certificate and any necessary intermediate certificates for secure operations.

In a Secure Firewall Management Center high availability setup, you must use the active peer when importing the audit log client certificate. This procedure is relevant when configuring audit log security or after obtaining a new signed client certificate.

Before you begin

- Obtain a signed audit log client certificate for the Firewall Management Center, on page 59.
- Make sure you are importing the signed certificate for the correct Firewall Management Center.

- If the signing authority requires an intermediate CA, ensure you provide the certificate chain. The same CA must sign both the client certificate and the intermediate certificates.

Follow these steps to import an audit log client certificate into the Secure Firewall Management Center.

Procedure

Step 1 On the Firewall Management Center, choose **System** (⚙) > **Configuration**.

Step 2 Click **Audit Log Certificate**.

Step 3 Click **Import Audit Client Certificate**.

Step 4 Open the client certificate in a text editor, copy the entire block of text, including the `BEGIN CERTIFICATE` and `END CERTIFICATE` lines. Paste this text into the **Client Certificate** field.

Step 5 To upload a private key, open the private key file and copy the entire block of text, including the `BEGIN RSA PRIVATE KEY` and `END RSA PRIVATE KEY` lines. Paste this text into the **PRIVATE KEY** field.

Step 6 Open any required intermediate certificates, copy the entire block of text for each, and paste it into the **Certificate Chain** field.

Step 7 Click **Save**.

The audit log client certificate and any required intermediate certificates are successfully imported into the Secure Firewall Management Center.

Require valid audit log server certificates

The system supports validating audit log server certificates using imported CRLs in Distinguished Encoding Rules (DER) format.

Note If you choose to verify certificates using CRLs, the system uses the same CRLs to validate both audit log server certificates and certificates used to secure the HTTP connection between an appliance and a web browser.

Important You cannot perform this procedure on the standby Firewall Management Center in a high availability pair.

Before you begin

- Understand the ramifications of requiring mutual authentication and of using certificate revocation lists (CRLs) to ensure that certificates are still valid. Refer to Audit log certificates, on page 57.
- Obtain and import the client certificate following the steps in Secure audit log streaming, on page 58 and the topics referenced in that procedure.

Procedure

Step 1 On the Firewall Management Center , choose **System** (⚙) > **Configuration** .

Step 2 Click **Audit Log Certificate** .

Step 3 To use Transport Layer Security to securely stream the audit log to an external server, select **Enable TLS** .

When TLS is enabled, the syslog client (Firewall Management Center) verifies the certificate received from the server. The connection between the client and the server succeeds only if server certificate verification is successful. For this verification process, these conditions must be met:

- Configure the syslog server to send the certificate to the client.
- Add (import) a CA certificate to the client to verify the server certificate:
  - You must import the CA certificate during the import of the client certificate.
  - If the issuing CA is a subordinate CA, add the issuing CA before adding the signing CA from the subordinate CA (Root CA).

Step 4 If you do not want the client to authenticate itself against the server, but wish to accept the server certificate when the certificate is issued by the same CA (not recommended), complete these steps.

a) Deselect **Enable Mutual Authentication** .

Important
Ensure that the server is configured to trust the client without verifying any client certificates.

b) Click **Save** and skip the remainder of this procedure.

Step 5 (Optional) To enable client certificate verification by the audit log server, select **Enable Mutual Authentication**.

Important
The **Enable Mutual Authentication** option is applicable only when TLS is enabled.

When mutual authentication is enabled, the syslog client (Firewall Management Center) sends a client certificate to the syslog server for verification. The client uses the same CA certificate of the CA who signed the server certificate of the syslog server. The connection succeeds only if client certificate verification is successful. For this verification process, these conditions must be met:

- Configure the syslog server to verify the certificate received from the client.
- Add a client certificate to be sent to the syslog server. This certificate must be signed by the same CA who signed the server certificate of the syslog server.

Note
To use mutual authentication for streaming Audit Log to the Syslog server, use PKCS#8 format for the private key instead of PKCS#1 format. Use this command line to convert PKCS#1 keys to PKCS#8 format:

```
openssl pkcs8 -topk8 -inform PEM -outform PEM
-nocrypt -in

 PKCS1 key file name
```

```
-out

 PKCS8 key filename
```

Step 6 (Optional) To automatically recognize server certificates that are no longer valid:

a) Select **Enable Fetching of CRL** .

Important
This option is displayed only when you select the **Enable Mutual Authentication** check box. However, the **Enable Fetching of CRL** option is applicable only when the TLS option is enabled. The use of CRL is for server certification verification, and it is not dependent on the use of Mutual Authentication which is for enabling client certificate verification.

When you enable fetching of the CRL, the client creates a scheduled task to regularly update (download) the CRL or CRLs. The CRLs are used for server certificate verification. Verification fails if a CRL from the CA specifies that the server certificate has been revoked.

b) Enter a valid URL to an existing CRL file and click **Add CRL**.

Repeat to add up to 25 CRLs.

c) Click **Refresh CRL** to load the current CRL or CRLs from the specified URL or URLs.

Step 7 Verify that you have a valid server certificate generated by the same certificate authority that created the client certificate.

Step 8 Click **Save**.

What to do next

(Optional) Set the frequency of CRL updates.

View the audit log client certificate on the Firewall Management Center

View the audit log client certificate to confirm its presence and details on the appliance you are logged in to. This helps ensure that audit log operations are secured with the correct certificate.

You can view the audit log client certificate only for the appliance that you are logged in to.

In Firewall Management Center high availability pairs, you can view the certificate only on the active peer.

Before you begin

Follow these steps to view the audit log client certificate on the appliance:

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **Audit Log Certificate**.

The appliance you are logged in to displays the audit log client certificate details. For high availability pairs, only the active peer shows the certificate.

Change reconciliation

A change reconciliation report is a configuration monitoring tool that

- captures snapshots whenever a user saves changes to the system configuration,
- combines information from these snapshots to present a clear summary of recent system changes, and
- provides a detailed report of changes made over the past 24 hours through email.

User change summaries

The change reconciliation report displays both the previous value for each configuration and the value after changes.

Users can view summaries of each distinct change in chronological order, beginning with the most recent.

- Lists both previous and updated values for each configuration.
- Summarizes multiple changes to the same configuration.
- Orders changes chronologically, starting with the most recent.

Change reconciliation report example

An example change reconciliation report includes a User section that lists configuration changes, showing both the previous and updated values for each item.

Non-reconciled changes

Changes made without reconciliation are not summarized or tracked in a report, making it difficult to monitor user activity.

Change reconciliation as a transaction log

Change reconciliation is similar to a transaction log, where each modification is recorded and summarized for review.

Configure change reconciliation

Configure change reconciliation to enable the system to generate and send daily reports of configuration changes, supporting audit and compliance requirements. This helps administrators monitor and review all changes made to the system within a 24-hour period.

Use this task when you need to track and review configuration changes made to your system for security, compliance, or operational awareness.

Change reconciliation is especially useful in environments where multiple administrators make changes or where audit trails are required for regulatory purposes. Perform this configuration after setting up your email server to ensure reports are delivered successfully.

Before you begin

Configure an email server to receive emailed reports of changes made to the system over a 24-hour period; refer to Configure a mail relay host and notification address, on page 71 for more information.

Follow these steps to configure change reconciliation.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **Change Reconciliation**.

Step 3 Check the **Enable** check box.

Step 4 Choose the time of day you want the system to send out the change reconciliation report from the **Time to Run** drop-down lists.

Step 5 Enter email addresses in the **Email to** field.

Tip
Once you have added email addresses, click **Resend Last Report** to send recipients another copy of the most recent change reconciliation report.

Step 6 If you want to include policy changes, check the **Include Policy Configuration** check box.

Step 7 If you want to include all changes over the past 24 hours, check the **Show Full Change History** check box.

Step 8 Click **Save**.

After you complete this task, the system will automatically generate and email daily reports of all configuration changes, allowing you to monitor and review system modifications for compliance and troubleshooting.

Related Topics

Use audit log to examine changes, on page 428

Change reconciliation options

The change reconciliation options control whether the system includes records of policy changes and the full change history in the change reconciliation report. These options determine the scope and detail of reported changes.

The **Include Policy Configuration** option controls whether the system includes records of policy changes in the change reconciliation report. It includes changes to access control, intrusion, system, health, and network discovery policies. If you do not select this option, the report will not show changes to any policies. This option is available only on Firewall Management Centers.

The **Show Full Change History** option controls whether the system includes records of all changes over the past 24 hours in the change reconciliation report. If you do not select this option, the report includes only a consolidated view of changes for each category.

Note The change reconciliation report does not include changes to Firewall Threat Defense interfaces and routing settings.

DNS caches

A DNS cache is a network appliance feature that

- enables automatic IP address resolution on event view pages,
- stores previously resolved IP addresses to avoid additional lookups, and
- reduces network traffic and improves page display speed when IP address resolution is enabled.

Configure DNS cache properties

Configure DNS cache properties to manage DNS resolution caching on your system. You can enable or disable DNS caching and set the cache timeout to optimize DNS lookup performance.

DNS resolution caching is a system-wide setting that allows the caching of previously resolved DNS lookups.

Before you begin

Follow these steps to configure DNS cache properties.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Choose **DNS Cache**.

Step 3 From the **DNS Resolution Caching** drop-down list, choose one of these options:

- **Enabled**—Enable caching.
- **Disabled**—Disable caching.

Step 4 In the **DNS Cache Timeout (in minutes)** field, enter the number of minutes a DNS entry remains cached in memory before it is removed for inactivity.

The default setting is 300 minutes (five hours).

Step 5 Click **Save**.

After completing this task, DNS cache properties are updated according to your configuration. DNS lookups will be cached based on your configuration. Entries remain in cache for the specified timeout period.

Related Topics

Configure event view settings, on page 234

Dashboards

A dashboard is a system interface that

- provides at-a-glance views of current system status through widgets,
- uses small, self-contained components to give insight into different aspects of the system, and
- includes several predefined dashboard widgets.

Dashboard configuration options

You can configure the Firewall Management Center so that Custom Analysis widgets are enabled on the dashboard.

Related Topics

Dashboards, on page 357

Enable custom analysis widgets for dashboards

Enable custom analysis widgets to provide users with the ability to create visual representations of events based on flexible, user-configurable queries.

Custom Analysis dashboard widgets are used to visually represent events according to user-defined queries, offering flexibility in monitoring and analyzing dashboard data.

Before you begin

Follow these steps to enable custom analysis widgets for dashboards.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **Dashboard**.

Step 3 Check the **Enable Custom Analysis Widgets** check box to allow users to add Custom Analysis widgets to dashboards.

Step 4 Click **Save**.

Related Topics

Dashboards, on page 357

Database

To manage disk space, the Firewall Management Center periodically prunes the oldest intrusion events, audit records, Security Intelligence data, and URL filtering data from the event database. For each event type, you can specify how many records the Firewall Management Center retains after pruning; never rely on the event database containing more records of any type than the retention limit configured for that type. To improve

performance, tailor the event limits to the number of events you regularly work with. You can optionally choose to receive email notifications when pruning occurs. For some event types, you can disable storage.

To manually delete individual events, use the event viewer. (Note that in Versions 6.6.0+, you cannot manually delete connection or security Intelligence events in this way.) You can also manually purge the database; refer to Data Purge and Storage, on page 521.

Configure database event limits

Configure database event limits to manage the number of records retained in each database and to receive notifications when data pruning occurs. This helps maintain optimal database performance and ensures you are informed about important data retention events.

This task is relevant when you need to control database size and receive alerts about event pruning in the Secure Firewall Management Center.

Use this procedure as part of regular database maintenance or when adjusting event retention policies. This configuration ensures that your system retains only the necessary amount of event data and notifies you when records are pruned.

Before you begin

If you want to receive email notifications when events are pruned from the Firewall Management Center's database, you must configure an email server; refer to Configure a mail relay host and notification address, on page 71.

Follow these steps to configure database event limits and set up pruning notifications:

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Choose **Database**.

Step 3 Enter the number of records to store for each database.

For information on how many records each database can maintain, refer to Database event limits, on page 68.

Step 4 Optionally, in the **Data Pruning Notification Address** field, enter the email address where you want to receive pruning notifications.

Step 5 Click **Save**.

After completing this task, the database will store only the specified number of records for each database, and you will receive email notifications if data pruning occurs (if configured).

Database event limits

This table lists the minimum and maximum number of records for each event type that you can store per Firewall Management Center .

Table 4: Database Event Limits

Event Type	Upper Limit	Lower Limit
Intrusion events	10 million (Firewall Management Center Virtual) 30 million (Firewall Management Center 1000, Firewall Management Center 1600,) 60 million (Firewall Management Center 2500, Firewall Management Center 2600, , FMCv 300) 300 million (Firewall Management Center 4500, Firewall Management Center 4600)	10,000
Discovery events	10 million (Firewall Management Center Virtual) 20 million (Firewall Management Center 2500, Firewall Management Center 2600, Firewall Management Center 4500, Firewall Management Center 4600, FMCv 300)	Zero (disables storage)
Connection events Security Intelligence events	50 million (Firewall Management Center Virtual) 100 million (Firewall Management Center 1000, Firewall Management Center 1600,) 300 million (Firewall Management Center 2500, Firewall Management Center 2600, , FMCv 300) 1 billion (Firewall Management Center 4500, Firewall Management Center 4600) Limit is shared between connection events and security intelligence events . The sum of the configured maximums cannot exceed this limit.	Zero (disables storage) If you set the **Maximum Connection Events** value to zero, then connection events that are not associated with security intelligence , intrusion, file, and malware events are not stored on the Firewall Management Center . **Caution** Setting **Maximum Connection Events** to zero immediately purges existing connection events other than security intelligence events . This setting affects the Maximum Flow Rate, as described in the next section. These settings do not affect connection summaries.

Event Type	Upper Limit	Lower Limit
Connection summaries (aggregated connection events)	50 million (Firewall Management Center Virtual) 100 million (Firewall Management Center 1000, Firewall Management Center 1600,) 300 million (Firewall Management Center 2500, Firewall Management Center 2600, , FMCv 300) 1 billion (Firewall Management Center 4500, Firewall Management Center 4600)	Zero (disables storage)
Correlation events and compliance allow list events	1 million (Firewall Management Center Virtual) 2 million (Firewall Management Center 2500, Firewall Management Center 2600, Firewall Management Center 4500, Firewall Management Center 4600, FMCv 300)	One
Malware events	10 million (Firewall Management Center Virtual, Firewall Management Center 1600,) 20 million (Firewall Management Center 2500, Firewall Management Center 2600, Firewall Management Center 4500, Firewall Management Center 4600, FMCv 300)	10,000
File events	10 million (Firewall Management Center Virtual, Firewall Management Center 1600) 20 million (Firewall Management Center 2500, Firewall Management Center 2600, , Firewall Management Center 4500, Firewall Management Center 4600, FMCv 300)	Zero (disables storage)
Health events	1 million	Zero (disables storage)
Audit records	100,000	One
Remediation status events	10 million	One
Allow list violation history	a 30-day history of violations	One day’s history
User activity (user events)	10 million	One
User logins (user history)	10 million	One

Event Type	Upper Limit	Lower Limit
Intrusion rule update import log records	1 million	One
VPN Troubleshooting database	10 million	Zero (disables storage)

Maximum Flow Rate

The **Maximum flow rate** (flows per second) value for your Firewall Management Center hardware model is specified in the **Platform Specifications** section of the Firewall Management Center datasheet at https://www.cisco.com/c/en/us/products/collateral/security/firesight-management-center/datasheet-c78-736775.html?cachemode=refresh

If you set the **Maximum Connection Events** value in platform settings to zero, then connection events that are not associated with security intelligence events , intrusion, file, and malware events are not counted toward the maximum flow rate for your Firewall Management Center hardware.

Any non-zero value in this field causes ALL connection events to be counted against the maximum flow rate.

Other event types on this page do not count against the maximum flow rate.

Email notifications

An email notification is a system alert mechanism that enables automated emailing of event-based, status, change reconciliation, and data-pruning reports.

Email notification configuration options

Configure a mail host if you plan to:

- Email event-based reports
- Email status reports for scheduled tasks
- Email change reconciliation reports
- Email data-pruning notifications
- Use email for discovery event, impact flag, correlation event alerting, intrusion event alerting, and health event alerting.

When you configure email notification, you can select an encryption method for the communication between the system and mail relay host, and can supply authentication credentials for the mail server if needed. After you configure email notification, you can test the connection.

Configure a mail relay host and notification address

Email notifications help administrators track important events and maintain appliance security. Configuring your appliance to use a relay mail server ensures that messages are sent securely.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **Email Notification**.

Step 3 In the **Mail Relay Host** field, enter the hostname or IP address of the mail server you want to use. Ensure the mail host you enter allows access from the appliance.

Step 4 In the **Port Number** field, enter the port number to use on the email server.

Typical ports include:

- 25, when using no encryption
- 465, when using SSLv3
- 587, when using TLS

Step 5 Choose an **Encryption Method**:

- **TLS**: Encrypt communications using Transport Layer Security.
- **SSLv3**: Encrypt communications using Secure Socket Layers.
- **None**: Allow unencrypted communication.

Note
Certificate validation is not required for encrypted communication between the appliance and mail server.

Step 6 In the **From Address** field, enter the valid email address you want to use as the source email address for messages sent by the appliance.

Step 7 To supply a username and password when connecting to the mail server, choose **Use Authentication**. Enter the username in the **Username** field. Enter a password in the **Password** field.

Step 8 To send a test email using the configured mail server, click **Test Mail Server Settings**.

A message appears next to the button indicating the success or failure of the test.

Step 9 Click **Save**.

External database access

An external database access is a database connectivity feature that

- allows read-only access to the database by third-party clients,
- enables querying the database using SQL with industry-standard reporting tools and applications, and
- supports secure connections through JDBC SSL and provides Cisco-provided tools for database access.

Database access tools and configuration

You can configure the Firewall Management Center to allow read-only access to its database by a third-party client. This allows you to query the database using SQL with these tools and applications:

- industry-standard reporting tools such as Actuate BIRT, JasperSoft iReport, or Crystal Reports,
- any other reporting application (including a custom application) that supports JDBC SSL connections,
- the Cisco-provided command-line Java application called RunQuery, which you can either run interactively or use to obtain comma-separated results for a single query.

Use the Firewall Management Center's system configuration to enable database access and create an access list so selected hosts can query the database. This access list does not control appliance access.

You can also download a package that contains these tools:

- RunQuery, the Cisco-provided database query tool,
- InstallCert, a tool that you can use to retrieve and accept the SSL certificate from the Firewall Management Center that you want to access,
- the JDBC driver you must use to connect to the database.

Refer to the *Secure Firewall Management Center Database Access Guide* for information on using the tools in the package you downloaded to configure database access.

Enable external access to the database

Enable external access to the database so that third-party applications and remote hosts can connect as required.

Perform this configuration when integrating with third-party tools or enabling remote database access for operational needs. Ensure you understand your organization's security requirements before enabling external access.

Before you begin

Verify that you have administrative access and know the required hostnames or IP addresses for external access.

Follow these steps to enable external access to the database.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **External Database Access**.

Step 3 Select the **Allow External Database Access** check box.

Step 4 Enter an appropriate value in the **Server Hostname** field. Depending on your third-party application requirements, this value can be either the fully qualified domain name (FQDN), IPv4 address, or IPv6 address of the Firewall Management Center.

Note
In Firewall Management Center high availability setups, enter only the active peer details. We do not recommend entering details of the standby peer.

Step 5 Next to **Client JDBC Driver**, click **Download** and follow your browser's prompts to download the `client.zip` package.

Step 6 To add database access for one or more IP addresses, click **Add Hosts**. An **IP Address** field appears in the **Access List** field.

Step 7 In the **IP Address** field, enter an IP address or address range, or `any`.

Step 8 Click **Add**.

Step 9 Click **Save**.

Tip
If you want to revert to the last saved database settings, click **Refresh**.

External access to the database is enabled. Specified hosts can now connect using the configured settings, and the JDBC driver is available for client applications.

What to do next

Review database access logs to monitor external connections and ensure compliance with security policies. Update access lists when you need to add or remove hosts.

Related Topics

IP address conventions, on page 27

HTTPS certificates

An HTTPS certificate is a security credential that

- enables Firewall Management Center to establish an encrypted channel to a web browser
- includes a default certificate with all firewall devices, and
- may require replacement with a custom certificate signed by a globally recognized or internally trusted certificate authority to improve trust.

Certificate support information

The Firewall Management Center supports 4096-bit HTTPS certificates. If the certificate used by the Firewall Management Center was generated using a public server key larger than 4096 bits, you will not be able to log in to the Firewall Management Center web interface. If this happens, contact Cisco TAC.

Note HTTPS certificates are not supported on the Firewall Management Center REST API.

Default HTTPS server certificates

A default HTTPS server certificate is a security credential that

- is provided with an appliance for web interface access,

- is not signed by the CA that signs client certificates, and
- expires 20 years from when it was first generated.

Default HTTPS server certificate reference information

Default HTTPS server certificates have specific characteristics and limitations.

- If you use the default server certificate provided with an appliance, do not configure the system to require a valid HTTPS client certificate for web interface access. The default server certificate is not signed by the CA that signs your client certificate.
- The default server certificate provided with an appliance expires 20 years from when it was first generated.
- The lifetime of the default server certificate depends on when the certificate was generated. To view your default server certificate expiration date, choose **System** (⚙) > **Configuration** > **HTTPS Certificate**.
- Some Secure Firewall software upgrades can automatically renew the certificate. For more information, refer to the appropriate version of the Release Notes for Secure Firewall Threat Defense.
- On the Firewall Management Center, you can renew the default certificate on the **System** (⚙) > **Configuration** > **HTTPS Certificate** page.

Custom HTTPS server certificates

A custom HTTPS server certificate is a security credential that

- is generated based on system and identification information you supply,
- can be signed by an internal certificate authority (CA) trusted by your browser, and
- can be imported after being signed by a certificate authority.

Meet HTTPS server certificate requirements

To secure the connection between your web browser and the Secure Firewall appliance web interface using HTTPS, you must use server certificates that comply with the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile (RFC 5280). The system rejects certificates that are not compliant with X.509 version 3 of that standard.

Import an HTTPS server certificate only if it includes all required fields.

Certificate Field	Description
Version	Version of the encoded certificate. Use version 3. Refer to RFC 5280, section 4.1.2.1.
Serial number	A positive integer assigned to the certificate by the issuing CA. Issuer and serial number together uniquely identify the certificate. Refer to RFC 5280, section 4.1.2.2.

Certificate Field	Description
Signature	Identifier for the algorithm used by the CA to sign the certificate. Must match the signatureAlgorithm field. Refer to RFC 5280, section 4.1.2.3.
Issuer	Identifies the entity that signed and issued the certificate. Refer to RFC 5280, section 4.1.2.4.
Validity	Interval during which the CA warrants that it will maintain information about the status of the certificate. Refer to RFC 5280, section 4.1.2.5.
Subject	Identifies the entity associated with the public key stored in the subject public key field; must be an X.500 distinguished name (DN). Refer to RFC 5280, section 4.1.2.6.
Subject Alternative Name	The certificate secures domain names and IP addresses. Subject Alternative Name is defined in RFC 5280, section 4.2.1.6. We recommend using this field if the certificate is used for multiple domains or IP addresses.
Subject Public Key Info	Public key and an identifier for its algorithm. Refer to RFC 5280, section 4.1.2.7.
Authority Key Identifier	Provides a means of identifying the public key corresponding to the private key used to sign a certificate. Refer to RFC 5280, section 4.2.1.1.
Subject Key Identifier	Provides a means of identifying certificates that contain a particular public key. Refer to RFC 5280, section 4.2.1.2.
Key Usage	Defines the purpose of the key contained in the certificates. Refer to RFC 5280, section 4.2.1.3.
Basic Constraints	Identifies whether the certificate Subject is a CA, and the maximum depth of validation certification paths that include this certificate. Refer to RFC 5280, section 4.2.1.9. For server certificates used in Secure Firewall appliances, use `critical CA:FALSE`.
Extended Key Usage extension	Indicates one or more purposes for which the certified public key may be used, in addition to or in place of the basic purposes indicated in the Key Usage extension. Refer to RFC 5280, section 4.2.1.12. Make sure to import certificates that can be used as server certificates.

Certificate Field	Description
signatureAlgorithm	Identifier for the algorithm the CA used to sign the certificate. Must match the Signature field. Refer to RFC 5280, section 4.1.1.2.
signatureValue	Digital signature. Refer to RFC 5280, section 4.1.1.3.

HTTPS client certificates

An HTTPS client certificate is a security credential that

- enables the web server to check that a user's browser client has a valid certificate selected,
- requires the certificate to be generated by the same trusted certificate authority as the server certificate, and
- prevents the browser from loading the web interface if the certificate is invalid or revoked.

Certificate validation methods

HTTPS client certificates are validated using OCSP or certificate revocation lists (CRLs).

- OCSP allows the web server to communicate with the certificate authority to confirm the client certificate's validity before establishing the connection.
- CRLs enable the web server to compare the client certificate against certificates listed in the CRLs. The system blocks access if the certificate is revoked.

Note If you choose to verify certificates using CRLs, the system uses the same CRLs to validate both client browser certificates and audit log server certificates.

Invalid certificate scenarios

The browser cannot load the web interface if the user selects a certificate that is not valid, not generated by the certificate authority that signed the server certificate, or not generated by a certificate authority in the certificate chain on the device.

Certificates not accepted by the web server

If a user selects a certificate listed in a CRL as revoked, the browser cannot load the web interface.

Certificate validation compared to ID checks

Validating HTTPS client certificates is similar to checking a person's ID against a trusted list to ensure authenticity and prevent access if the ID is invalid or revoked.

View the current HTTPS server certificate

This task enables you to verify the HTTPS server certificate currently in use on your appliance. This process helps ensure secure communications by allowing you to review certificate details.

You can only view server certificates for the appliance you are logged in to. This task is relevant when you need to confirm the certificate used for HTTPS connections on your device.

Before you begin

Access to the appliance interface is necessary. Follow these steps to view the current HTTPS server certificate.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **HTTPS Certificate**.

After completing these steps, the current HTTPS server certificate details for the appliance are displayed, allowing you to verify secure communication settings.

What to do next

Review certificate information when necessary for compliance or troubleshooting.

Generate an HTTPS server certificate signing request

Configure an HTTPS server certificate signing request (CSR) to enable secure communication between the appliance and client browsers. This task helps you generate a unique CSR for your device, which is required to obtain a signed certificate from a certificate authority.

If you install a certificate that is not signed by a globally recognized or internally trusted certificate authority, the browser displays a security warning when you attempt to connect to the web interface.

A certificate signing request (CSR) is unique to the appliance or device from which you generated it. You cannot generate a CSR for multiple devices from a single appliance. All fields are optional. However, we recommend that you enter values for CN, Organization, Organization Unit, City/Locality, State/Province, Country/Region, and Subject Alternative Name.

The key generated for the certificate request is in Base-64 encoded PEM format.

Before you begin

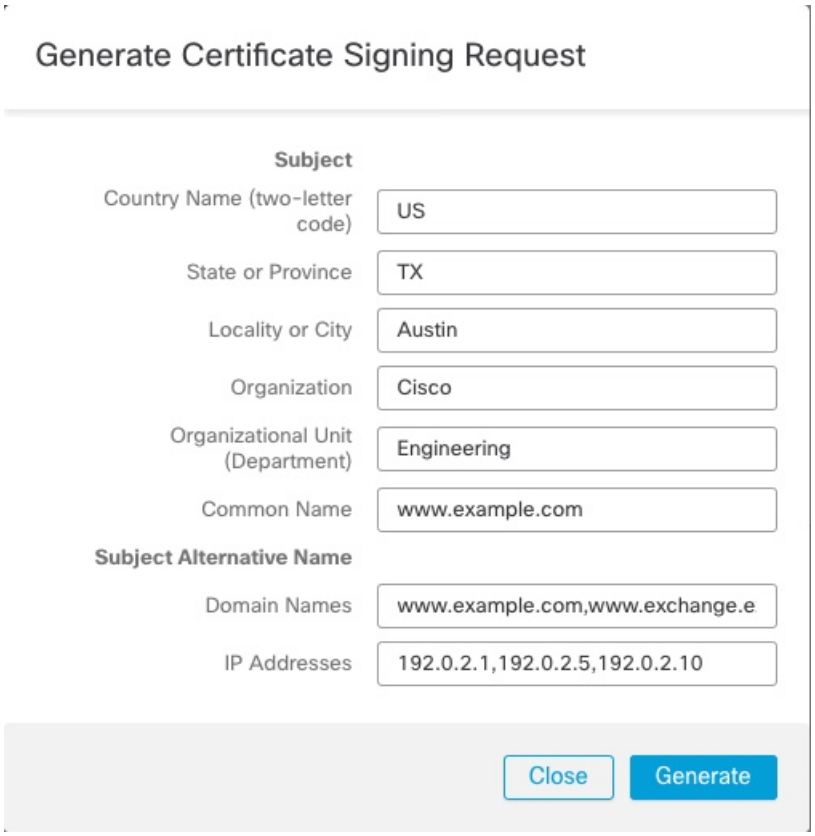
Follow these steps to generate an HTTPS server certificate signing request:

Procedure

Step 1 Choose **System** (⚙) > **Configuration** and click **HTTPS Certificate**.

Step 2 Click **Generate New CSR**.

This figure shows an example.



Step 3 Enter a country code in the **Country Name (two-letter code)** field and a state or province postal abbreviation in the **State or Province** field.

Step 4 Enter a **Locality or City**, **Organization** name, and an **Organizational Unit (Department)** name.

Step 5 Enter the fully qualified domain name of the server for which you want to request a certificate in the **Common Name** field.

Note
Enter the fully qualified domain name of the server exactly as it should appear in the certificate in the **Common Name** field. If the common name and the DNS hostname do not match, you receive a warning when connecting to the appliance.

Step 6 To request a certificate that secures multiple domain names or IP addresses, enter this information in the Subject Alternative Name section:

a) **Domain Names**: Enter the fully qualified domains and subdomains (if any) secured by the Subject Alternative Name.

b) **IP Addresses**: Enter the IP addresses secured by the Subject Alternative Name.

Step 7 Click **Generate**. Then, open a text editor.

Step 8 Copy the entire block of text in the certificate request, including the `BEGIN CERTIFICATE REQUEST` and `END CERTIFICATE REQUEST` lines. Paste it into a blank text file.

Step 9 Save the file as *`servername`*`.CSR`, where *servername* is the name of the server where you plan to use the certificate. Click **Close**.

After completing this task, you will have a Base-64 encoded PEM format certificate signing request file that you can submit to a certificate authority for signing.

What to do next

- Submit the certificate request to the certificate authority.
- When you receive the signed certificate, import it to the Firewall Management Center; refer to Import HTTPS server certificates, on page 80.

Import HTTPS server certificates

Importing HTTPS server certificates enables secure communication between users and the appliance web interface by ensuring that the server presents a trusted certificate.

If the signing authority that generated the certificate requires you to trust an intermediate CA, you must also supply a certificate chain (or certificate path).

If you require client certificates, accessing an appliance via the web interface will fail when the server certificate does not meet either of these criteria:

- The certificate is signed by the same CA that signed the client certificate.
- The certificate is signed by a CA that has signed an intermediate certificate in the certificate chain.

Caution The Firewall Management Center supports 4096-bit HTTPS certificates. If the certificate used by the Firewall Management Center was generated using a public server key larger than 4096 bits, you will not be able to log in to the web interface. For more information about updating HTTPS Certificates to Version 6.0.0, refer to the release notes. If you generate or import an HTTPS certificate and cannot log in to the Firewall Management Center web interface, contact support.

Before you begin

- Generate a certificate signing request; refer to Generate an HTTPS server certificate signing request, on page 78.
- Upload the CSR file to the certificate authority where you want to request a certificate or use the CSR to create a self-signed certificate.
- Confirm that the certificate meets the requirements described in Meet HTTPS server certificate requirements, on page 75.

Follow these steps to import HTTPS server certificates.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **HTTPS CERTIFICATE**.

Step 3 Click **Import HTTPS Server Certificate**.

Note
You cannot import an encrypted HTTPS certificate.

Step 4 Open the server certificate in a text editor, copy the entire block of text, including the `BEGIN CERTIFICATE` and `END CERTIFICATE` lines. Paste this text into the **Server CERTIFICATE** field.

Step 5 Whether you must supply a **PRIVATE KEY** depends on how you generated the certificate Signing Request:

- If you generated the certificate Signing Request using the Secure Firewall Management Center web interface (as described in Generate an HTTPS server certificate signing request, on page 78), the system already has the PRIVATE KEY and you need not enter one here.
- If you generated the certificate Signing Request using some other means, you must supply the private key here. Open the private key file and copy the entire block of text, include the `BEGIN RSA PRIVATE KEY` and `END RSA PRIVATE KEY` lines. Paste this text into the **Private Key** field.

Step 6 Open each required intermediate certificate, copy the entire block of text, and paste it into the **Certificate Chain** field. If you have a root certificate, paste it into the same field. For all certificates, copy the entire block of text, including the `BEGIN CERTIFICATE` and `END CERTIFICATE` lines.

Step 7 Click **Save**.

After completing these steps, the HTTPS server certificate is imported and the appliance web interface can be accessed securely using the new certificate.

Require valid HTTPS client certificates

This task configures the system to require users connecting to the web interface to present a valid HTTPS client certificate, enhancing authentication security. Ensure only authorized users with valid certificates can access the web interface.

Use this procedure to require users connecting to the Firewall Management Center web interface to supply a user certificate. The system supports validating HTTPS client certificates using either OCSP or imported CRLs in Privacy-enhanced Electronic Mail (PEM) format.

If you choose to use CRLS, create a scheduled task to ensure that the list of revoked certificates stays current. The system displays the most recent refresh of the CRLS. To access the web interface after enabling client certificates, you **must** have a valid client certificate present in your browser (or a CAC inserted in your reader).

Note To access the web interface after enabling client certificates, you **must** have a valid client certificate present in your browser (or a CAC inserted in your reader).

Before you begin

- Import a server certificate signed by the same certificate authority as the client certificate that will be used for the connection ; refer to Import HTTPS server certificates, on page 80.
- Import the server certificate chain if needed; refer to Import HTTPS server certificates, on page 80.

Follow these steps to require valid HTTPS client certificates.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **HTTPS Certificate**.

Step 3 Choose **Enable Client Certificates**. If prompted, select the appropriate certificate from the drop-down list.

Step 4 You have three options:

- To verify client certificates using one or more CRLS, select **Enable Fetching of CRL** and continue with Step 5.
- To verify client certificates using OCSP, select **Enable OCSP** and skip to Step 7.
- To accept client certificates without checking for revocation, skip to Step 8.

Step 5 Enter a valid URL to an existing CRL file and click **Add CRL**. Repeat this step to add up to 25 CRLs.

Step 6 Click **Refresh CRL** to load the current CRL or CRLS from the specified URLs.

Note
Enabling fetching of the CRL creates a scheduled task to regularly update the CRL or CRLS. Edit the task to set the frequency of the update.

Step 7 Verify that the client certificate is signed by the certificate authority loaded onto the appliance and the server certificate is signed by a certificate authority loaded in the browser certificate store. (These should be the same certificate authority.)

Caution
Saving a configuration with enabled client certificates, with no valid client certificate in your browser certificate store, disables all web server access to the appliance. Make sure that you have a valid client certificate installed before saving settings.

Step 8 Click **Save**.

After completing this task, only users with valid HTTPS client certificates will be able to access the web interface, and certificate validation will be enforced according to your configuration.

Related Topics

Scheduling, on page 505

Renew the default HTTPS server certificate

Renewing the default HTTPS server certificate updates the certificate used for secure communications on the appliance. This task helps maintain security by ensuring the certificate is current and valid.

You can only view and renew server certificates for the appliance you are logged in to. This procedure is relevant when the system is configured to use the default HTTPS server certificate. The renewal option appears only if the default certificate is in use.

Before you begin

Ensure you are logged in to the appliance whose certificate you wish to renew. System must be configured to use the default HTTPS server certificate.

Follow these steps to renew the default HTTPS server certificate:

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **HTTPS Certificate**.

The button appears only if your system is configured to use the default HTTPS server certificate.

Step 3 Click **Renew HTTPS Certificate**. (This option appears on the display under the certificate information only if your system is configured to use the default HTTPS server certificate.)

Step 4 (Optional) In the **Renew HTTPS Certificate** dialog box, select **Generate New Key** to generate a new key for the certificate.

Step 5 In the **Renew HTTPS Certificate** dialog box, click **Save**.

After completing these steps, the default HTTPS server certificate is renewed. If you generated a new key, the renewed certificate uses the updated key. The appliance will display updated certificate validity dates.

What to do next

You can confirm that the certificate has been renewed by checking that the certificate validity dates displayed on the **HTTPS Certificate** page have updated. Verify the new certificate information on the HTTPS Certificate page.

View system information fields

The **System** (⚙) > **Configuration** page of the web interface displays appliance information fields, most of which are read-only. These fields provide details such as name, model, serial number, software version, and network addresses.

Note Refer also to the **Help** (?) > **About** page, which includes similar but slightly different information.

Field	Description
Name	A descriptive name you assign to the Firewall Management Center appliance. Although you can use the host name as the name of the appliance, entering a different name in this field does not change the host name. This name is used in certain integrations. For example, it appears in the **Devices** list in Security Services Exchange when you integrate Firewall Management Center with Cisco XDR. If you change the name, all registered devices become outdated. A deployment is required to update the devices with the new name.

Field	Description
Product Model	The model name of the appliance.
Serial Number	The serial number of the appliance.
Software Version	The version of the software currently installed on the appliance.
Operating System	The operating system currently running on the appliance.
Operating System Version	The version of the operating system currently running on the appliance.
IPv4 Address	The IPv4 address of the default (`eth0`) management interface. If IPv4 management is disabled, this field indicates that.
IPv6 Address	The IPv6 address of the default (`eth0`) management interface. If IPv6 management is disabled, this field indicates that.
Current Policies	The system-level policies currently deployed. If a policy has been updated since it was last deployed, the name of the policy appears in italics.
Model Number	The appliance-specific model number stored on the internal flash drive. This number may be important for troubleshooting.

Intrusion policy preferences

An intrusion policy preference is a configuration option that

- enables monitoring of critical policies,
- tracks changes to those policies, and
- supports management of policy settings in a deployment.

Set intrusion policy preferences

Configure intrusion policy preferences to manage how policy changes are tracked, logged, and monitored for security intelligence. You can enable or disable comments on policy changes, audit logging, user overrides for deleted Snort 3 rules, and Talos Threat Hunting Telemetry.

This task is relevant when you need to control how intrusion policy changes are documented, audited, and monitored for advanced threat detection and compliance.

Use this configuration to ensure that policy modifications are properly tracked and that security events are sent to Cisco Talos for analysis when required. Perform this task during initial setup or when updating security and compliance requirements for your deployment.

Before you begin

Follow these steps to set intrusion policy preferences:

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **Intrusion Policy Preferences**.

Step 3 You have these options:

- **Comments on policy change**: Check this check box to track policy-related changes using the comment functionality when users modify intrusion policies. With policy change comments enabled, administrators can quickly assess why critical policies in a deployment were modified.

  If you enable comments on policy changes, you can make the comment optional or mandatory. The Firewall Management Center prompts the user for a comment when each new change to a policy is saved.

- **Write changes in Intrusion Policy to audit log**: Check this check box to record the changes to the intrusion policies to the audit logs. This option is enabled by default.

- **Retain user overrides for deleted Snort 3 rules**: Check this check box to get notifications for changes to any *overridden* system-defined rules during LSP updates. When you enable this feature, the system retains the rule overrides in the new replacement rules that are added as part of the LSP update. On the Firewall Management Center menu bar, click **Notifications** (message center). Then click **Tasks** to view the notifications. This option is enabled by default.

After completing this task, your intrusion policy preferences are set according to your selections. Policy changes are tracked, logged, and, if enabled, threat-hunting telemetry is sent to Cisco Talos for analysis.

Languages

A Language page is a configuration option that allows you to select a preferred language for the web interface.

Set the language for the web interface

Set the language for the web interface to ensure all users interact with the system in the selected language. This helps provide a consistent and localized user experience for every user accessing the web interface.

The language you specify here is used for the web interface for every user.

- English
- French
- Chinese (simplified)
- Chinese (traditional)
- Japanese
- Korean

Before you begin

Follow these steps to set the language for the web interface.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.
Step 2 Click **Language**.
Step 3 Choose the language you want to use.
Step 4 Click **Save**.

The web interface displays in the selected language for all users.

Login banners

A login banner is a security appliance feature that

- enables you to specify session, login, or custom message banners for a security appliance or shared policy,
- allows the use of ASCII characters and carriage returns to create custom messages, and
- does not preserve tab spacing and may cause Telnet or SSH sessions to fail if the banner is too large or causes errors.

Customize the login banner

Configure a custom login banner to present a specific message to users upon login.

This task is relevant when you want to customize the login experience for users accessing the system.

Before you begin

Ensure you have administrative access to the Secure Firewall Management Center.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.
Step 2 Choose **Login Banner**.
Step 3 In the **Custom Login Banner** field, enter the login banner text you want to use.
Step 4 Click **Save**.

After completing these steps, the custom login banner will be displayed to users when they access the system. During login, the new message appears and provides the intended information or instructions.

What to do next

Verify the login banner displays correctly for all users.

Management interfaces

A management interface is a network configuration element that

- allows you to change management network settings,
- enables adding more management interfaces, hostname, search domains, DNS servers, and HTTP proxy, and
- operates on the Firewall Management Center.

Firewall Management Center management interfaces

A Firewall Management Center management interface is a network interface that

- enables the Firewall Management Center to manage all devices on a single interface,
- allows initial setup and administrator login on the interface, and
- facilitates communication with the Smart Licensing server, enables downloading of updates, and supports other management functions.

Device management interfaces reference information

For information about device management interfaces, refer to *About Device Management Interfaces* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Device management

A device management system is a network administration solution that

- establishes a two-way, SSL-encrypted communication channel between the Management Center and managed devices
- enables configuration of policies, installation of software updates, and monitoring of device health status from a centralized location, and
- aggregates and correlates intrusion events, network discovery information, and device performance data for comprehensive network monitoring.

Management center device management features

The Firewall Management Center manages a device, it sets up a two-way, SSL-encrypted communication channel between itself and the device. The Firewall Management Center uses this channel to send information to the device about how you want the device to analyze and manage your network traffic. As the device evaluates the traffic, it generates events and sends them to the Firewall Management Center using the same channel.

By using the Firewall Management Center to manage devices, you can:

- configure policies for all your devices from a single location, making it easier to change configurations
- install various types of software updates on devices
- push health policies to your managed devices and monitor their health status from the Firewall Management Center

Note If you have a Security Cloud Control-managed device and are using the on-prem Firewall Management Center for analytics only, then the on-prem Firewall Management Center does not support policy configuration or upgrading. Chapters and procedures in this guide related to device configuration and other unsupported features do not apply to devices whose primary manager is Security Cloud Control.

The Firewall Management Center aggregates and correlates intrusion events, network discovery information, and device performance data, allowing you to monitor the information that your devices are reporting in relation to one another, and to assess the overall activity occurring on your network.

You can use the Firewall Management Center to manage nearly every aspect of a device's behavior.

Note Although the Firewall Management Center can manage devices running certain previous releases as specified in the compatibility matrix available at http://www.cisco.com/c/en/us/support/security/defense-center/products-device-support-tables-list.html, new features that require the latest version of Firewall Threat Defense software are not available to these previous-release devices. Some Firewall Management Center features may be available for earlier versions.

Management connections

A management connection is a communication channel that

- enables secure, TLS-1.3-encrypted communication between a device and the Firewall Management Center
- can be initiated by either the device or the Firewall Management Center, depending on initial setup, and
- requires a stable connection with at least 5 Mbps throughput, using TCP port 8305 by default.

Management connection initiation and interface selection

Devices and the Firewall Management Center can establish management connections after configuration and addition to the management center.

Initiation always begins with eth0 on the Firewall Management Center or with the lowest-numbered management interface on the device. If the connection is not established, the system tries the additional management interfaces. Having multiple management interfaces on the Firewall Management Center lets you connect to discrete networks or segregate management and event traffic. The initiator does not use the routing table to select the interface.

To prevent potential management disruption, exempt management traffic from deep inspection by applying a prefilter policy for it.

Note The management connection is a secure, TLS-1.3-encrypted communication channel between the management center and the device. You do not need to run this traffic over an additional encrypted tunnel such as Site-to-Site VPN for security purposes. If the VPN goes down, for example, you will lose your management connection, so we recommend a simple management path.

Management connection initiation scenarios

Either the device or the Firewall Management Center can initiate the management connection, or only one side may initiate, depending on the initial setup.

Management interfaces on the Firewall Management Center

A management interface is a network interface that

- enables initial setup, HTTP access for administrators, and device management on the Firewall Management Center,
- supports additional management functions such as licensing and updates, and
- can be expanded with additional interfaces to improve throughput and performance for managing large numbers of devices on different networks.

Management interface reference information

The Firewall Management Center uses the eth0 interface for initial setup, HTTP access for administrators, device management, and other management functions such as licensing and updates.

Adding more management interfaces can improve throughput and performance when managing large numbers of devices on different networks. Each management interface can be used for particular functions, such as HTTP administrator access or device management.

For device management, the management interface carries two separate traffic channels:

- The *management traffic channel* carries all internal traffic (such as inter-device traffic specific to managing the device).
- The *event traffic channel* carries all event traffic (such as web events).

You can optionally configure a separate event-only interface on the Firewall Management Center to handle event traffic; only one event interface can be configured. A management interface for the management traffic channel is always required. Separating event traffic from management traffic can improve performance, as event traffic can use a large amount of bandwidth. For example, a 10 GigabitEthernet interface can be assigned as the event interface, while 1 GigabitEthernet interfaces are used for management. An event-only interface can be configured on a secure, private network, while the regular management interface can be used on a network with Internet access. Although both management and event interfaces can be used on the same network, placing each interface on a separate network is recommended to avoid potential routing problems, including routing problems from other devices to the management center.

Managed devices send management traffic to the Firewall Management Center's management interface and event traffic to the event-only interface. If the managed device cannot reach the event-only interface, it will

fall back to sending events to the management interface. Management connections cannot be made through the event-only interface.

Management connection initiation from the Firewall Management Center is always attempted first from eth0, then other interfaces are tried in order; the routing table is not used to determine the best interface.

Note All management interfaces support HTTP administrator access as controlled by your Access List configuration (Configure an access list, on page 51). You cannot restrict an interface to only HTTP access; management interfaces always support device management (management traffic, event traffic, or both).

Note Only the eth0 interface supports DHCP IP addressing. Other management interfaces only support static IP addresses.

Default management interface functions

View management interface support per Firewall Management Center model

Refer to the hardware installation guide for your model for the management interface locations. This reference lists the supported management interfaces on each Firewall Management Center model.

Refer to the table for supported management interfaces available on each model.

Table 5: Management interface support on the Firewall Management Center

Model	Management interfaces
MC1600, MC2600, MC4600	eth0 (Default) eth1 eth2 eth3 CIMC (Supported for Lights-Out Management only.)
Firewall Management Center Virtual	eth0 (Default)

Network routes on Firewall Management Center management interfaces

A network route on a Firewall Management Center management interface is a routing configuration that

- supports only static routes to reach remote networks,
- creates a default route to the gateway IP address specified during setup, and
- uses the lowest-numbered management interface as the egress interface for the default route.

Default route behavior on management interfaces

Management interfaces can be configured with multiple static routes to access remote networks. The default route does not include an egress interface, so the interface chosen depends on the gateway address and the network to which the gateway belongs.

- Multiple management interfaces can be configured on some platforms.
- The device uses the lower-numbered interface as the egress interface when multiple interfaces exist on the default network.
- At least one static route is recommended per management interface to access remote networks.

1. Set up the Firewall Management Center and specify the gateway IP address.
2. The setup process creates a default route to the specified gateway.
3. You can modify the gateway address but cannot delete the default route.

- A default route uses the lowest-numbered management interface, such as eth0.
- A static route is recommended for each management interface to access remote networks.

Table 6: Comparison of default and static routes on management interfaces

Attribute	Default route	Static route
Interface selection	Lowest-numbered interface (e.g., eth0)	Specified interface for the route
Gateway address	Set during setup	Can be set per static route



Note The interface used for management connections is not determined by the routing table. Connections are always tried using eth0 first, and then subsequent interfaces are tried in order until the managed device is reached.

Static route configuration example for management interfaces

If you want to use interfaces on the same network, configure static routes correctly. For example, on the Firewall Management Center both eth0 and eth1 are on the same network, but you want to manage a different group of devices on each interface. The default gateway is 192.168.45.1. To manage devices on the remote 10.6.6.0/24 destination network through eth1, create a static route for 10.6.6.0/24 through eth1 with the same gateway of 192.168.45.1. Traffic to 10.6.6.0/24 will hit this route before the default route, so eth1 will be used as expected.

Static routing limitations for same network management interfaces

If you want to use two Firewall Management Center interfaces to manage remote devices that are on the same network, static routing on the Firewall Management Center may not scale well, because you need separate static routes per device IP address.

Management and event-only interface routing analogy

Separate management and event-only interfaces on both the Firewall Management Center and the managed device can be compared to having dedicated paths for different types of traffic. Event-only interfaces are on a separate network from management interfaces. In this case, add a static route through the event-only interface for traffic destined for the remote event-only network, and vice versa.

NAT environments

A NAT environment is a network configuration that

- uses network address translation (NAT) to transmit and receive network traffic by reassigning source or destination IP addresses,
- enables private networks to communicate with the internet, and
- requires NAT IDs and registration keys for device authentication and registration when IP addresses are unavailable, especially behind port address translation (PAT) routers.

NAT ID usage in device registration

Devices and management centers use NAT IDs and registration keys to establish trust and authenticate when only one IP address is available.

- Both IP addresses and registration keys are typically needed for routing and authentication.
- If only one IP address is known, a unique NAT ID must be specified on both sides to establish trust and look up the correct registration key.
- The management center and device use the registration key and NAT ID (instead of IP addresses) to authenticate and authorize for initial registration.

To simplify adding multiple devices, specify a unique NAT ID for each device while leaving the IP address blank on the management center, and specify both the management center IP address and NAT ID on each device.

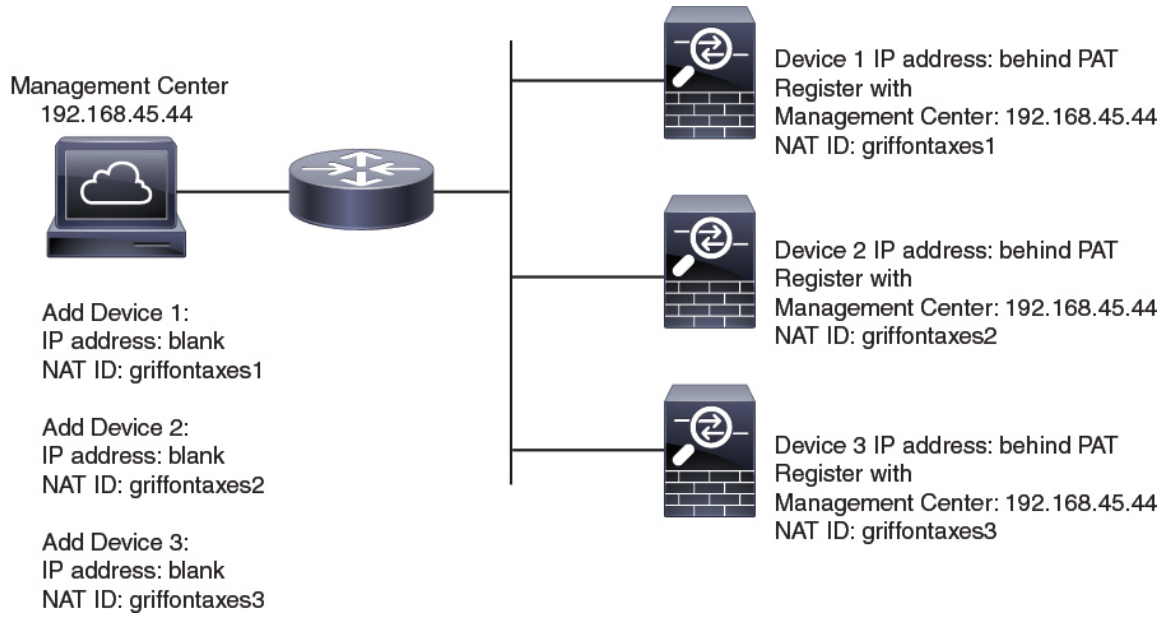


Note The NAT ID must be unique for each device to ensure proper authentication and registration.

Examples of NAT ID usage in PAT environments

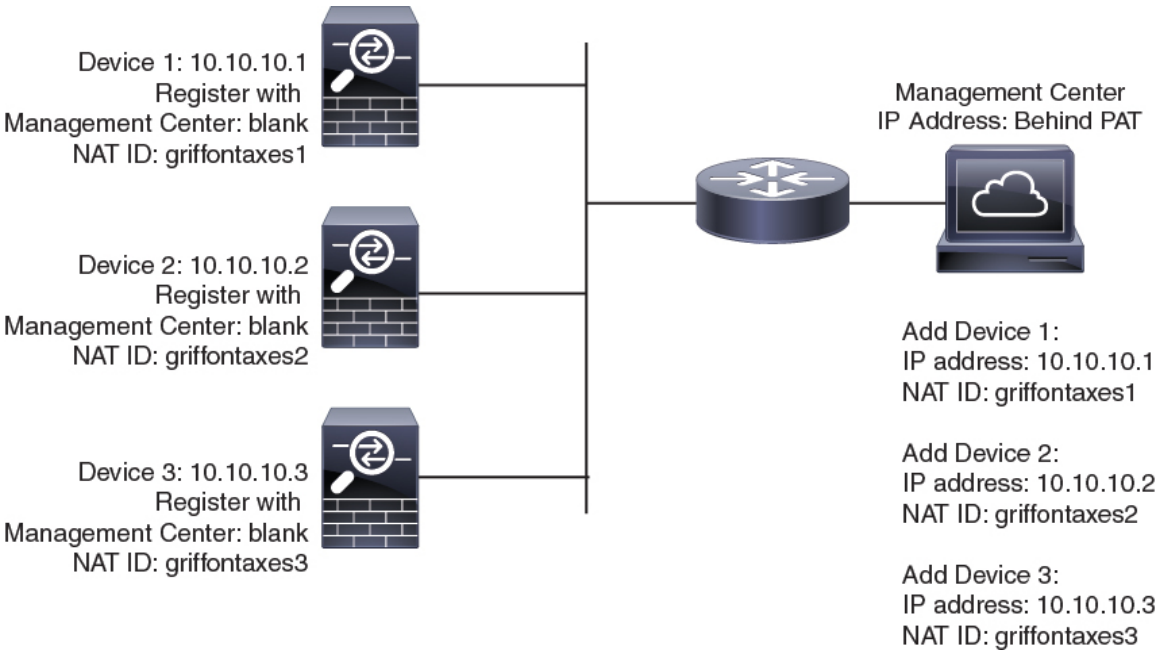
Each of the three devices behind a PAT IP address requires a unique NAT ID on both the management center and the devices. The management center IP address must also be specified on each device.

Figure 6: NAT ID for managed devices behind PAT



The management center behind a PAT IP address requires a unique NAT ID for each device on both the management center and the devices. The device IP addresses must also be specified on the management center.

Figure 7: NAT ID for Firewall Management Center behind PAT

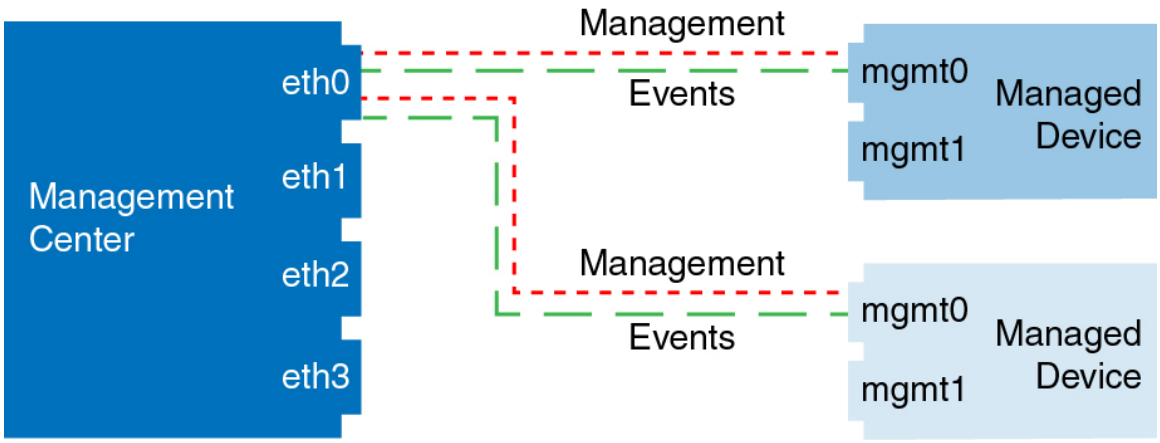


Recommendation: management and event traffic channel configuration

- Configure a separate, dedicated event interface on both the Firewall Management Center and the managed device to segregate event traffic from management traffic. This separation improves performance and security by isolating event data flows.
- If Firewall Management Center uses a single interface for both management and event traffic, the managed device must send event traffic on that same interface. If the Firewall Management Center has a dedicated event interface, the managed device must also have a dedicated event interface configured.
- If management and event channels are mixed on the same interface on the Firewall Management Center, and the managed device uses a separate event interface, connectivity or event delivery issues can occur. Maintain consistent interface pairing between the Firewall Management Center and managed devices.
- If you use a data interface for management on a Firewall Threat Defense, you cannot use separate management and event interfaces for that device.

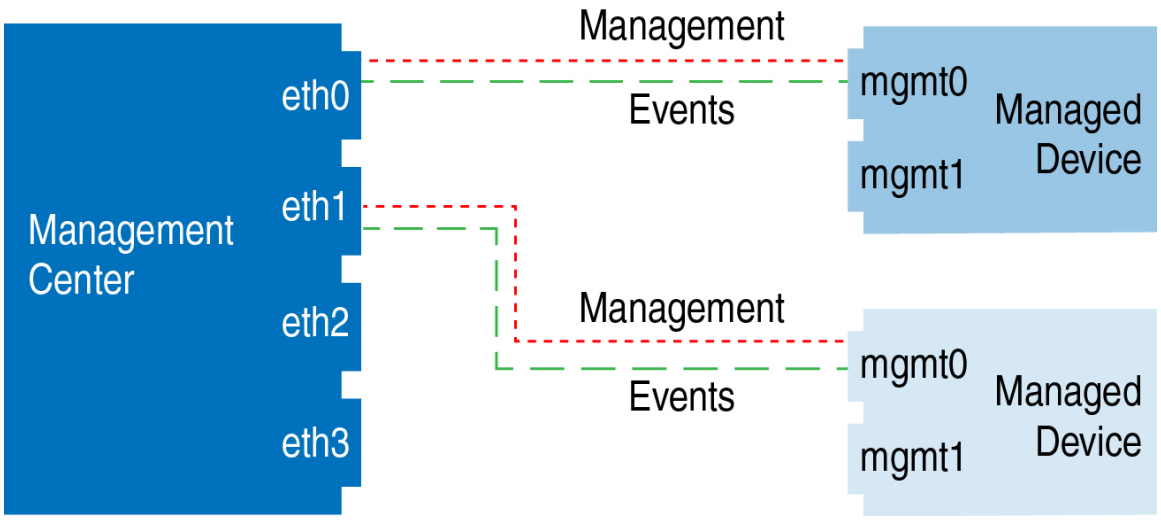
This example illustrates multiple management interfaces and a separate event interface on the Firewall Management Center. Managed devices may use a separate event interface or a single management interface.

Figure 8: Single Management Interface on the Secure Firewall Management Center



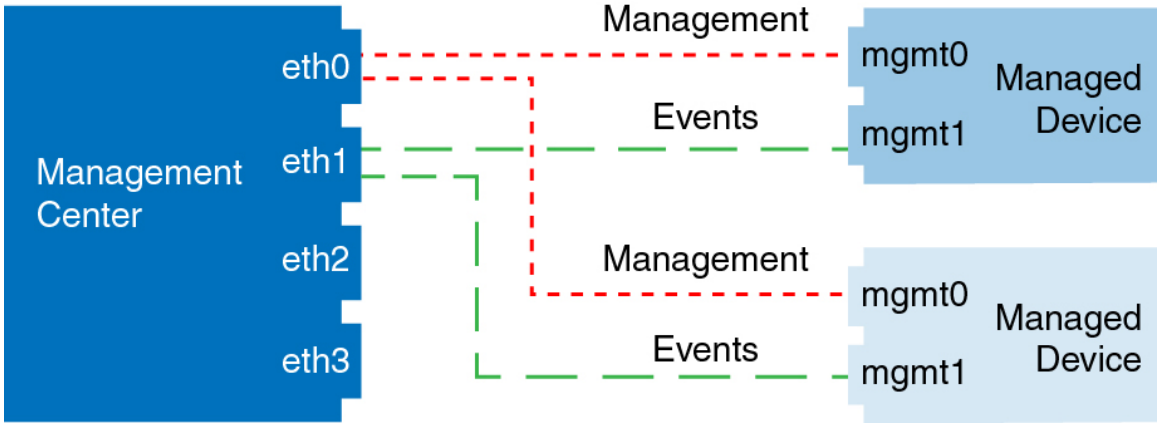
This example shows the Firewall Management Center using separate management interfaces for devices; and each managed device using 1 management interface.

Figure 9: Multiple Management Interfaces on the Secure Firewall Management Center



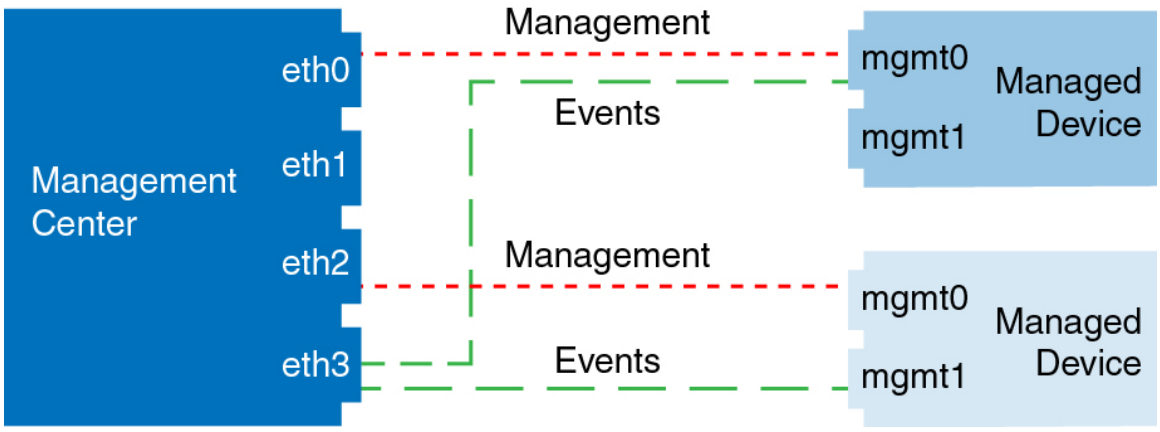
This example shows the Firewall Management Center and managed devices using a separate event interface.

Figure 10: Separate Event Interface on the Secure Firewall Management Center and Managed Devices



This example shows a mix of multiple management interfaces and a separate event interface on the Firewall Management Center and a mix of managed devices using a separate event interface, or using a single management interface.

Figure 11: Mixed management and event interface usage



Modify Firewall Management Center management interfaces

Modify the management interface settings on the Firewall Management Center. You can optionally enable additional management interfaces or configure an event-only interface.

Caution Be careful when making changes to the management interface to which you are connected; if you cannot reconnect because of a configuration error, you must access the Firewall Management Center console port to reconfigure the network settings in the Linux shell. You must contact Cisco TAC to guide you in this operation.

If you change the Firewall Management Center IP address, then refer to *Edit the Firewall Management Center IP Address or Hostname on the Device* in the *Cisco Secure Firewall Management Center Device Configuration Guide*. If you change the Firewall Management Center IP address or hostname, you should also change the value at the device CLI so the configurations match. Although in most cases, the management connection will be reestablished without changing the Firewall Management Center IP address or hostname on the device, in at least one case, you must perform this task for the connection to be reestablished: when you added the device to the Firewall Management Center and you specified the NAT ID only. For better network resiliency, keep the Firewall Management Center's IP address or hostname current in all cases.

In a high-availability configuration, when you modify the management IP address of a registered device from the device CLI or from the Firewall Management Center, the standby Firewall Management Center does not reflect the changes even after a high-availability synchronization. To ensure that the standby Firewall Management Center is also updated, modify the management IP address of the registered device on the **Device Management** page of the standby Firewall Management Center.

If you modify the management IP address of one peer Firewall Management Center in a high availability configuration, the remote peer does not reflect the changes even after an high availability synchronization. To update the remote peer Firewall Management Center, log in to the remote peer Firewall Management Center, navigate to **Integration** > **Other Integrations** > **High Availability**, click **Peer Manager**, and then manually update its peer manager IP address. For more detailed instructions, refer to Change the IP address of the Firewall Management Center in a high availability pair, on page 335.

Before you begin

- For information about how device management works, refer to *About Device Management Interfaces* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Procedure

Step 1 Choose **System** (⚙) > **Configuration** > **Management Interfaces**.

Step 2 In the **Interfaces** area, click **Edit** next to the interface that you want to configure.

All available interfaces are listed in this section. You cannot add more interfaces.

You can configure these options on each management interface:

- **Enabled** —Enable the management interface. Do **not** disable the default eth0 management interface. Some processes require the eth0 interface.
- **Channels** —You must always have at least one interface with **Management Traffic** enabled. You can optionally configure an event-only interface. You can configure only one event interface on the Firewall Management Center. To do so, uncheck the **Management Traffic** check box, and leave the **Event Traffic** check box checked. You can optionally disable **Event Traffic** for the remaining management interfaces. In either case, the device tries to send events to the event-only interface, and if that interface is down, it sends events on the management interface even if you disable the event channel. You cannot disable both event and management channels on an interface.
- **Mode** —Specify a link mode. Note that any changes you make to auto-negotiation are ignored for Gigabit Ethernet interfaces.
- **MDI/MDIX** —Set the **Auto-MDIX** setting.
- **MTU** —Set the maximum transmission unit (MTU) between 1280 and 1500. The default is 1500.
- **IPv4 Configuration** —Set the IPv4 IP address. Choose:
  - **Static** —Manually enter the **IPv4 Management IP** address and **IPv4 Netmask**.
  - **DHCP** —Set the interface to use DHCP (eth0 only).

    If you use DHCP, you must use DHCP reservation, so the assigned address does not change. If the DHCP address changes, device registration will fail because the Firewall Management Center network configuration gets out of sync. To recover from a DHCP address change, connect to the Firewall Management Center (using the hostname or the new IP address) and navigate to **System** (⚙) > **Configuration** , and then click **Management Interfaces** to reset the network.
  - **Disabled** —Disable IPv4. Do **not** disable both IPv4 and IPv6.
- **IPv6 Configuration** —Set the IPv6 IP address. Choose:
  - **Static** —Manually enter the **IPv6 Management IP** address and **IPv6 Prefix Length**.
  - **DHCP** —Set the interface to use DHCPv6 (eth0 only).
  - **Router Assigned** —Enable stateless autoconfiguration.
  - **Disabled** —Disable IPv6. Do **not** disable both IPv4 and IPv6.

- **IPv6 DAD** —When you enable IPv6, enable or disable duplicate address detection (DAD). You might want to disable DAD because the use of DAD opens up the possibility of denial-of-service attacks. If you disable this setting, you need check manually that this interface is not using an already-assigned address.

Step 3 In the **Routes** area, edit a static route by clicking **Edit** ( ), or add a route by clicking **Add** ( ).

Click the **View** ( ) icon to view the route table.

You need a static route for each additional interface to reach remote networks. For more information about when new routes are needed, refer to Network routes on Firewall Management Center management interfaces, on page 90.

Note
For the default route, you can change only the gateway IP address. The egress interface is chosen automatically by matching the specified gateway to the interface's network.

You can configure these settings for a static route:

- **Destination** —Set the destination address of the network to which you want to create a route.
- **Netmask** or **Prefix Length** —Set the netmask (IPv4) or prefix length (IPv6) for the network.
- **Interface** —Set the egress management interface.
- **Gateway** —Set the gateway IP address.

Step 4 In the **Shared Settings** area, set network parameters shared by all interfaces.

Note
If you selected **DHCP** for the eth0 interface, you cannot manually specify some shared settings derived from the DHCP server.

You can configure these shared settings:

- **Hostname**—Set the Firewall Management Center hostname. The hostname can have a maximum of 64 characters, must start and end with a letter or digit, and have only letters, digits, or a hyphen. If you change the hostname, reboot the Firewall Management Center if you want the new hostname reflected in syslog messages. Syslog messages do not reflect a new hostname until after a reboot.
- **Domains**—Set one or more search domains for the Firewall Management Center, separated by commas. The system adds these domains to hostnames if a command does not specify a fully-qualified domain name, for example, **ping system**. These domains are used only on the management interface or for commands routed through the management interface.
- **Primary DNS Server**, **Secondary DNS Server**, **Tertiary DNS Server**—Set the DNS servers to be used in order of preference.
- **Remote Management Port**—Set the remote management port for communication with managed devices. The Firewall Management Center and managed devices communicate using a two-way, SSL-encrypted communication channel, which by default is on port 8305.

  Note
  Cisco **strongly** recommends that you keep the default settings for the remote management port, but if the management port conflicts with other communications on your network, you can choose a different

port. If you change the management port, you must change it for **all** devices in your deployment that need to communicate with each other.

Step 5 In the **ICMPv6** area, configure ICMPv6 settings.

- **Allow Sending Echo Reply Packets** —Enable or disable Echo Reply packets. You might want to disable these packets to guard against potential denial of service attacks. Disabling Echo Reply packets means you cannot use IPv6 ping to the Firewall Management Center management interfaces for testing purposes.
- **Allow Sending Destination Unreachable Packets**—Enable or disable Destination Unreachable packets. You might want to disable these packets to guard against potential denial of service attacks.

Step 6 In the **Proxy** area, configure HTTP proxy settings.

The Firewall Management Center is configured to directly connect to the internet on ports TCP/443 (HTTPS) and TCP/80 (HTTP). You can use a proxy server, to which you can authenticate via HTTP Digest.

Changing proxy settings causes all network-related services to be restarted.

Proxies that use NT LAN Manager (NTLM) authentication are not supported.

a) Check the **Enabled** check box.
b) In the **HTTP Proxy** field, enter the IP address or fully qualified domain name of your proxy server.
   - For Smart Licensing, the proxy FQDN cannot have more than 64 characters .
   - For IPv6, you can only specify an FQDN for the proxy, not an IP address.
c) In the **Port** field, enter a port number.
d) Supply authentication credentials by choosing **Use Proxy Authentication**, and then provide a **User Name** and **Password**.

   For Smart Licensing, the proxy password should not contain special characters other than hyphen (-), underscore (_) and period (.) .

Step 7 Click **Save**.

Step 8 If you change the Firewall Management Center IP address, then refer to *Edit the Firewall Management Center IP Address or Hostname on the Device* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

If you change the Firewall Management Center IP address or hostname, you should also change the value at the device CLI so the configurations match. Although in most cases, the management connection will be reestablished without changing the Firewall Management Center IP address or hostname on the device, in at least one case, you must perform this task for the connection to be reestablished: when you added the device to the Firewall Management Center and you specified the NAT ID only. Even in other cases, we recommend keeping the Firewall Management Center IP address or hostname up to date for extra network resiliency.

Change both Firewall Management Center and Threat Defense IP addresses

You might want to change both Firewall Management Center and Firewall Threat Defense IP addresses if you need to move them to a new network.

Procedure

Step 1 Disable the management connection.

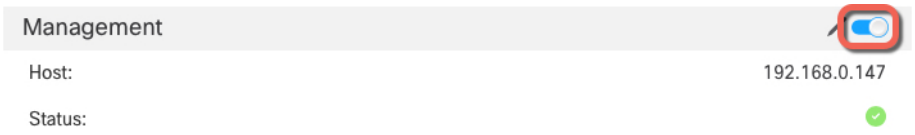
For a high-availability pair or cluster, perform these steps on all units.

a) Choose **Devices** > **Device Management**.

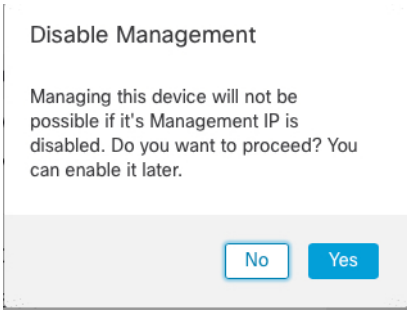
b) Next to the device, click **Edit** ( ).

c) Click **Device**, and view the **Management** area.

d) Disable management temporarily by clicking the slider so it is disabled ( ).



You are prompted to proceed with disabling management; click **Yes**.

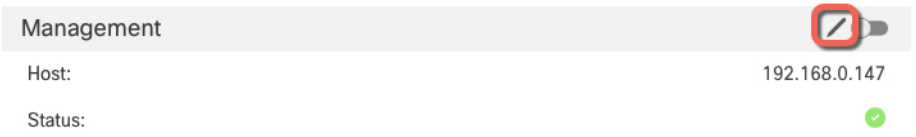


Step 2 Change the device IP address in the Firewall Management Center to the new device IP address.

You will change the IP address on the device later.

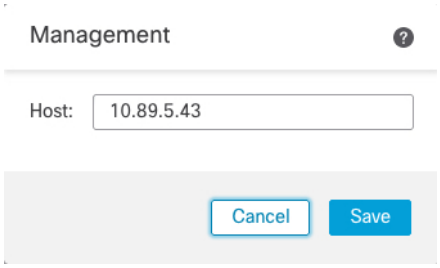
For a high-availability pair or cluster, perform these steps on all units.

a) Edit the **Host** IP address or hostname by clicking **Edit** ( ).



b) In the **Management** dialog box, modify the name or IP address in the **Host** field , and click **Save**.

Figure 12: Management IP Address



Step 3 Change the Firewall Management Center IP address.

Caution
Be careful when making changes to the Firewall Management Center interface to which you are connected. If you cannot re-connect because of a configuration error, you need to access the Firewall Management Center console port to re-configure the network settings in the Linux shell. You must contact Cisco TAC to guide you in this operation.

a) Choose **System** (⚙) > **Configuration** > **Management Interfaces**.
b) In the **Interfaces** area, click **Edit** next to the interface that you want to configure.
c) Change the IP address, and click **Save**.

Step 4 Change the manager IP address on the device.

For a high-availability pair or cluster, perform these steps on all units.

a) At the Firewall Threat Defense CLI, view the Firewall Management Center identifier.

show managers

Example:

```
> show managers
Type                      : Manager
Host                      : 10.10.1.4
Display name              : 10.10.1.4
Identifier                : f7ffad78-bf16-11ec-a737-baa2f76ef602
Registration              : Completed
Management type           : Configuration
```

b) Edit the Firewall Management Center IP address or hostname.

configure manager edit *identifier* { **hostname** { *ip_address* | *hostname* } | **displayname** *display_name* }

If the Firewall Management Center was originally identified by **DONTRESOLVE** and a NAT ID, you can change the value to a hostname or IP address using this command. You cannot change an IP address or hostname to **DONTRESOLVE**.

Example:

```
                                > configure manager edit
f7ffad78-bf16-11ec-a737-baa2f76ef602
                                hostname
                                10.10.5.1
```

Step 5 Change the IP address of the manager access interface at the console port.

For a high-availability pair or cluster, perform these steps on all units.

If you use the dedicated Management interface:

configure network ipv4

configure network ipv6

If you use the dedicated Management interface:

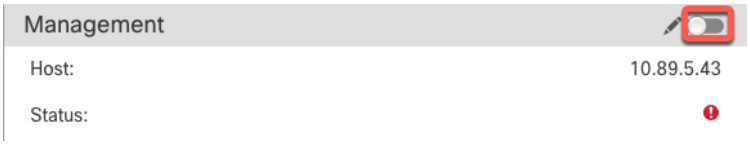
configure network management-data-interface disable

configure network management-data-interface

Step 6 Re-enable management by clicking the slider ( ).

For a high-availability pair or cluster, perform these steps on all units.

Figure 13: Enable Management Connection



Step 7 (If you use a data interface for manager access) Refresh the data interface settings in the Firewall Management Center.

For a high-availability pair, perform this step on both units.

a) Choose **Devices** > **Device Management**, and click **Manager Access - Configuration Details** , and then click **Refresh**.
b) Choose **Devices** > **Device Management**, and click the **Interfaces** tab and set the IP address to match the new address.
c) Return to the **Manager Access - Configuration Details** dialog box, and click **Acknowledge** to remove the deployment block.

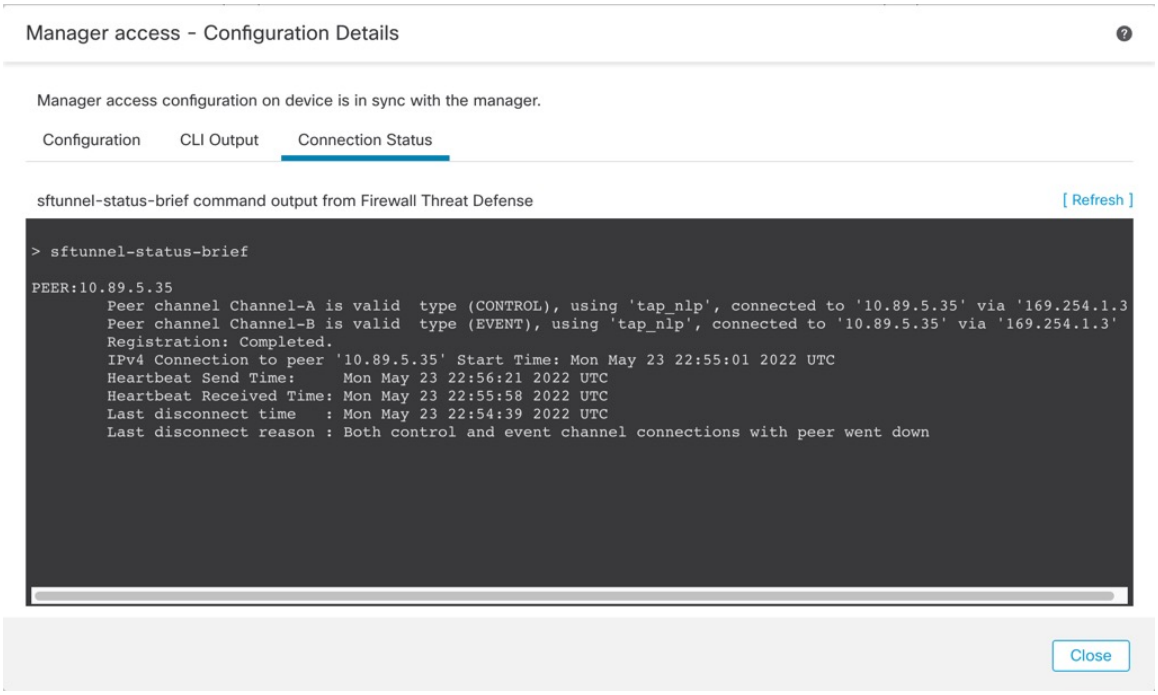
Step 8 Ensure the management connection is reestablished.

In the Firewall Management Center, check the management connection status. Navigate to the **Devices** > **Device Management**, and click **Management** section under the **Device** tab. Then, click **Manager Access - Configuration Details** to view the **Connection Status** page.

At the Firewall Threat Defense CLI, enter the **sftunnel-status-brief** command to view the management connection status.

This status indicates a successful connection for a data interface and displays the internal "tap_nlp" interface.

Figure 14: Connection Status



Step 9 (For a high-availability Firewall Management Center pair) Repeat configuration changes on the secondary Firewall Management Center.

a) Change the secondary Firewall Management Center IP address.
b) Specify the new peer addresses on both units.
c) Make the secondary unit the active unit.
d) Disable the device management connection.
e) Change the device IP address in the Firewall Management Center.
f) Re-enable the management connection.

Network analysis policy preferences

Network analysis policy preferences are configuration settings that enable tracking and documentation of policy modifications within the system.

Policy Change Tracking Options

You can configure the system to track policy-related changes using the comment functionality when users modify network analysis policies. With policy change comments enabled, administrators can quickly assess why critical policies in a deployment were modified.

If you enable comments on policy changes, you can make the comment optional or mandatory. The system prompts the user for a comment when each new change to a policy is saved.

Optionally, you can have changes to network analysis policies written to the audit log.

Processes

A process is a system operation that

- can be shut down gracefully using the web interface
- can be rebooted to restart the appliance, and
- can restart communications, database, and HTTP server processes for troubleshooting.

Shutdown and restart options

Use the web interface to control the shutdown and restart of processes on the Firewall Management Center. You can perform these actions:

- Shut down: Initiate a graceful shutdown of the appliance.

Caution Do **not** shut off Secure Firewall appliances using the power button; it may cause a loss of data. Using the web interface (or CLI) prepares the system to be safely powered off and restarted without losing configuration data.

- Reboot: Shut down and restart gracefully.
- Restart the console: Restart the communications, database, and HTTP server processes. This is typically used during troubleshooting.

Tip For virtual devices, refer to the documentation for your virtual platform. For VMware in particular, custom power options are part of VMware Tools.

Process control example

For example, shutting down the appliance using the web interface ensures configuration data is preserved, while restarting the console can help resolve communication or database issues.

Improper shutdown counter-example

Shutting off the appliance using the power button may cause a loss of data and should not be used as a shutdown method.

Process management analogy

Managing processes through the web interface is similar to safely closing applications on a computer before turning it off to prevent data loss.

Shut down or restart the Firewall Management Center

This task enables you to shut down, reboot, or restart the console of the Secure Firewall Management Center for maintenance or troubleshooting purposes.

Perform this task when you need to power off, reboot, or restart the console, for example, during maintenance, troubleshooting, or system updates. Be aware that rebooting or restarting may temporarily disrupt access and system operations.

Before you begin

Follow these steps to shut down, reboot, or restart the Secure Firewall Management Center:

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Choose **Process**.

Step 3 Do one of these actions:

Shut down	Click **Run Command** next to **Shutdown Management Center**.
Reboot	Click **Run Command** next to **Reboot Management Center**. **Note** Rebooting logs you out, and the system runs a database check that can take up to an hour to complete.
Restart the console	Click **Run Command** next to **Restart Management Center Console**. **Note** Restarting may cause deleted hosts to reappear in the network map.

The Secure Firewall Management Center will shut down, reboot, or restart the console as selected. The chosen action will affect system availability.

REST API preferences

A REST API preference is a configuration option that

- provides a lightweight interface for third-party applications to view and manage device configuration using a REST client and standard HTTP methods,
- enables access control for applications using the REST API, and
- allows administrators to block or permit requests from applications as needed.

REST API reference information

The management center REST API provides a lightweight interface for third-party applications to view and manage device configuration using a REST client and standard HTTP methods.

For more information on the management center REST API, refer to the *Secure Firewall Management Center REST API Quick Start Guide*.

- REST API preferences allow requests from applications by default.
- You can configure the Firewall Management Center to block this access.

Note The REST API for the management center does not support HTTPS certificates.

Enable REST API access

Enable or disable REST API access to allow or restrict API-based management and automation.

Note In deployments using the Firewall Management Center high availability, this feature is available only in the active Firewall Management Center.

Use this task when you need to manage or automate your Secure Firewall Management Center using REST APIs. This is relevant for deployments where API access must be controlled for security or operational reasons.

Before you begin

Follow these steps to enable or disable REST API access:

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **REST API Preferences**.

Step 3 To enable or disable REST API access to the Firewall Management Center, check or uncheck the **Enable REST API** check box.

Step 4 Click **Save**.

Step 5 Access the REST API Explorer at:

```
https://<management_center_IP_or_name>:<https_port>/API/API-explorer
```

REST API access is enabled or disabled as configured. You can now use the REST API Explorer to interact if access is enabled.

Remote console access management

A remote console access management system is a hardware management feature that

- enables users to perform limited tasks, such as viewing the chassis serial number,
- allows monitoring of hardware conditions, including fan speed and temperature, and
- provides access through a Linux command line interface using multiple interfaces.

Console access interfaces

Remote console access management supports these interfaces for physical Firewall Management Centers:

- VGA port (the default)
- Serial port
- Lights-Out Management (LOM) on a Serial Over LAN (SOL) on the CIMC port

Configure remote console settings

You can configure remote console settings to enable access to the system console through VGA, serial port, or Lights-Out Management (LOM). Administrators can select and set up a preferred remote console access method to manage the system effectively.

By default, you can access the Firewall Management Center console using the VGA port. To use the serial port or Lights-Out Management (LOM) serial-over-LAN (SOL) with the CIMC port, follow these steps.

Before you begin

- If you plan to enable LOM, refer to the *Getting Started Guide* for your Firewall Management Center for information about installing and using an Intelligent Platform Management Interface (IPMI) utility.
- Disable Spanning Tree Protocol (STP) on any third-party switching equipment connected to the Firewall Management Center CIMC interface.
- You must be an Admin user to perform this procedure.

Follow these steps to configure remote console settings:

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **Console Configuration**.

Step 3 Choose a remote console access option:

- Choose **VGA** to use the VGA port. This is the default.
- Choose **Physical Serial Port** to use the serial port.

- Choose **Lights-Out Management** to use an SOL connection on the Firewall Management Center CIMC port.

Step 4 Configure the CIMC port network settings for use with LOM.

- Choose the address **Configuration** for the system (**DHCP** or **Manual**).
- If you chose manual configuration, enter the necessary IPv4 settings:
  - Enter the **IP Address**.

    Note
    We recommend putting the CIMC port on a different network from the Firewall Management Center Management interface.

  - Enter the **Netmask**.
  - Enter the **Default Gateway**.

Step 5 Click **Save**.

Step 6 The system displays this warning: "You will have to reboot your system for these changes to take effect." Click **OK** to reboot the system immediately or **Cancel** to to postpone the reboot.

After completing these steps, the remote console settings are configured, and you can access the system console using the selected method. Changes take effect after a system reboot if prompted.

What to do next

- If you configured serial access, ensure the rear panel serial port connects to a local computer, terminal server, or other device that supports remote serial access over ethernet. Details are available in the *Getting Started Guide* for your Firewall Management Center model.
- If you configured Lights-Out Management, enable a Lights-Out Management user; refer to Enable Lights-Out Management user access, on page 108.

Enable Lights-Out Management user access

This task enables you to grant Lights-Out Management (LOM) access to an existing user, allowing them to perform remote management operations on the system. Only authorized users with the Administrator role can access LOM features. Enforcing username and password requirements for LOM users helps maintain system security.

Use this task to grant LOM access to an existing user. To grant LOM access to a new user, refer to Add or edit an internal user, on page 148.

You must explicitly grant Lights-Out Management permissions to users who use the feature. LOM users also face these restrictions.

- You must assign the Administrator role to the user.
- The username may have up to 16 alphanumeric characters. Hyphens and longer usernames are not supported for LOM users.

- A user's LOM password is the same as that user's system password. The password must comply with the requirements described in User passwords, on page 142. Cisco recommends that you use a complex, non-dictionary-based password of the maximum supported length for your appliance and change it every three months.
- Physical Firewall Management Centers can have up to 13 LOM users.

If you deactivate and then reactivate a user with LOM while the user is logged in, the user may need to log back into the web interface to regain access to `ipmitool` commands.

Note High Availability synchronization is not applicable for LOM users and hence they are not replicated on high availability Firewall Management Centers. You must create different admin users with LOM enabled on the active Firewall Management Center.

In a high-availability configuration, when you create a local user or reset the password for a local user with LOM privilege enabled, from the UCS-based active Firewall Management Center, the changes get synced to both the active and standby Firewall Management Centers and the active Firewall Management Center CIMC. The new password is not synced with the standby Firewall Management Center for CIMC login. To ensure that the standby Firewall Management Center is also updated, reset the CIMC login password for the local user on the standby Firewall Management Center.

Before you begin

You must be an Admin user to perform this procedure. Follow these steps to enable Lights-Out Management user access:

Procedure

Step 1 Choose **System** (⚙) > **Users** > **Users**.

Step 2 To grant LOM user access to an existing user, click **Edit** (✎) next to a user name in the list.

Step 3 Under **User Configuration**, enable the Administrator role.

Step 4 Check the **Allow Lights-Out Management Access** check box.

Step 5 Click **Save**.

After completing this task, the selected user will have Lights-Out Management access with the Administrator role, subject to username and password restrictions. The user can now perform remote management operations as permitted.

Use IPMI with serial over LAN

To create a Serial Over LAN connection to the Firewall Management Center, use a third-party IPMI utility on your computer. On Linux-like or macOS systems, use IPMItool. On Windows, use IPMIutil or IPMItool depending on your Windows version.

Note Cisco recommends using IPMItool version 1.8.12 or greater.

Linux

IPMItool is standard with many distributions and is ready to use.

Mac

You must install IPMItool on a Mac. First, confirm that your Mac has Apple's XCode Developer tools installed. Ensure the optional components for command line development are present: UNIX Development and System Tools in newer versions, or Command Line Support in older versions. Next, install MacPorts and IPMItool. Search online for more information or refer to these websites:

```
https://developer.apple.com/technologies/tools/
http://www.macports.org/
http://github.com/ipmitool/ipmitool/
```

Windows

For Windows Versions 10 and later with Windows Subsystem for Linux (WSL) enabled, as well as some older versions of Windows Server, you can use IPMItool. If your version is not supported, compile IPMIutil on your Windows system. You can use IPMIutil itself to complete the compilation. For more information, refer to this website:

```
http://ipmiutil.sourceforge.net/man.html#ipmiutil
```

Understand IPMI utility commands

Commands for IPMI utilities are composed of segments, for example, when using IPMItool on a Mac:

```
ipmitool -I lanplus -H IP_address -U user_name command
```

Definitions:

- `ipmitool` invokes the utility.
- `-I lanplus` specifies to use an encrypted IPMI v2.0 RMCP+ LAN Interface for the session.
- `-H IP_address` indicates the IP address you have configured for Lights-Out Management on the Firewall Management Center you want to access.
- `-U user_name` is the name of an authorized remote session user.
- *`command`* is the name of the command you want to use.

Note Cisco recommends using IPMItool version 1.8.12 or later.

To connect on Windows using IPMIutil, use this command:

```
ipmiutil command -V 4 -J 3 -N IP_address -Uuser_name
```

This command connects you to the command line on the Firewall Management Center where you can log in as if you are physically present near the appliance. You may be prompted to enter a password.

Configure serial over LAN with IPMItool

Configure Serial Over LAN (SOL) with IPMItool to enable remote console access to the server for management and troubleshooting purposes.

You must have Admin user privileges and LOM access to perform this procedure. Use this task to access a server remotely through Serial Over LAN with IPMItool.

Before you begin

You must have Admin privileges and LOM (Lights Out Management) access. Follow these steps to configure Serial Over LAN with IPMItool:

Procedure

Using IPMItool, enter the following command, and a password if prompted:

Example:

```
ipmitool -I lanplus -H IP_address -U user_name sol activate
```

After you complete this task, Serial Over LAN activates and allows remote access to the server console with IPMItool.

Configure serial over LAN with IPMIutil

This task enables you to configure Serial Over LAN (SOL) access using IPMIutil, allowing remote management of the server's serial console.

You must be an Admin user with LOM access to perform this procedure. Use this procedure when you need to access the server's serial console remotely via the network using IPMIutil.

Before you begin

Make sure that IPMIutil is installed. Verify network access to the target device. Confirm that you have Admin privileges and Lights Out Management (LOM) access.

Follow these steps to configure Serial Over LAN with IPMIutil:

Procedure

Use IPMIutil to enter this command. Enter a password if prompted.

Example:

```
ipmiutil -J 3 -N IP_address -U username sol -a
```

After completing this task, you will have established Serial Over LAN access to the server using IPMIutil, enabling remote serial console management.

Recommendation: using IPMI with Lights-Out management

Recommendation: secure and reliable use of Lights-Out management

Lights-Out Management (LOM) allows you to perform limited actions remotely over a SOL connection without logging into the system. After completing a command, the connection ends. If all attempts to access your system fail, you can use LOM to restart your system remotely. However, restarting while the SOL connection is active may disconnect or time out the session.

Caution In rare cases, if your computer is on a different subnet than the LOM interface, and the Firewall Management Center gets its IP address using DHCP, attempting to access LOM features can fail. If this occurs, you can either disable and then re-enable LOM on the Firewall Management Center, or use a computer on the same subnet as the Firewall Management Center to ping its LOM interface. You should then be able to use LOM.

Caution Cisco is aware of a vulnerability inherent in the Intelligent Platform Management Interface (IPMI) standard (CVE-2013-4786). Enabling Lights-Out Management (LOM) on a system exposes this vulnerability. To mitigate this vulnerability, deploy your systems on a secure management network accessible only to trusted users. Use a complex, non-dictionary-based password of the maximum supported length for your system, and change the password every three months. To prevent exposure to this vulnerability, do not enable LOM.

Caution Do **not** restart your system unless it does not respond to any other attempts to restart. Remotely restarting does not gracefully reboot the system and you may lose data.

- Lights-Out Management commands for IPMItool and IPMIutil include options for admin privileges, encryption, specifying LOM IP address or hostname, authorized username, starting and ending SOL sessions, power cycling, powering up or down, and displaying system information such as fan speeds and temperatures.
- For example, to display a list of Firewall Management Center information, the IPMItool command is:

```
ipmitool -I lanplus -H IP_address -U user_name sdr
```

The same command with the IPMIutil utility is:

```
ipmiutil sensor -V 4 -J 3 -N IP_address -U user_name
```

Note Cisco recommends using IPMItool version 1.8.12 or a later version.

These principles apply to systems using Lights-Out Management (LOM) with IPMI for remote management and control.

Ensuring secure and reliable use of LOM prevents unauthorized access, mitigates vulnerabilities, and avoids data loss during remote operations.

Following these recommendations enables safe remote management, reduces risk, and maintains system integrity.

Deploy systems on secure networks, use strong passwords, and follow proper restart procedures to minimize exposure and operational issues.

Configure lights-out management with IPMItool

This task enables you to configure lights-out management (LOM) on a server using IPMItool, allowing remote management and monitoring of server hardware.

You must be an Admin user with LOM access to perform this procedure. Use this task when you need to manage server hardware remotely without physical access.

Before you begin

Verify that you have Admin privileges and LOM access before starting.

Follow these steps to configure lights-out management with IPMItool:

Procedure

Run the IPMItool command and enter the password if you are prompted.

```
ipmitool -I lanplus -H IP_address -U user_name command
```

After you complete this task, you can remotely manage and monitor the server hardware using IPMItool commands.

Configure lights-out management with IPMIutil

This task enables you to configure lights-out management (LOM) on a server using the IPMIutil utility. It allows remote management and monitoring of server hardware through IPMI commands.

You must be an Admin user with LOM access to perform this procedure.

Use this task to manage server hardware remotely for purposes such as troubleshooting or maintenance with IPMIutil.

This procedure applies in environments that require remote server management.

Before you begin

Ensure you have Admin privileges and LOM access before proceeding. Verify that IPMIutil is installed on your system.

Follow these steps to configure lights-out management with IPMIutil:

Procedure

Enter this command for IPMIutil and a password if prompted:

Example:

```
ipmiutil -J 3 -N IP_address -U username command
```

After you complete this task, you will have enabled lights-out management on the server, allowing remote hardware management using IPMIutil commands.

Remote storage devices

You can store Firewall Management Center backups and reports locally or use one of these remote systems:

- Network File System (NFS)
- Server Message Block (SMB)/Common Internet File System (CIFS)
- Secure Shell Filesystem (SSHFS)

Remote storage device characteristics

You can switch among the storage systems (NFS, SMB, SSHFS) at any time. However, the system does not retain the previous configuration settings for these storage types when you switch between them. You must re-enter the configuration details whenever you change to a different storage system.

You cannot send backups to one remote system and reports to another. However, you can send either backups or reports to a remote system and store the other on the Firewall Management Center.

Tip After configuring and selecting remote storage, you can switch back to local storage only if you have not increased the connection database limit. Refer to Database, on page 67.

Configure local storage

Configure local storage to ensure that all data is stored directly on the device, without utilizing any remote storage solutions. This task helps you set up local storage for your system, which may be required for environments where remote storage is not available or desired.

This is relevant in scenarios where remote storage is not configured or not required for your deployment. Perform this configuration if your organization’s policy or infrastructure requires local-only storage.

Before you begin

Follow these steps to configure local storage.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Choose **Remote Storage Device**.

Step 3 Choose **Local (No Remote Storage)** from the **Storage Type** drop-down list.

Step 4 Click **Save**.

When you complete this task, local storage is configured and all data is stored on the device without using a remote storage device.

Configure NFS for remote storage

Configure NFS for remote storage to allow the Secure Firewall Management Center to store backups and reports externally.

The Firewall Management Center can store backups and reports on an NFS mount. This task is relevant when you need to configure remote storage for backup and reporting purposes in Secure Firewall Management Center.

Before you begin

Ensure you have access to the Secure Firewall Management Center and the necessary credentials for the NFS storage system.

- Verify the NFS server is reachable from the management center.
- Obtain the IPv4 address or hostname and directory path for the NFS storage.

Follow these steps to configure NFS for remote storage.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **Remote Storage Device**.

Step 3 Choose **NFS** from the **Storage Type** drop-down list.

Step 4 Add the connection information:

- Enter the IPv4 address or hostname of the storage system in the **Host** field.
- Enter the path to your storage area in the **Directory** field.

Step 5 Optionally, check the **Use Advanced Options** check box and enter any required mount options in **Command Line Options**.

Note
You can use white space, commas, and colons in Command Line Options.

You can specify the version number of the NFS storage using this format:

vers=*version*

For example, to select NFSv4, enter:

vers=4.0

Note
Ensure you validate the values by clicking **Test** button (as mentioned in Step 7). This prevents duplicate or invalid entries.

Step 6 Under **System Usage**:

- Choose **Use for Backups** to store backups on the designated host.
- Choose **Use for Reports** to store reports on the designated host.
- Enter **Disk Space Threshold** for backup to remote storage. Default is 90%.

Step 7 Click **Test** to test the connection.

Step 8 Click **Save**.

After completing this task, the Secure Firewall Management Center will use the configured NFS mount for storing backups and reports. The connection will be tested and saved, ensuring reliable remote storage.

What to do next

Refer to troubleshooting procedures if you experience random latency in the NFS connection with the firewall device.

- Collect troubleshooting file before or after the issue from the device. You can generate the troubleshoot file from the web interface or using CLI commands. For information on how to generate the troubleshoot file, refer to Troubleshoot Firepower File Generation Procedures.
- Collect the incoming and exiting traffic PCAP records. For information on the procedure, refer to Packet capture, on page 460.
- Collect system-support trace data while NFS application fails using the following command in the device (CLISH mode):

```
> system support trace
```

- Collect snort counters twice during the failure using the **show snort counters** command to view the statistics for the Snort preprocessor connections. For information on this command, refer to show snort counters.

Configure SMB for remote storage

Configure SMB for remote storage to allow the Secure Firewall Management Center to store backups and reports on an external SMB mount.

The Firewall Management Center can store backups and reports on an SMB mount. Use this task to configure external storage for backup and reporting.

Before you begin

Ensure that your external remote storage system is functional and accessible from your Firewall Management Center:

- The system recognizes top-level SMB shares, not full file paths. You must use Windows to share the exact directory you want to use.
- Make sure the Windows user you will use to access the SMB share from the Firewall Management Center has ownership of and read/change access to the share location.
- For improved security, install SMB version 2.0 or later.

Follow these steps to configure SMB for remote storage.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **Remote Storage Device**.

Step 3 Choose **SMB** from the **Storage Type** drop-down list.

Step 4 Add the connection information:

- Enter the IPv4 address or hostname of the storage system in the **Host** field.
- Enter the share of your storage area in the **Share** field.
- Optionally, enter the domain name for the remote storage system in the **Domain** field.
- Enter the user name for the storage system in the **Username** field and the password for that user in the **Password** field.

Step 5 Optionally, check the **Use Advanced Options** check box and enter any required mount options in **Command Line Options**.

Note
You can use white space, commas, and colons in Command Line Options.

You can specify the version number of the SMB storage using this format:

vers=*version*

If SMB encryption is enabled for a file server, only SMB version 3.0 clients are allowed to access the file server. In this case, enter:

vers=3.0

Note
Ensure you validate the values by clicking **Test** button (as mentioned in Step 7). This prevents duplicate or invalid entries.

Step 6 Under **System Usage**:

- Choose **Use for Backups** to store backups on the designated host.
- Choose **Use for Reports** to store reports on the designated host.

Step 7 To test the settings, click **Test**.

Step 8 Click **Save**.

After completing this task, the Secure Firewall Management Center will use the configured SMB remote storage device for backups and reports. The system will be able to securely access and store data externally.

What to do next

Verify that backups and reports are successfully stored on the SMB remote storage device. Monitor storage usage and access logs to ensure ongoing reliability and security.

Configure SSH for remote storage

Configure SSH as a remote storage device so that the Secure Firewall Management Center can store backups and reports on an SSHFS mount.

The Firewall Management Center can store backups and reports on an SSHFS mount. Use this task when you need to configure remote storage for backups or reports using SSHFS.

Before you begin

Ensure you have access to the Secure Firewall Management Center and the remote storage system supports SSHFS. Gather the IP address or host name, directory path, and user credentials for the remote storage system.

Follow these steps to configure SSH for remote storage:

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **Remote Storage Device**.

Step 3 Choose **SSH** from the **Storage Type** drop-down list.

Step 4 Add the connection information:

- Enter the IP address or host name of the storage system in the **Host** field.
- Enter the path to your storage area in the **Directory** field.
- Enter the storage system's user name in the **Username** field and the password for that user in the **Password** field. To specify a network domain as part of the connection user name, precede the user name with the domain followed by a forward slash (/).
- To use SSH keys, copy the content of the **SSH Public Key** field and place it in your authorized_keys file.

Step 5 Optionally, check the **Use Advanced Options** check box and enter any required mount options in **Command Line Options**.

Step 6 Under System Usage:

- Choose **Use for Backups** to store backups on the designated host.
- Choose **Use for Reports** to store reports on the designated host.

Step 7 Click **Test** to test the connection.

Step 8 Click **Save**.

After you complete these steps, Secure Firewall Management Center uses the SSHFS mount to store backups and reports. The system tests and saves the connection, making remote storage available for your system data.

What to do next

Verify that your backups and reports are stored on the remote SSHFS mount as expected. Monitor storage usage and the connection status regularly to ensure continued operation.

SNMP

Simple Network Management Protocol (SNMP) is a network management protocol that:

- enables polling to access the standard management information base (MIB) for system details,
- supports SNMP protocol versions 1, 2, and 3, and
- provides information such as contact, administrative, location, and service information; IP addressing and routing information; and transmission protocol usage statistics.

SNMP polling reference information

Your network management system can access information in the MIBs through SNMP polling without sending SNMP traps.

- SNMPv2 only supports read-only communities.
- SNMPv3 only supports read-only users and supports encryption with AES128.

Note When selecting SNMP versions, SNMPv2 only supports read-only communities and SNMPv3 only supports read-only users. SNMPv3 also supports encryption with AES128.

Configure SNMP polling

Configure SNMP polling to allow network management systems to monitor and retrieve information from your deployment using SNMP. This task enables you to set up SNMP access and select the appropriate SNMP version and credentials for secure monitoring.

SNMP polling is used by network management systems to collect information from your deployment. Proper configuration ensures only authorized hosts can access SNMP data.

Restrict SNMP access to trusted hosts and use SNMPv3 with strong passwords for enhanced security. SNMP MIB contains sensitive information. Improper access to this data can pose security risks.

Before you begin

Add SNMP access for each computer you plan to use to poll the system. Refer to Configure an access list, on page 51.

- Restrict your access list for SNMP access to the specific hosts that will be used to poll for the MIB.
- Use SNMPv3 and strong passwords for network management access.

Follow these steps to configure SNMP polling.

Procedure

Step 1 Choose **System** (⚙) > **Configuration** and click **SNMP**.

Step 2 From the **SNMP Version** drop-down list, choose the SNMP version you want to use:

- **Version 1** or **Version 2**: Enter a read-only SNMP community name in the **Community String** field, then skip to the end of the procedure.

  Note
  Do not include special characters (< > / % # & ? ', etc.) in the SNMP community string name.

- **Version 3**: Click **Add User** to display the user definition page. SNMPv3 only supports read-only users and encryption with AES128.

Step 3 Enter a **Username**.

Step 4 Choose the protocol you want to use for authentication from the **Authentication Protocol** drop-down list.

Step 5 Enter the password required for authentication with the SNMP server in the **Authentication Password** field. Re-enter the authentication password in the **Verify Password** field.

Step 6 Choose the privacy protocol you want to use from the **Privacy Protocol** list, or choose **None** to not use a privacy protocol.

Step 7 Enter the SNMP privacy key required by the SNMP server in the **Privacy Password** field.

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

Step 9 Click **Add**. Click **Save**.

SNMP polling is configured. Your network management system can now securely poll the device using the selected SNMP version and credentials.

Session timeouts

A session timeout is a security feature that

- automatically ends idle login sessions,
- reduces security risks from unattended sessions, and
- allows configuration of idle time thresholds for user logins.

Session timeout configuration options

You can configure session timeouts to control how long a user's login session remains active during periods of inactivity.

- You can exempt specific web interface users from timeout in secure, passive monitoring scenarios.
- Users with the Administrator role cannot be exempt from session timeouts, because their accounts pose an increased risk if compromised.

Note Unattended login sessions may be security risks.

Configure session timeouts

Configure session timeout settings to control how long user sessions remain active in the web interface and CLI before automatic logout occurs.

Procedure

Step 1 Choose **System** (⚙) > **Configuration** > **Session Timeout**.

Step 2 Configure session timeouts:

- Web interface (Firewall Management Center only): Configure the **Browser Session Timeout (Minutes)**. The default value is `60`; the maximum value is `1440` (24 hours).

  To exempt users from this session timeout, refer to Add or edit an internal user, on page 148.

- CLI: Configure the **CLI Timeout (Minutes)** field. The default value is `0`; the maximum value is `1440` (24 hours).

Step 3 Click **Save**.

The session timeout settings are saved and will be applied to new user sessions. Existing sessions continue until they reach their timeout or users log out manually.

Configure time settings

Configure time settings so that displayed times reflect your local time zone preference, while system times are stored in UTC. This ensures consistent time management across the appliance.

- Maintain accurate time records for logs and events.
- Prevent unsupported system states by keeping the system clock set to UTC.

Time settings appear on most pages in local time, based on the time zone you set in User Preferences (the default is America/New York). The appliance stores all times using UTC.

The Time Zone function in User Preferences assumes that the default system clock is set to UTC. If you change the system time from UTC, support for the system is not guaranteed. You must reimage the device to recover from this state.

- Do not attempt to change the system time from UTC.
- Set your preferred time zone in User Preferences to display local time accurately.

Before you begin

Ensure you have access to User Preferences and the System Configuration menu in Secure Firewall Management Center. Your account must have permission to view and modify time settings.

Follow these steps to configure time settings.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **Time**.

The current time is displayed using the time zone specified for your account in User Preferences.

If your appliance uses an NTP server: For information about the table entries, refer to View NTP server status, on page 122.

After completing these steps, the appliance displays times in your preferred local time zone while maintaining system time in UTC. This configuration ensures that event and log timestamps remain accurate.

What to do next

Review time zone settings periodically to ensure they match your requirements. If you need to change your time zone preference, update it in User Preferences.

View NTP server status

If you are synchronizing time from an NTP server, you can view connection status on the **Time** page (choose **System** (⚙) > **Configuration**).

This table describes the fields related to NTP server status.

Table 7: NTP status

Column	Description
NTP Server	The IP address or name of the configured NTP server.

Column	Description
Status	The status of the NTP server time synchronization: • **Being Used** indicates that the appliance is synchronized with the NTP server. • **Available** indicates that the NTP server is available for use, but time is not yet synchronized. • **Not Available** indicates that the NTP server is in your configuration, but the NTP daemon is unable to use it. • **Pending** indicates that the NTP server is new or the NTP daemon was recently restarted. Over time, its value should change to **Being Used**, **Available**, or **Not Available**. • **Unknown** indicates that the status of the NTP server is unknown.
Authentication	The authentication status for communication between the Firewall Management Center and the NTP server: • **none** indicates no authentication is configured. • **bad** indicates authentication is configured but has failed. • **ok** indicates authentication is successful. If authentication has been configured, the system displays the key number and key type (SHA-1, MD5, or AES-128 CMAC) following the status value. For example: **bad, key 2, MD5**.
Offset	The number of milliseconds of difference between the time on the appliance and the configured NTP server. Negative values indicate that the appliance is behind the NTP server, and positive values indicate that it is ahead.
Last Update	The number of seconds that have elapsed since the time was last synchronized with the NTP server. The NTP daemon automatically adjusts the synchronization times based on a number of conditions. For example, if you see larger update times such as 300 seconds, that indicates that the time is relatively stable and the NTP daemon has determined that it does not need to use a lower update increment.

Time synchronization

A time synchronization is a system operation that

- ensures consistent system time across the Secure Firewall Management Center and managed devices,
- uses Network Time Protocol (NTP) servers to synchronize time securely, and
- supports authentication methods such as MD5, SHA-1, or AES-128 CMAC for secure communications.

Time synchronization reference information

Synchronizing system time is critical for proper operation and security of the Secure Firewall Management Center and its managed devices. Key points about time synchronization:

- Specify NTP servers during initial configuration to establish secure time synchronization.
- The Firewall Management Center supports secure communications with NTP servers using MD5, SHA-1, or AES-128 CMAC symmetric key authentication.
- Configuring connections solely with authenticated NTP servers improves security when environments use a mix of authentication methods or during migration.

You can synchronize time using one of these options:

1. Configure the Firewall Management Center to synchronize with an NTP server, or with multiple NTP servers.
2. Set the time manually and configure the Firewall Management Center to serve as an NTP server for managed devices.

Time synchronization options and descriptions:

- Recommended: Synchronize time on the Firewall Management Center with an NTP server, on page 125 Configure the Firewall Management Center to synchronize with NTP servers and link managed devices to the same NTP servers.
- Alternatively, set the time manually, configure the Firewall Management Center as an NTP server, and link managed devices to that server for synchronization.



Note If you specified an NTP server for the Firewall Management Center during initial configuration, the connection with that NTP server is not secured. You must edit the configuration for that connection to specify MD5, SHA-1, or AES-128 CMAC keys.



Caution Unintended consequences can occur when time is not synchronized between the Firewall Management Center and managed devices.

Time synchronization example

For example, configuring the Firewall Management Center to synchronize with an NTP server ensures all managed devices maintain consistent time, which is critical for logging and security events.

Time synchronization counter-example

If time is not synchronized, logs and security events may be inconsistent, leading to troubleshooting difficulties and potential security risks.

Time synchronization analogy

Synchronizing system time is like setting all clocks in a building to the same time, ensuring everyone operates on a consistent schedule.

Synchronize time on the Firewall Management Center with an NTP server

Synchronizing the Firewall Management Center with an NTP server ensures accurate timekeeping. Accurate timekeeping supports system logs, event correlation, and network operations. Proper time synchronization is critical for system operation and troubleshooting.

Ensure time synchronization among all components of your system because it is critically important.

Use an NTP server on your network to ensure proper time synchronization between Firewall Management Center and all managed devices.

The Firewall Management Center supports NTPv4.

You must have Admin or Network Admin privileges to do this procedure.

Before you begin

Note these pointers:

- If your Firewall Management Center and managed devices cannot access a network NTP server, do not use this procedure. Instead, refer to Synchronize time without access to a network NTP server, on page 126.
- Do not specify an untrusted NTP server.
- If you plan to establish a secure connection with an NTP server (recommended for system security), obtain an SHA-1, MD5, or AES-128 CMAC key number and value configured on that NTP server.
- Connections to NTP servers do not use configured proxy settings.
- Firepower 4100 Series devices and Firepower 9300 devices cannot use this procedure to set the system time. Instead, configure those devices to use the same NTP server(s) that you configure using this procedure. For instructions, refer to the documentation for your hardware model.

Caution If the Firewall Management Center is rebooted and your DHCP server sets an NTP server record different than the one you specify here, the DHCP-provided NTP server will be used instead. To avoid this situation, configure your DHCP server to use the same NTP server.

Follow these steps to synchronize time on the Firewall Management Center with an NTP server.

Procedure

Step 1 Choose **System** (⚙) > **Configuration** and click **Time Synchronization**.

Step 2 If **Serve Time via NTP** is **Enabled**, choose **Disabled** to disable the Firewall Management Center as an NTP server.

Step 3 For the **Set My Clock** option, choose **Via NTP**. Click **Add**.

Step 4 In the **Add NTP Server** dialog box, enter the host name or the IPv4 or IPv6 address of an NTP server.

Step 5 (Optional) To secure communication between your Firewall Management Center and the NTP server:

a) Select **MD5**, **SHA-1** or **AES-128 CMAC** from the **Key Type** drop-down list.

b) Enter the corresponding MD5, SHA-1, or AES-128 CMAC **Key Number** and **Key Value** from the specified NTP server.

Step 6 Click **Add**.

Step 7 Configure at least three NTP servers. When only two NTP servers are configured, the offset difference between them becomes high. As a result, the Firewall Management Center uses Local Time.

To add more NTP servers, repeat Steps 5 through 8.

Step 8 (Optional) To force the Firewall Management Center to use only an NTP server that successfully authenticates, check the **Use the authenticated NTP server only** check box.

Step 9 Click **Save**.

The Firewall Management Center synchronizes its system time with the specified NTP server(s), ensuring accurate and consistent time across your deployment.

What to do next

Set managed devices to synchronize with the same NTP server or servers:

- Configure device platform settings: *Configure NTP Time Synchronization for Threat Defense* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

  Even if you force the Firewall Management Center to make a secure connection with an NTP server (**Use the authenticated NTP server only**), device connections to that server do not use authentication.

- Deploy configuration changes; see the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Synchronize time without access to a network NTP server

This task enables you to configure a physical-hardware Firewall Management Center as an NTP server for time synchronization when a network NTP server is not accessible. Use this procedure only if no other NTP server is available in your environment.

If your devices cannot directly reach the network NTP server, or your organization does not have a network NTP server, a physical-hardware Firewall Management Center can serve as an NTP server.

- Do not use this procedure unless you have no other NTP server. Instead, use the procedure in Synchronize time on the Firewall Management Center with an NTP server, on page 125.
- Do not use a virtual Firewall Management Center as an NTP server.

To change the time manually **after** configuring the Firewall Management Center as an NTP server, you must disable the NTP option, change the time manually, and then re-enable the NTP option.

Before you begin

Make sure you have access to a physical-hardware Firewall Management Center, and verify that no other NTP server is available. Do not use a virtual management center for this procedure.

Follow these steps to synchronize time without access to a network NTP server:

Procedure

Step 1 Manually set the system time on the Firewall Management Center:

a) Choose **System** (⚙) > **Configuration**.
b) Click **Time Synchronization**.
c) If **Serve Time via NTP** is **Enabled**, choose **Disabled**.
d) Click **Save**.
e) For **Set My Clock**, choose **Manually in Local Configuration**.
f) Click **Save**.
g) In the navigation panel at the left side of the screen, click **Time**.
h) Use the **Set Time** drop-down lists to set the time.

Note
When you change the time on the management center by more than 2 hours, you must reboot the device as soon as possible, for example in a maintenance window, to avoid any malfunction.

i) If the time zone displayed is not UTC, click it and set the time zone to **UTC**.
j) Click **Save**.
k) Click **Done**.
l) Click **Apply**.

Step 2 Set the Firewall Management Center to serve as an NTP server:

a) In the navigation panel at the left side of the screen, click **Time Synchronization**.
b) For **Serve Time via NTP**, choose **Enabled**.
c) Click **Save**.

Step 3 Set managed devices to synchronize with the Firewall Management Center NTP server:

a) In the Time Synchronization settings for the platform settings policy assigned to your managed devices, set the clock to synchronize **Via NTP from Management Center**.
b) Deploy the change to managed devices.

For Firewall Threat Defense devices, refer to *Configure NTP Time Synchronization for Threat Defense* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

After completing this task, your devices will synchronize their time with the physical-hardware management center acting as an NTP server, ensuring accurate time settings even without access to a network NTP server.

Restrictions for changing time synchronization settings

Your Firewall Management Center and its managed devices are heavily dependent on accurate time. The system clock maintains system time. The system clock is set to Universal Coordinated Time (UTC), which is the primary time standard by which the world regulates clocks and time.

- Do not attempt to change the system time. Changing the system time zone from UTC is not supported, and doing so will require you to reimage the device to recover from an unsupported state.

- If you configure the Firewall Management Center to serve time using NTP, and then later disable it, the NTP service on managed devices still attempts to synchronize time with the Firewall Management Center. You must update and redeploy any applicable platform settings policies to establish a new time source.
- To change the time manually after configuring the Firewall Management Center as an NTP server, disable the NTP option, change the time manually, and then re-enable the NTP option.

UCAPL/CC compliance

UCAPL/CC compliance is a security requirement that

- mandates the use of equipment and software that meet specific security standards,
- is established by the U.S. Department of Defense, and
- is recognized by global certification organizations.

Security certification standards

Organizations must ensure their equipment and software comply with security standards for regulatory and operational requirements.

For more information about this setting, refer to Security certifications compliance modes, on page 347.

Configure user settings

Global user configuration settings affect all users on the Firewall Management Center. Configure these settings on the User Configuration page (**System** (⚙) > **Configuration** > **User Configuration**).

- **Password Reuse Limit**: The number of passwords in a user's most recent history that cannot be reused. This limit applies to web interface access for all users. For the `admin` user, this also applies to CLI access. The system maintains separate password lists for web interface and CLI access. If the limit is set to zero (the default), password reuse is unrestricted. Refer to Set a password reuse limit, on page 129.
- **Track Successful Logins**: The number of days that the system tracks successful logins to the Firewall Management Center, per user, per access method (web interface or CLI). When users log in, the system displays their successful login count for the interface being used. When **Track Successful Logins** is set to zero (the default), the system does not track or report successful login activity. Refer to Track successful logins, on page 130.
- **Max Number of Login Failures**: The number of times in a row that users can enter incorrect web interface login credentials before the system temporarily blocks the account from access for a configurable time period. If a user continues to attempt logins while the temporary lockout is in force, the system will:
  - The system refuses access for that account (even with a valid password) without informing the user that a temporary lockout is in force.
  - The system continues to increment the failed login count for that account with each login attempt.
  - If the user exceeds the **Maximum Number of Failed Logins** configured for that account on the individual User Configuration page, the account is locked out until an admin user reactivates it.

- **Set Time in Minutes to Temporarily Lockout Users**: The duration in minutes for a temporary web interface user lockout if **Max Number of Failed Logins** is non-zero.
- **Max Concurrent Sessions Allowed**: Maximum sessions for users: The number of sessions of a particular type (read-only or read/write) that can be open at the same time. The session type depends on the user's assigned roles. A user assigned only read-only roles has sessions counted toward the **(Read Only)** session limit. If the user has any write privileges, sessions count toward the **Read/Write** session limit. If a user has any roles with write privileges, the session is counted toward the **Read/Write** session limit. For example, if a user is assigned the Admin role and the **Maximum sessions for users with Read/Write privileges/CLI users** is set to 5, the user will not be allowed to log in if there are already five other users logged in that have read/write privileges.



Note Predefined user roles and custom user roles that the system considers read-only for the purposes of concurrent session limits, are labeled with **(Read Only)** in the role name on the **System** (⚙) > **Users** > **Users** and the **System** (⚙) > **Users** > **User Roles**. If a user role does not contain **(Read Only)** in the role name, the system considers the role to be read/write. The system automatically applies **(Read Only)** to roles that meet the required criteria. You cannot make a role read-only by adding that text string manually to the role name.

For each type of session, you can set a maximum limit ranging from 1 to 1024. When **Max Concurrent Sessions Allowed** is set to zero (the default), concurrent sessions are unlimited.

If you set a more restrictive concurrent session limit, the system keeps existing sessions open but blocks new sessions that exceed the specified limit.

Set a password reuse limit

Set a password reuse limit to enhance account security by preventing users from reusing previous passwords.

If you enable the **Password Reuse Limit**, the system keeps encrypted *password histories* for Firewall Management Center users. Users cannot reuse passwords in their histories. You can specify the number of stored passwords for each user using each access method, such as the web interface or CLI . A user's current password counts towards this number. If you lower the limit, the system deletes older passwords from the history. Increasing the limit does not restore deleted passwords.

Before you begin

Follow these steps to set a password reuse limit.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **User Configuration**.

Step 3 Set the **Password Reuse Limit** to the number of passwords you want to maintain in the history (maximum 256).

To disable password reuse checking, enter 0.

Step 4 Click **Save**.

After completing these steps, the system enforces the specified password reuse limit, preventing users from reusing previous passwords according to your configuration.

Track successful logins

Enable tracking of successful logins for each user to monitor login activity over a specified period. Helps administrators review user access patterns and security events.

Tracking successful logins shows the login count when users access the web interface or CLI. When you adjust the tracking period, the retention of login records changes. Lowering the number of days deletes records of older logins from the system, and increasing the limit does not restore the count from those days. As a result, the reported number of successful logins may be temporarily lower than the actual number.

Before you begin

Follow these steps to track successful logins for users:

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **User Configuration**.

Step 3 Set **Track Successful Login Days** to the number of days to track successful logins (maximum 365).

To disable login tracking, enter 0.

Step 4 Click **Save**.

After completing this task, the system tracks and displays the number of successful logins for each user based on the configured retention period.

Enable temporary lockouts

Enable temporary lockouts to protect your system from repeated unauthorized login attempts by automatically locking out users after a specified number of failed logins. Prevent brute-force attacks by limiting consecutive failed login attempts.

Enable the temporary timed lockout feature by specifying the number of failed login attempts in a row that the system allows before the lockout goes into effect.

Before you begin

Ensure you have access to the system configuration interface. Follow these steps to enable temporary lockouts.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **User Configuration**.

Step 3 Set the **Max Number of Login Failures** to the maximum number of consecutive failed login attempts before the user is temporarily locked out.

To disable the temporary lockout, enter zero.

Step 4 Set the **Time in Minutes to Temporarily Lockout Users** to the number of minutes to lock out users who have triggered a temporary lockout.

When this value is zero, users do not have to wait to retry to log in, even if the **Max Number of Login Failures** is non-zero.

Step 5 Click **Save**.

After completing this task, users who exceed the allowed number of failed login attempts will be temporarily locked out for the configured duration, enhancing system security.

Set the maximum number of concurrent sessions

Set limits on the number of concurrent sessions for users based on their assigned roles to control system access and resource usage.

You can specify the maximum number of sessions of a particular type (read-only or read/write) that can be open at the same time. The type of session is determined by the roles assigned to a user.

If a user is assigned only read-only roles, that user's session is counted toward the **Read Only** session limit. If a user has any roles with write privileges, the session is counted toward the **Read/Write** session limit.

Before you begin

Ensure you have access to the system configuration interface.

Follow these steps to set the maximum number of concurrent sessions.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **User Configuration**.

Step 3 For each type of session, set the **Max Concurrent Sessions Allowed** to the maximum number of sessions that can be open at the same time.

To apply no limits on concurrent sessions per users, enter zero.

Note
If you set a more restrictive concurrent session limit, the system keeps current sessions open but stops new sessions from being started when the limit is reached.

Step 4 Click **Save**.

The system enforces the specified maximum number of concurrent sessions for each session type, preventing new sessions from exceeding the set limits.

VMware tools

A VMware Tools suite is a collection of utilities that

- enhances performance for virtual machines,
- enables full use of convenient features in VMware products, and
- supports plugins for Secure Firewall virtual appliances running on VMware.

Plugin support for VMware tools

VMware Tools supports these plugins for virtual appliances running on VMware:

- guestInfo
- powerOps
- timeSync
- vmbackup

You can also enable VMware Tools on all supported versions of ESXi.

For information on the full functionality of VMware Tools, refer to the VMware website (http://www.vmware.com/).

Enable VMware Tools on the Secure Firewall Management Center for VMware

Enable VMware Tools on the Secure Firewall Management Center for VMware to improve integration and management of the virtual appliance.

This task is relevant when you want to enhance the performance, monitoring, and manageability of your Secure Firewall Management Center running on VMware platforms.

Before you begin

Follow these steps to enable VMware Tools on the Secure Firewall Management Center for VMware:

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **VMware Tools**.

Step 3 Click **Enable VMware Tools**.

Step 4 Click **Save**.

VMware Tools is enabled on the Secure Firewall Management Center for VMware. This provides improved monitoring and management of the virtual appliance.

Vulnerability mapping

A vulnerability mapping is a security feature that

- automatically maps vulnerabilities to a host IP address for any application protocol traffic received or sent from that address,
- requires the server to have an application ID in the discovery event database and the packet header to include a vendor and version, and
- allows configuration for associating vulnerabilities with server traffic for servers lacking vendor or version information.

Configuration options for vendor and versionless servers

You can configure whether the system associates vulnerabilities with server traffic for servers that do not include vendor or version information in their packets.

SMTP server vulnerability mapping example

For example, a host serves SMTP traffic that does not have a vendor or version in the header. If you enable the SMTP server on the Vulnerability Mapping page of a system configuration, then save that configuration to the Firewall Management Center managing the device that detects the traffic, all vulnerabilities associated with SMTP servers are added to the host profile for the host.

Application protocol detector exclusion

Although detectors collect server information and add it to host profiles, the application protocol detectors are not used for vulnerability mapping. You cannot specify a vendor or version for a custom application protocol detector, and you also cannot select the server for vulnerability mapping.

Mapping vulnerabilities for servers

This procedure requires any Smart License or the Protection classic license.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Choose **Vulnerability Mapping**.

Step 3 Select from these options:

- Clear the check box for a server to prevent its vulnerabilities from being mapped to hosts that receive application protocol traffic without vendor or version information.
- Check the box for a server to map its vulnerabilities to hosts that receive application protocol traffic without vendor or version information.

Tip
You can check or clear all check boxes at once using the check box next to **Enabled**.

Step 4 Click **Save**.

Configure web analytics

Configure web analytics to manage the collection of non-personally-identifiable usage data by Cisco from your appliances.

By default, in order to improve firewall products, Cisco collects non-personally-identifiable usage data, including but not limited to page interactions, browser versions, product versions, user location, and management IP addresses or hostnames of your Firewall Management Center appliances.

Data collection begins after you accept the End User License Agreement. If you do not want Cisco to continue to collect this data, you can opt out using this procedure.

Before you begin

Ensure you have access to the system configuration menu in your appliance interface. Follow these steps to configure web analytics.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.
Step 2 Click **Web Analytics**.
Step 3 Make your choice and click **Save**.

After completing these steps, your web analytics data sharing preference is updated. Cisco will collect or stop collecting usage data according to your selection.

What to do next

(Optional) Determine whether to share data through the Configure Cisco Success Network Enrollment.

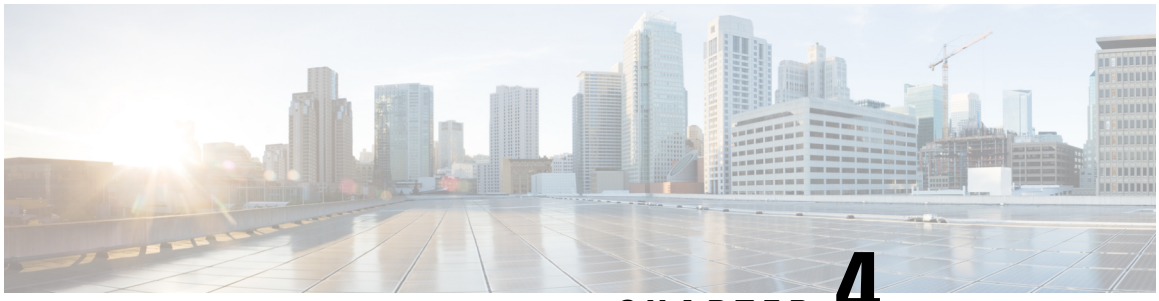
History for system configuration

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Access control performance improvements (object optimization).	7.2.4 7.4.0	Any	**Upgrade impact. First deployment after Firewall Management Center upgrade to 7.2.4–7.2.5 or 7.4.0 can take a long time and increase CPU use on managed devices.** Access control object optimization improves performance and consumes fewer device resources when you have access control rules with overlapping networks. The optimizations occur on the *managed device* on the first deploy after the feature is enabled on the Firewall Management Center (including if it is enabled by an upgrade). If you have a high number of rules, the system can take several minutes to an hour to evaluate your policies and perform object optimization. During this time, you may also see higher CPU use on your devices. A similar thing occurs on the first deploy after the feature is disabled (including if it is disabled by upgrade). After this feature is enabled or disabled, we recommend you deploy when it will have the least impact, such as a maintenance window or a low-traffic time. New/modified screens (requires Version 7.2.6): **System**(⚙) > **Configuration** > **Access Control Preferences** > **Object-group optimization**.
French language option.	7.2	Any	You can now switch the management center web interface to French. New/modified screens: **System** > **Configuration** > **Language**.
Exempt most connection events from event rate limits.	7.0	Any	Setting the **Maximum Connection Events** value for the Connection Database to zero now exempts low priority connection events from counting towards the flow rate limit for your FMC hardware. Previously, setting this value to zero applied only to event storage, and did not affect the flow rate limit. New/modified screens: **System** > **Configuration** > **Database**. Supported platforms: Hardware FMCs.
Support for AES-128 CMAC authentication for NTP servers.	7.0	Any	Connections between the FMC and NTP servers can be secured with AES-128 CMAC keys as well as previously-supported MD5 and SHA-1 keys. New/modified screens: **System** > **Configuration** > **Time Synchronization**.
Subject Alternative Name (SAN).	6.6	Any	When creating an HTTPS certificate for the FMC, you can specify SAN fields. Cisco recommends using a subject alternative name (SAN) if the certificate secures multiple domain names or IP addresses. For more information about SAN, refer to RFC 5280, section 4.2.1.6. New/modified screens: **System** > **Configuration** > **HTTPS Certificate**.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
HTTPS certificates.	6.6	Any	The default HTTPS server certificate provided with the system now expires in 800 days. If your appliance uses a default certificate that was generated before you upgraded to Version 6.6, the certificate life span varies depending on the Firepower version being used when the certificate was generated. Refer to Default HTTPS server certificates, on page 74 for more information. Supported platforms: Hardware FMCs.
Secure NTP.	6.5	Any	The FMC supports secure communications with NTP servers using SHA1 or MD5 symmetric key authentication. New/modified screens: **System** > **Configuration** > **Time Synchronization**.
Web analytics.	6.5	Any	Web analytics data collection begins after you accept the EULA. As before, you can opt not to continue to share data. Refer to Configure web analytics, on page 134.
Automatic CLI access for the FMC.	6.5	Any	When you use SSH to log into the FMC, you automatically access the CLI. Although strongly discouraged, you can then use the CLI `expert` command to access the Linux shell. **Note** This feature deprecates the Version 6.3 ability to enable and disable CLI access for the FMC. As a consequence of deprecating this option, the virtual FMC no longer displays the **System** > **Configuration** > **Console Configuration** page, which still appears on physical FMCs.
Configurable session limits for read-only and read/write access.	6.5	Any	Added the **Max Concurrent Sessions Allowed** setting. This setting allows the administrator to specify the maximum number of sessions of a particular type (read-only or read/write) that can be open at the same time. **Note** Predefined user roles and custom user roles that the system considers read-only for the purposes of concurrent session limits, are labeled with **(Read Only)** in the role name on **System** > **Users** > **Users** and **System** > **Users** > **User Roles**. If a user role does not contain **(Read Only)** in the role name, the system considers the role to be read/write. New/modified screens: • **System** > **Configuration** > **User Configuration**. • **System** > **Users** > **User Roles**.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Ability to disable Duplicate Address Detection (DAD) on management interfaces.	6.4	Any	When you enable IPv6, you can disable DAD. Disabling DAD can prevent denial of service attacks that exploit this protocol. If you disable this setting, manually verify that the interface is not using an address that is already assigned. New/modified screens: **System** > **Configuration** > **Management Interfaces** > **Interfaces** > **Edit Interface** > **IPv6 DAD**. Supported platforms: FMC
Ability to disable ICMPv6 Echo Reply and Destination Unreachable messages on management interfaces.	6.4	Any	When you enable IPv6, you can now disable ICMPv6 Echo Reply and Destination Unreachable messages. Disabling these packets can help protect against denial of service attacks. If you disable Echo Reply packets, you cannot use IPv6 ping to the device management interfaces for testing. New/modified screens: **System** > **Management Interfaces** > **ICMPv6**. New/modified commands: **configure network ipv6 destination-unreachable**, **configure network ipv6 echo-reply** Supported platforms: FMC (web interface only), FTD (CLI only)
Global User Configuration Settings.	6.3	Any	Added the **Track Successful Logins** setting. The system tracks the number of successful logins for each FMC account within a configured time period. When this feature is enabled, users see a message at login that reports how many times they have successfully logged in during that period. This applies to both the web interface and shell/CLI access. Added the **Password Reuse Limit** setting. The system can track the password history for each account for a configurable number of previous passwords. The system prevents all users from re-using passwords that appear in that history. (Applies to web interface as well as shell/CLI access.) Added the **Max Number of Login Failures** and **Set Time in Minutes to Temporarily Lockout Users** settings. These allow the administrator to limit the number of times in a row a user can enter incorrect web interface login credentials before the system temporarily blocks the account for a configurable period of time. New/modified screens: **System** > **Configuration** > **User Configuration**. Supported platforms: FMC
HTTPS Certificates.	6.3	Any	The default HTTPS server certificate provided with the system now expires in three years. If your appliance uses a certificate generated before upgrading to Version 6.3, it will expire 20 years from its creation date. The system now allows you to renew the default HTTPS server certificate. New/modified screens: **System** > **Configuration** > **HTTPS Certificate** > **Renew HTTPS Certificate**. Supported platforms: FMC

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Ability to enable and disable CLI access for the FMC.	6.3	Any	There is a new check box available to administrators in FMC web interface: **Enable CLI Access** on the **System** > **Configuration** > **Console Configuration**. • Checked: Logging into the FMC using SSH accesses the CLI. • Unchecked: Logging into FMC using SSH accesses the Linux shell. This is the default state for fresh Version 6.3 installations as well as upgrades to Version 6.3 from a previous release. Previous to Version 6.3, there was only one setting on the **Console Configuration** page, and it applied to physical devices only. So the **Console Configuration** page was not available on virtual FMCs. With the addition of this new option, the **Console Configuration** page now appears on virtual FMCs as well as physical. However, for virtual FMCs, this check box is the only thing that appears on the page. Supported platforms: FMC



CHAPTER 4

Users

The Firewall Management Center includes default **admin** accounts for web and CLI access. This chapter discusses how to create custom user accounts. Refer to Logging into the Management Center, on page 29 for detailed information about logging into the Firewall Management Center with a user account.

- CLI users, on page 139
- Guidelines and limitations for user accounts for Firewall Management Center, on page 146
- Requirements and prerequisites for user accounts for Firewall Management Center, on page 147
- Add or edit an internal user, on page 148
- External authentication configuration for Firewall Management Center, on page 150
- SAML Single Sign-On, on page 169
- Custom user roles for the web interface, on page 226
- Troubleshoot LDAP authentication connections, on page 232
- Configure User Preferences, on page 232
- History for Firewall Management Center user accounts, on page 241

CLI users

A CLI user is a user account that

- can be added as internal users or as external users on a LDAP or RADIUS server,
- is maintained separately on each managed device, and
- requires separate configuration for access to the Firewall Management Center versus managed devices.

User account separation

When you add a user to the Firewall Management Center, that user only has access to the Firewall Management Center; you cannot then use that username to log directly into a managed device. You must separately add a user on the managed device.

Internal and external users

Managed devices support two categories of users—internal and external—for user access control.

Internal users authenticate through a local database on the device. External users authenticate through LDAP or RADIUS servers when they are not found in the local database. Both user types enable secure access management for network devices.

User authentication types

Managed devices support these user authentication methods:

- Internal user: The device checks a local database for user authentication.
- External user: If the user is not present in the local database, the system queries an external LDAP or RADIUS authentication server.

Web interface and CLI access



Caution CLI users can access the Linux shell using the **expert** command. We strongly recommend that you do not use the Linux shell unless directed by Cisco TAC or provided with explicit instructions in the Firewall Management Center documentation. CLI users can obtain `sudoers` privileges in the Linux shell, which can present a security risk. For system security reasons, we strongly recommend that you:

- Restrict the list of external users with CLI access appropriately.
- Do not add users directly in the Linux shell. Use only the procedures in this chapter.

The Firewall Management Center provides access through a web interface, CLI (via console or SSH to the management interface), and Linux shell. For details about management UIs, refer to System user interfaces, on page 31.

- The Firewall Management Center supports two different internal **admin** users: one for the web interface and another for CLI access. The system initialization process synchronizes the passwords for these two **admin** accounts initially. However, they are managed separately, and the passwords may diverge after configuration. Refer to the *Getting Started Guide* for your model for more information on system initialization. To change the password for the web interface **admin**, use **System** (⚙) > **Users** > **Users**. To change the password for the CLI **admin**, use the Firewall Management Center CLI command **configure password**.
- Internal users added in the web interface have web interface access only.
- External users have web interface access, and you can optionally configure CLI access.
- SSO users have web interface access only.

User roles

CLI User Role

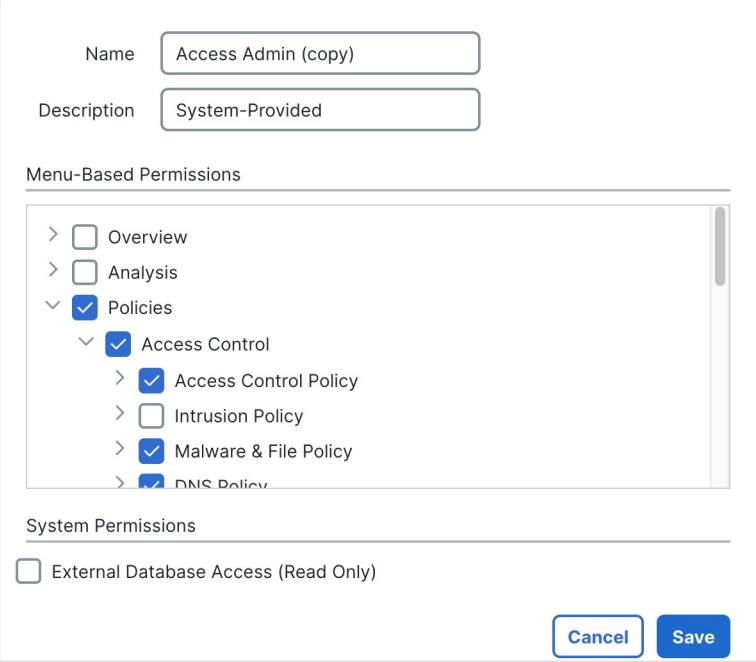
CLI external users on the Firewall Management Center do not have a user role; they can use all available commands.

Web Interface User Roles

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

To view the privileges assigned to predefined user roles, click **Copy** ( ) for a role to make a custom role based on the predefined role. You can then view all of the assigned privileges.

Figure 15: View User Role Privileges



The Firewall Management Center includes these predefined user roles:

Access Admin

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

Administrator

Administrators have access to everything in the product; their sessions present a higher security risk if compromised, therefore, you cannot make them exempt from login session timeouts.

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

Discovery Admin

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

External Database User (Read Only)

Provides read-only access to the database using an application that supports JDBC SSL connections. For the third-party application to authenticate to the appliance, you must enable database access in the system settings. On the web interface, External Database Users have access only to online help-related options

in the **Help** menu. This role's function does not involve the web interface. Access is provided only for support and password changes.

Intrusion Admin

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

Maintenance User

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

Network Admin

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

Security Analyst

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

Security Analyst (Read Only)

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

A user with this role can also:

- From the health monitor pages for specific devices, generate and download troubleshooting files.
- Under user preferences, set file download preferences.
- Under user preferences, set the default time window for event views (with the exception of the **Audit Log Time Window**).

Security Approver

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

Threat Intelligence Director (TID) User

Provides access to Threat Intelligence Director configurations in the **Intelligence** menu. Threat Intelligence Director (TID) Users can view and configure TID.

User passwords

This reference describes the rules and requirements for passwords on internal user accounts for the Firewall Management Center, with Lights-Out Management (LOM) enabled or disabled. Different password requirements apply to externally authenticated accounts and to systems with security certification compliance enabled. Refer to External authentication configuration for Firewall Management Center, on page 150 and Security certifications compliance, on page 347 for more information.

During Firewall Management Center initial configuration, the system requires the **admin** user to set the account password to comply with strong password requirements. For physical Firewall Management Centers, the system uses the strong password requirements with LOM enabled. For virtual Firewall Management

Centers, the system uses the strong password requirements with LOM not enabled. The system synchronizes the passwords for the web interface **admin** and CLI access **admin**. After initial configuration, the web interface **admin** can remove the strong password requirement. However, the CLI access **admin** must always comply with strong password requirements with LOM not enabled.

Table 8: Password requirements for internal user accounts

	LOM Not Enabled	LOM Enabled
Password Strength Checking On	Passwords must include: • At least eight characters or the number of characters configured for the user by the administrator, whichever is greater. • No more than two consecutively repeating characters • At least one lowercase letter. • At least one uppercase letter. • At least one digit. • At least one special character such as ! @ # * - _ +. The system checks passwords against a special dictionary containing many English words and other character strings that could be easily cracked with common password hacking techniques.	Passwords must include: • Between eight and twenty characters (On MC 1000, MC 2500, and MC 4500 the upper limit is 14 characters rather than 20.) • No more than two sequentially repeating characters • At least one lower case letter • At least one upper case letter • At least one digit • At least one special character such as ! @ # * - _ + The rules for special characters are different for each series of physical Firewall Management Centers. Choose special characters only from the list in the final bullet. Do not include the user name in the password. The system checks passwords against a special dictionary of English words and character strings that can be easily cracked with common password hacking techniques.

	LOM Not Enabled	LOM Enabled

Password Strength Checking Off	Passwords must include the minimum number of characters configured for the user by the administrator. (Refer to Add or edit an internal user, on page 148 for more information.)	Passwords must include:
		• Between 8 and 20 characters (On MC 1000, MC 2500, and MC 4500 the upper limit is 14 characters rather than 20.)
		• Characters from at least three of the four categories:
		  • Uppercase letters
		  • Lowercase letters
		  • Digits
		  • Special characters such as ! @ # * - _ +
		The rules for special characters vary between different series of physical Firewall Management Centers. Use only the special characters from the permitted list.
		Do not include the user name in the password.

Users and domains

A user account is an access entity that

- can be created in any domain where Administrator access is assigned
- is only visible in the domain in which it is created, and
- can have different privileges and roles in ancestor and descendant domains.

User roles and domain management

Users can have different privileges in each domain, and user roles can be assigned in both ancestor and descendant domains.

- Assign read-only privileges to a user in the Global domain.
- Assign Administrator privileges in a descendant domain.

Figure 16: User roles per domain



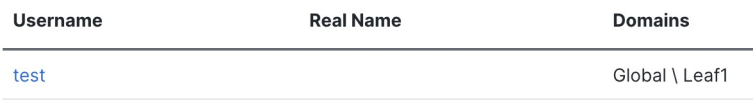
User visibility in domains

Users are only visible in the domain in which they are created. If a user is added in the current domain but assigned a user role in a subdomain, the user will only show on the current domain's **Users** page.

Figure 17: Global domain users



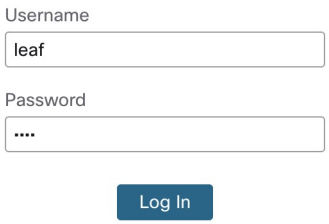
Figure 18: Subdomain users



User login procedures

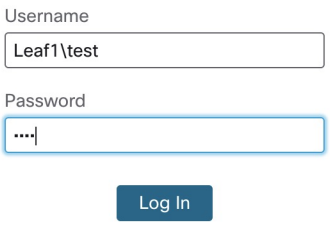
Users added from the Global domain log in with just their username, even if their roles are only in a subdomain.

Figure 19: User login when added in global



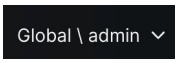
log into the Firewall Management Center with the subdomain or subdomains as part of the login name, depending on which domain their user was added from: *subdomain1\subdomain2\username*. You do not need to enter the **Global** parent domain.

Figure 20: Global domain user login



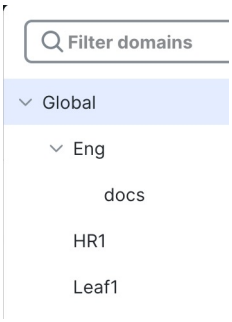
When you log in, you are placed in the domain where your username was added. For example, the admin user defaults to the Global domain.

Figure 21: User domain at login



After you log in, click the down arrow to change to a subdomain.

Figure 22: Change to a subdomain



User account examples in domains

For example, from the Global domain, you add user **leaf** and assign a role for **Leaf1**. Because the user was added from the Global domain, it is visible from the Global domain. If you change domains to **Leaf1**, you cannot view the user **leaf**, but you can view the user **test**, which was added directly from the Leaf1 subdomain.

Guidelines and limitations for user accounts for Firewall Management Center

The Firewall Management Center includes an **admin** user as a local user account for all forms of access. You cannot delete the **admin** user. The default initial password is **Admin123**. During initialization, the system requires you to change this password. It supports a maximum of 16 characters. For more information about system initialization, refer to the *Getting Started Guide* for your model.

Refer to these guidelines for default admin user and password management on the Firewall Management Center:

- There are no limits on password reuse.
- The system does not track successful logins.
- The system does not enforce a timed temporary lockout for users who enter incorrect login credentials.
- The system allows unlimited read-only and read or write sessions simultaneously.

You can change these settings for all users as a system configuration (**System** (⚙) > **Configuration** > **User Configuration**).For more information, refer to Configure user settings, on page 128.

Assign default access roles to users at initial setup according to the principles of least privilege. When a user logs in to the system for the first time, the system assigns the default access role to their account. Assign the minimum privilege required for users to log in to the system. For example, you can assign the Security Analyst (Read-Only) role as the default access role for common users. You can add administrators to a separate administrator group to give them full administrator rights. This guideline applies to all internal, external, and CAC users.

Note Assign the default access role using the principles of least privilege to prevent users from receiving unintended privilege levels on subsequent logins.

Refer to these best practices for assigning default access roles and privilege management:

- If a user with the default access role needs temporary elevated privileges, an administrator can assign a higher-privilege role to that user. The system revokes elevated privileges after 24 hours of inactivity. The user then returns to their default access role.
- To permanently assign a higher-privilege access role, such as System Admin, use the Group Controlled Access Roles method to provide admin access. This method ensures that the provided access role persists beyond 24 hours and users will have the correct privilege level as per the group assignment. For more information, refer to the Add an LDAP External Authentication Object for Management Center section.

Requirements and prerequisites for user accounts for Firewall Management Center

This reference lists the requirements and prerequisites for user accounts for Firewall Management Center, including supported models, domains, and user roles.

Model support

Firewall Management Center

Supported domains

- SSO configuration: Global only.
- All other features: Any.

User roles

- SSO configuration: Only users with the Admin role authenticated internally or by LDAP or RADIUS can configure SSO.
- All other features: Any user with the Admin role.
- Configure common access card authentication with LDAP, on page 167 also supports the Network Admin role.

Add or edit an internal user

This procedure allows you to create custom internal user accounts for the Firewall Management Center and modify existing user configurations including passwords, roles, and access permissions.

The **Users** tab in **System** (⚙) > **Users** displays both internal users that you add manually and external users who are added automatically when logging in with LDAP or RADIUS authentication. For external users, you can modify the user role on this screen if you assign a role with higher privileges; you cannot modify the password settings.

If you enable security certifications compliance or Lights-Out Management (LOM) on a device, different password restrictions apply. For more information on security certifications compliance, refer to Security certifications compliance, on page 347.

Note Avoid having multiple Admin users simultaneously creating new users on the Firewall Management Center, as this may cause an error resulting from a conflict in user database access.

Procedure

Step 1 Choose **System** (⚙) > **Users**.

Step 2 To create a new user:

a) Click **Create User**.

b) Enter a **User Name**.

The username must comply with these restrictions:

- Maximum 32 alphanumeric characters, plus hyphen (-) and underscore (_).
- Letters may be upper or lower case.
- Cannot include any punctuation or special characters other than period (.), hyphen (-), and underscore (_).

Step 3 To edit an existing user, click the **Edit** (✎) icon next to the user you want to edit.

a. **Real Name**: Enter descriptive information to identify the user or department to whom the account belongs.

b. (Optional) **Email Address**: Provide your email address. Cisco uses your email address to send product renewal information, newsletters about new releases, and other product updates.

c. The **Use External Authentication Method** check box is checked for users that were added automatically when they logged in with LDAP or RADIUS. You do not need to preconfigure external users, so you can ignore this field. For an external user, you can revert this user to an internal user by unchecking the check box.

d. Enter values in the **Password** and **Confirm Password** fields.

The values must conform to the password options you set for this user.

e. Set the **Maximum Number of Failed Logins**.

Enter an integer without spaces to specify the maximum number of times each user can try to log in after a failed login attempt before the account is locked. The default setting is 5 tries; use `0` to allow an unlimited number of failed logins. The **admin** account is exempt from being locked out after a maximum number of failed logins unless you enable security certification compliance.

f. Set the **Minimum Password Length**.

Enter an integer without spaces to specify the minimum required length, in characters, of a user's password. The default setting is `8`. A value of `0` indicates that no minimum length is required.

g. Set the **Days Until Password Expiration**.

Enter the number of days after which the user's password expires. The default setting is `0`, which indicates that the password never expires. If you change from the default, then the **Password Lifetime** column of the **Users** list indicates the days remaining on each user's password.

h. Set the **Days Before Password Expiration Warning**. Enter the number of warning days users have to change their password before their password actually expires. The default setting is `0` days.

i. Set these **Options**:

- **Force Password Reset on Login**: Forces users to change their passwords the next time they log in.
- **Check Password Strength**: Requires strong passwords. When password strength checking is enabled, passwords must comply with the strong password requirements described in User passwords, on page 142.
- **Exempt from Browser Session Timeout**: Exempts a user's login sessions from termination due to inactivity. Users with the Administrator role cannot be made exempt.

j. In the **User Role Configuration** area, assign the user roles. For more information about user roles, refer to Custom user roles for the web interface, on page 226.

For external users, if the user role is assigned through group membership (LDAP), or based on a user attribute (RADIUS), you cannot remove the minimum access rights. You can, however, assign additional rights. If the user role is the default user role that you set on the device, then you can modify the role in the user account without limitations. When you modify the user role, the **Authentication Method** column on the **Users** tab provides a status of **External - Locally Modified**.

The options that you view depend on whether the device is in a single domain or multidomain deployment.

- Single domain: Check the user roles you want to assign the user.
- Multidomain: In a multidomain deployment, you can create user accounts in any domain in which you have been assigned Administrator access. Users can have different privileges in each domain.

You can assign user roles in both ancestor and descendant domains. For example, you can assign read-only privileges to a user in the Global domain, but Administrator privileges in a descendant domain. For more information, including how to log in when a user is added in a subdomain, refer to Users and domains, on page 144. Refer to these steps:

1. Click **Add Domain**.
2. Choose a domain from the **Domain** drop-down list.
3. Check the user roles that you want to assign the user.
4. Click **Save**.

Step 4 (Optional, for physical Firewall Management Centers only) If you have assigned the user the Administrator role, the **Administrator Options** appear. You can select **Allow Lights-Out Management Access** to grant Lights-Out Management access to the user. Refer to #unique_198 for more information about Lights-Out Management.

Step 5 Click **Save**.

The system creates or updates the user account with the specified configuration settings. The new or modified user appears in the Users list, displaying the assigned roles and authentication method.

External authentication configuration for Firewall Management Center

An external authentication configuration is a security setup that

- adds one or more external authentication objects,
- enables users to authenticate through external identity providers, and
- replaces local credential authentication with external authentication.

External authentication objects for Firewall Management Center

An external authentication object is a security configuration element that

- enables the Firewall Management Center to verify user credentials with LDAP or RADIUS servers,
- supports multiple objects for web interface access, allowing users from any configured object to authenticate, and
- restricts CLI access to a single external authentication object, permitting authentication only through the first object in the list.

External authentication objects can be used by the Firewall Management Center and Firewall Threat Defense devices. You can either share the same object across different devices or create distinct objects for each device type.

Note The timeout range is different for the Firewall Threat Defense and the Firewall Management Center, so if you share an object, be sure not to exceed the Firewall Threat Defense's smaller timeout range (1 to 30 seconds for LDAP, and 1 to 300 seconds for RADIUS). If you set the timeout to a higher value, the Firewall Threat Defense external authentication configuration will not work.

For the Firewall Management Center, enable the external authentication objects directly on the **System** (⚙) > **Users** > **External Authentication** tab. This setting affects only Firewall Management Center usage. You do not need to enable it on this tab for managed device usage. For Firewall Threat Defense devices, enable the external authentication object in the platform settings that you deploy to the devices.

Define web interface users and CLI users separately in the external authentication object. For RADIUS, pre-configure the list of CLI usernames in the object. For LDAP, set up a filter on the LDAP server to match CLI users.

You cannot use an LDAP object for CLI access if that object is also configured for CAC authentication.

Note Users with CLI access can gain Linux shell access with the **expert** command. Linux shell users can obtain root privileges. This access creates a security risk. Make sure you:

- Restrict the list of users with CLI or Linux shell access.
- Do not create Linux shell users.

If you have five external authentication objects configured for web interface access, users from any of them can be authenticated to access the web interface. For CLI access, only the first external authentication object in the list is used for authentication.

You cannot use an LDAP object for CLI access if it is also configured for CAC authentication.

An external authentication object functions as a control point, permitting access only to users with valid credentials from specified sources.

LDAP

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

LDAP binding information

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

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

RADIUS

RADIUS is a network authentication protocol that

- provides centralized authentication and authorization for network access,
- operates using client-server architecture for user credential verification, and
- supports accounting functions to track user sessions and network usage.

Remote Authentication Dial In User Service (RADIUS) is an authentication protocol used to authenticate, authorize, and account for user access to network resources. You can create an authentication object for any RADIUS server that conforms to RFC 2865.

Secure Firewall devices support the use of SecurID tokens. When you configure authentication by a server using SecurID, users authenticated against that server append the SecurID token to the end of their SecurID PIN and use that as their password when they log in. You do not need to configure anything extra on the Secure Firewall device to support SecurID.

By default, RADIUS uses port 1812 for authentication and port 1813 for accounting.

If you change the RADIUS authentication port, the RADIUS accounting port changes accordingly. Ensure that the Firewall Management Center can connect to the RADIUS server on the new accounting port; otherwise, authentication delays may occur.

Although you cannot configure RADIUS accounting parameters, the Firewall Threat Defense uses port 1813 for RADIUS accounting on the same server used for authentication. If the RADIUS server is unreachable on port 1813, it can cause delays in logging in.

Add an LDAP external authentication object for the Firewall Management Center

Add an LDAP server to support external users for device management.

Before you begin

- You must specify DNS server(s) for domain name lookup on your device. Even if you specify an IP address for the LDAP server during this procedure, the server may return a URI for authentication that includes a hostname. A DNS lookup is required to resolve the hostname. Refer to Modify Firewall Management Center management interfaces, on page 96 to add DNS servers.
- If you are configuring an LDAP authentication object for use with CAC authentication, do not remove the CAC from your computer. You must have a CAC inserted at all times after enabling user certificates.

Procedure

Step 1 (Optional) Choose **System** (⚙) > **Users** and click the **External Authentication** tab. Click **Add icon** (+) **Add External Authentication Object** and set the **Authentication Method** to **LDAP**. Now, check the check box for **CAC** if you plan to use this authentication object for CAC authentication and authorization.

You must also follow the procedure in Configure common access card authentication with LDAP, on page 167 to fully configure CAC authentication and authorization. You cannot use this object for CLI users.

Step 2 Choose a **Server Type** from the drop-down list.

Tip
If you click **Set Defaults** , the device populates the **User Name Template** , **UI Access Attribute** , **CLI Access Attribute** , **Group Member Attribute** , and **Group Member URL Attribute** fields with default values for the server type.

Step 3 For the **Primary Server**, enter a **Host Name/IP Address**. If you are using a certificate to connect via TLS or SSL, the host name in the certificate must match the host name used in this field. In addition, IPv6 addresses are not supported for encrypted connections. Now, change the **Port** from the default and enter the **Backup Server** parameters and **LDAP-Specific Parameters**.

a) Enter the **Base DN** for the LDAP directory you want to access. For example, to authenticate names in the Security organization at the Example company, enter `ou=security,dc=example,dc=com` . Alternatively click **Fetch DNs** , and choose the appropriate base distinguished name from the drop-down list.

b) (Optional) Enter the **Base Filter** . For example, if the user objects in a directory tree have a `physicalDeliveryOfficeName` attribute and users in the New York office have an attribute value of `NewYork` for that attribute, to retrieve only users in the New York office, enter `(physicalDeliveryOfficeName=NewYork)` .

If you are using CAC authentication, to filter only active user accounts (excluding the disabled user accounts), enter `(!(userAccountControl:1.2.840.113556.1.4.803:=2))` . This criteria retrieves user accounts within AD belonging to `ldpgrp` group and with `userAccountControl` attribute value that is not `2` (disabled).

c) Enter a **User Name** for a user who has sufficient credentials to browse the LDAP server. For example, if you are connecting to an OpenLDAP server where user objects have a `uid` attribute, and the object for the administrator in the Security division at your example company has a `uid` value of `NetworkAdmin` , you might enter `uid=NetworkAdmin,ou=security,dc=example,dc=com.`

d) Enter the user password in the **Password** and the **Confirm Password** fields.

e) (Optional) Click **Show Advanced Options** to configure these advanced options.

- **Encryption** —Click **None** , **TLS** , or **SSL** .

  If you change the encryption method after specifying a port, you reset the port to the default value for that method. For **None** or **TLS** , the port resets to the default value of 389

  - **None** — LDAP communication is unencrypted.
  - **TLS** — LDAP communication is upgraded to TLS using the STARTTLS operation on default port 389
  - **SSL** — LDAP communication is encrypted using TLS immediately upon connection (LDAPS) on default port 636.

- **SSL Certificate Upload Path** —For SSL or TLS encryption, click **Choose File** and choose the complete CA chain certificate.

  Note
  Do not choose a binary certificate (PKCS12, DER, and alike) file because Firewall Threat Defense does not support them.

  To remove the uploaded certificate, check the **Clear loaded certificate** check box. This option only appears when you have uploaded a certificate, and when you are in the Edit mode of the external authentication object.

If you had previously uploaded a certificate and want to replace it, reupload the new certificate (complete CA chain), and redeploy the configuration to your devices to copy over the new certificate.

Note
TLS encryption requires a certificate on all platforms. We recommend that you **always** upload a certificate for SSL to prevent adversary-in-the-middle attacks.

- **User Name Template** —Provide a template that corresponds with your **UI Access Attribute** . For example, to authenticate all users who work in the Security organization of the Example company by connecting to an OpenLDAP server where the UI access attribute is `uid` , you might enter `uid=%s,ou=security,dc=example,dc=com` in the **User Name Template** field. For a Microsoft Active Directory server, you could enter `%s@security.example.com` .

  This field is required for CAC authentication.

- **Shell User Name Template** —Provide a template that corresponds with your **CLI Access Attribute** to authenticate CLI users. For example, to authenticate all users who work in the Security organization by connecting to an OpenLDAP server where the CLI access attribute is `sAMAccountName` , you might enter `%s` in the **Shell User Name Template** field.

- **Timeout (Seconds)** —Enter the number of seconds before rolling over to the backup connection, between 1 and 1024. The default is 30.

  Note
  The timeout range is different for Firewall Threat Defense and the Firewall Management Center , so if you share an object, be sure not to exceed the Firewall Threat Defense 's smaller timeout range (1-30 seconds). If you set the timeout to a higher value, the Firewall Threat Defense LDAP configuration will not work.

Step 4 Configure **Attribute Mapping** to retrieve users based on an attribute.

- Enter a **UI Access Attribute** , or click **Fetch Attrs** to retrieve a list of available attributes. For example, on a Microsoft Active Directory Server, you may want to use the UI access attribute to retrieve users, because there may not be a `uid` attribute on Active Directory Server user objects. Instead, you can search the `userPrincipalName` attribute by typing `userPrincipalName` in the **UI Access Attribute** field.

  This field is required for CAC authentication.

- Set the **CLI Access Attribute** if you want to use a shell access attribute other than the user distinguished type. For example, on a Microsoft Active Directory Server, use the `sAMAccountName` CLI access attribute to retrieve CLI access users by typing `sAMAccountName` .

Step 5 (Optional) Configure **Group Controlled Access Roles**.

If you do not configure a user's privileges using group-controlled access roles, a user has only the privileges granted by default in the external authentication policy.

a) (Optional) In the fields that correspond to user roles, enter the distinguished name for the LDAP groups that contain users who should be assigned to those roles.

Any group you reference must exist on the LDAP server. You can reference static LDAP groups or dynamic LDAP groups. Static LDAP groups are groups where membership is determined by group object attributes that point to specific users, and dynamic LDAP groups are groups where membership is determined by creating an LDAP search that retrieves group users based on user object attributes. Group access rights for a role only affect users who are members of the group.

If you use a dynamic group, the LDAP query runs exactly as it is configured on the LDAP server. To prevent infinite loops caused by search syntax errors, the device limits recursions of a search to four.

Example:

Enter this in the **Administrator** field to authenticate names in the information technology organization at the Example company:

```
cn=itgroup,ou=groups, dc=example,dc=com
```

b) Choose a **Default User Role** for users that do not belong to any of the specified groups.

c) If you use static groups, enter a **Group Member Attribute** .

Example:

If the `member` attribute is used to indicate membership in the static group for default Security Analyst access, enter `member` .

d) If you use dynamic groups, enter a **Group Member URL Attribute** .

Example:

If the `memberURL` attribute contains the LDAP search that retrieves members for the dynamic group you specified for default Admin access, enter `memberURL` .

If you change a user's role, you must save/deploy the changed external authentication object and also remove the user from the **Users** screen. The user will be re-added automatically the next time they log in.

Step 6 (Optional) Set the **CLI Access Filter** to allow CLI users.

To prevent LDAP authentication of CLI access, leave this field blank. To specify CLI users, choose one of these methods:

- To use the same filter you specified when configuring authentication settings, check the check box of **Same as Base Filter** .
- To retrieve administrative user entries based on attribute value, enter the attribute name, a comparison operator, and the attribute value you want to use as a filter, enclosed in parentheses. For example, if all network administrators have a `manager` attribute which has an attribute value of `shell` , you can set a base filter of `(manager=shell)` .

The usernames must be Linux-valid:

- Maximum 32 alphanumeric characters, plus period (.), hyphen (-), and underscore (_)
- All lowercase
- Cannot start with hyphen (-); cannot be all numbers; cannot include at sign (@) or slash (/)

Note
Users with CLI access can gain Linux shell access with the **expert** command. "Linux shell users can obtain root privileges, which presents a security risk. Make sure that you restrict the list of users with CLI or Linux shell access.

Note
Do not create any internal users that have the same user name as users included in the **CLI Access Filter**. The only internal Firewall Management Center user should be **admin** ; do not include an **admin** user in the **CLI Access Filter**.

Step 7 (Optional) Click **Test** to test connectivity to the LDAP server.

The test output lists valid and invalid user names. Valid user names are unique, and can include underscores (_), periods (.), hyphens (-), and alphanumeric characters. Note that testing the connection to servers with more than 1000 users only returns 1000 users because of UI page size limitations. If the test fails, refer to Troubleshoot LDAP authentication connections, on page 232 .

Step 8 (Optional) You can also enter **Additional Test Parameters** to test user credentials for a user who should be able to authenticate: enter a **User Name** `uid` and **Password** , and then click **Test**. Click **Save**. Enable use of this server. Refer to Enable external authentication for users on the Firewall Management Center, on page 166.

If you are connecting to a Microsoft Active Directory Server and supplied a UI access attribute in place of `uid` , use the value for that attribute as the user name. You can also specify a fully qualified distinguished name for the user.

Tip
If you mistype the name or password of the test user, the test fails even if the server configuration is correct. To verify that the server configuration is correct, click **Test** without entering user information in the **Additional Test Parameters** field first. If that succeeds, supply a user name and password to test with the specific user.

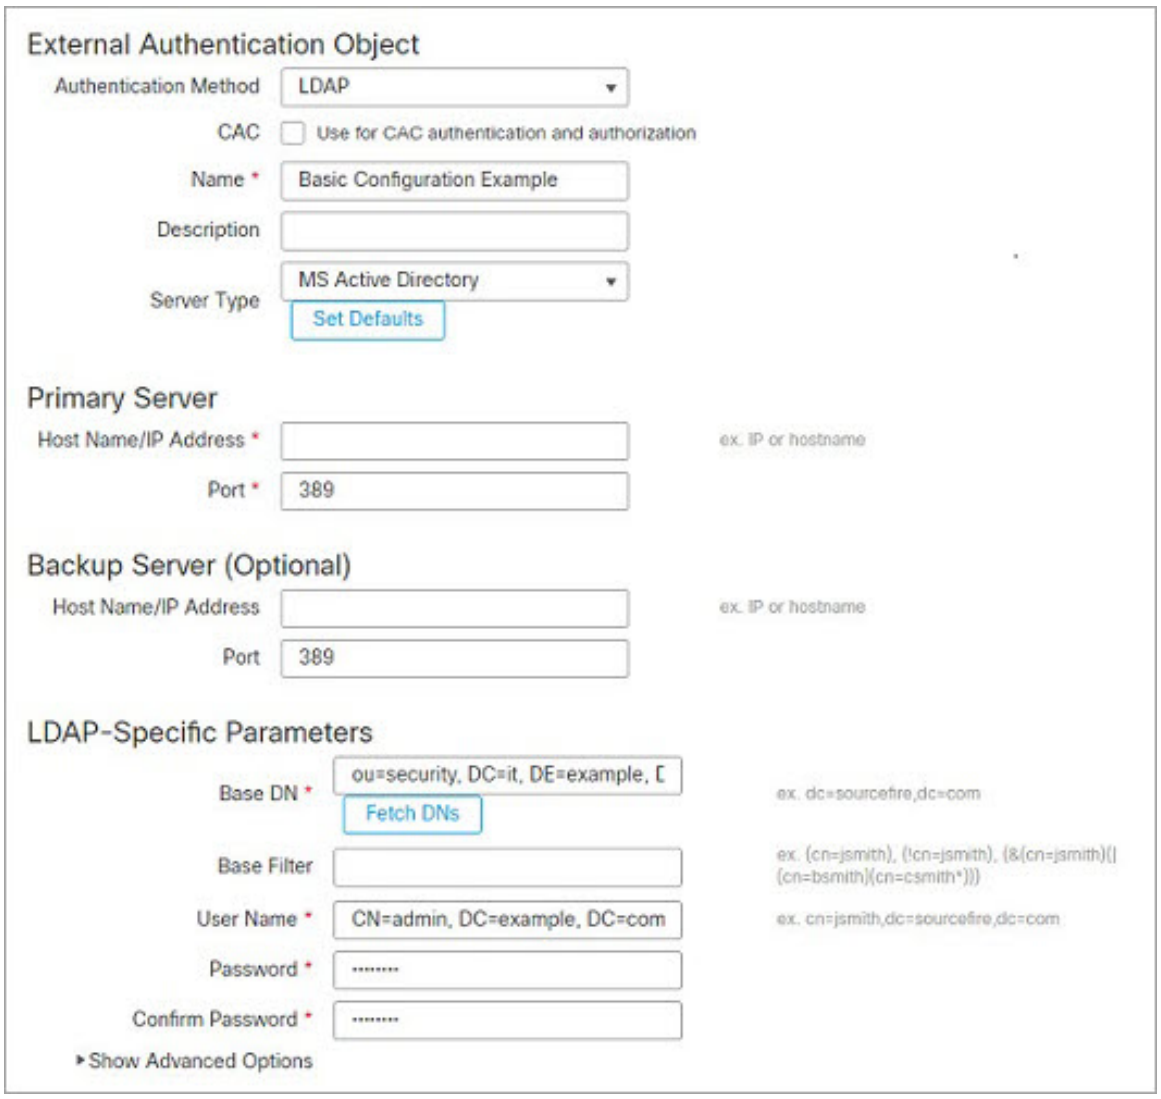
Example:

To test if you can retrieve the `JSmith` user credentials at the Example company, enter `JSmith` and the correct password.

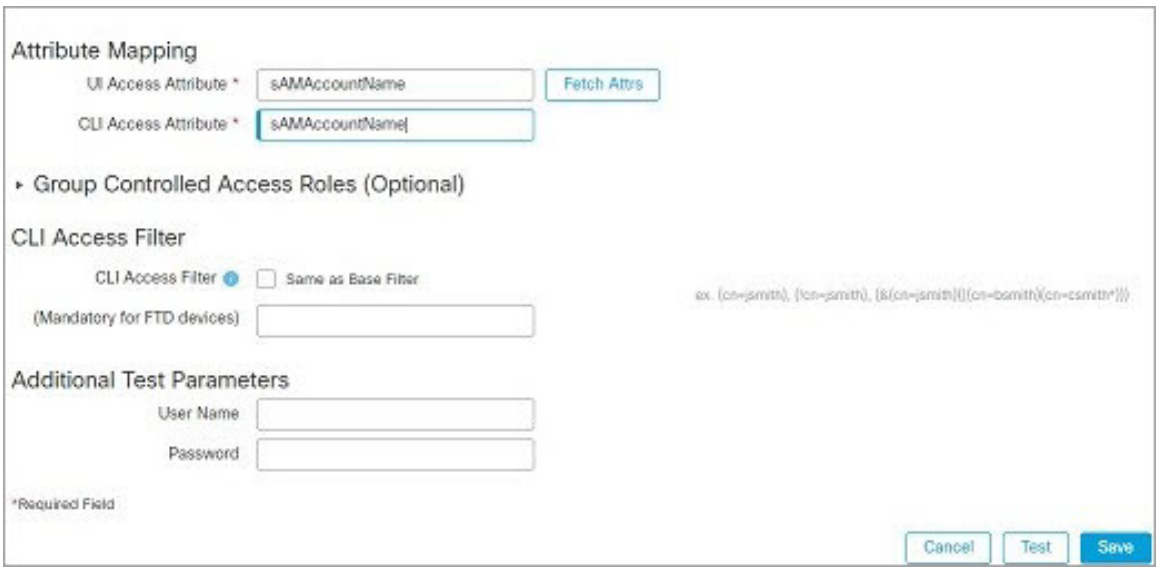
Examples

Basic Example

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



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



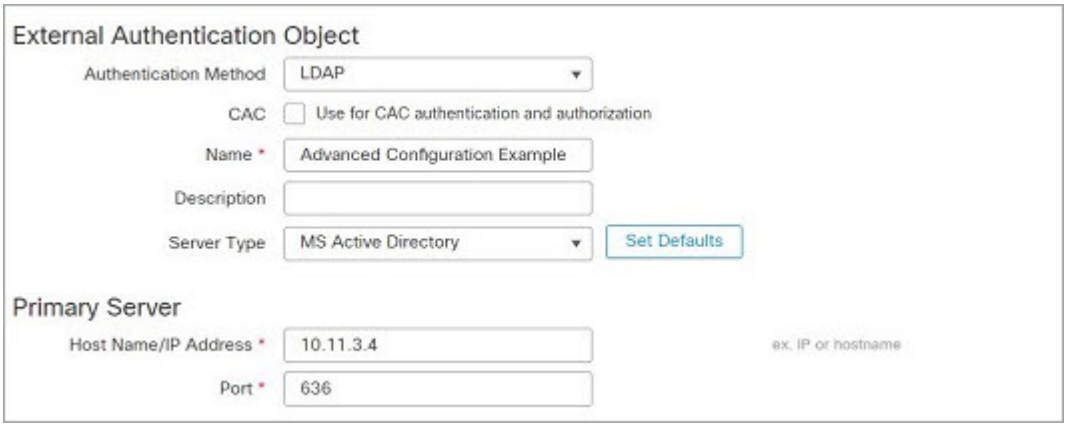
However, because this server is a Microsoft Active Directory server, it uses the `sAMAccountName` attribute to store user names rather than the `uid` attribute. Choosing the MS Active Directory server type and clicking **Set Defaults** sets the UI Access Attribute to `sAMAccountName` . As a result, the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

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

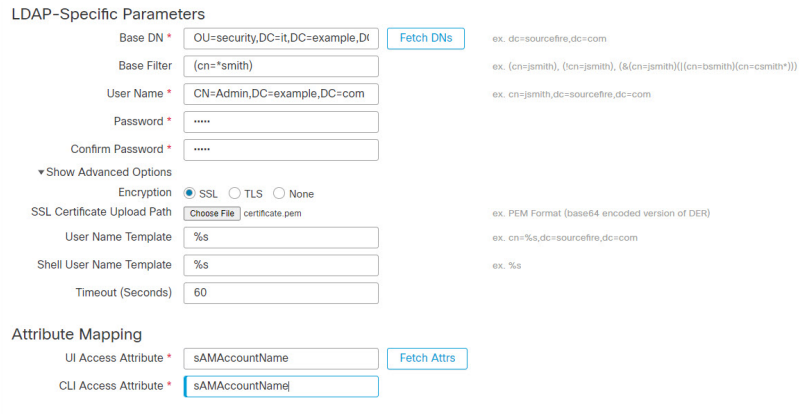
Note that because no base filter is applied to this server, the system checks attributes for all objects in the directory indicated by the base distinguished name. Connections to the server time out after the default time period (or the timeout period set on the LDAP server).

Advanced Example

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



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

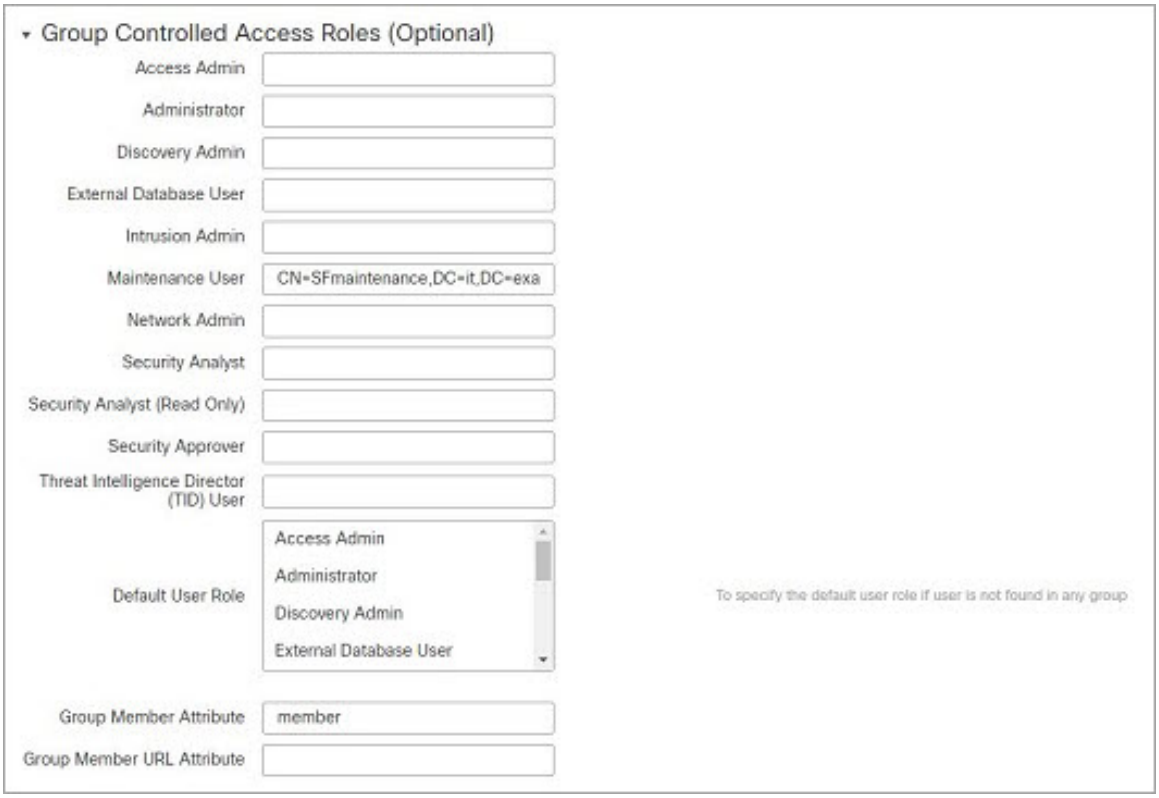


The connection to the server is encrypted using SSL and a certificate named `certificate.pem` is used for the connection. In addition, connections to the server time out after 60 seconds because of the **Timeout (Seconds)** setting.

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

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

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



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



Add a RADIUS external authentication object for Firewall Management Center

Add a RADIUS server to support external users for device management.

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

Note that Firewall Management Center supports only these six attributes in the access request of the RADIUS authentication:

- `NAS-IP-Address` : Used in UI and CLI authentication

- `NAS-Identifier` : Used in CLI authentication
- `NAS-Port` : Used in UI and CLI authentication
- `NAS-Port-Type` : Used in CLI authentication
- `Service-Type` : Used in UI and CLI authentication
- `Calling-Station-Id` : Used in CLI authentication

Note

- `NAS-Port` attribute: The attribute value is dynamically generated and cannot be configured.
- `NAS-Port-Type` attribute: The server always sets the attribute value to `Virtual`, regardless of whether the device is virtual or physical. of whether the Firewall Management Center is virtual or physical.
- `NAS-Identifier` attribute: The attribute value is always `sshd`.
- `Service-Type` attribute: Firewall Management Center supports only `Login` and `Authenticate Only` service types

Procedure

Step 1 Choose **System** (⚙) > **Users** and click **External Authentication**.

Step 2 Click **Add icon** (+). In the **Add External Authentication Object** window, set the **Authentication Method** to **RADIUS**, and enter a unique **Name** and optional **Description**.

Step 3 Check the **RADIUS Server-Enabled Message Authenticator** check box. This requires the Message-Authenticator attribute in all RADIUS responses and ensures that every response from the RADIUS server is securely verified by the Firewall Threat Defense.

The feature is enabled by default for new RADIUS servers. Enable it for existing servers after the upgrade. Enable message authenticators to protect your firewalls from potential attacks. Ensure that your RADIUS server has the Message-Authenticator configuration.

Note
The Message-Authenticator attribute persists whether the feature is enabled or disabled. The server always requires this attribute. When the check box is enabled in the configuration, it also becomes a requirement from the Firewall Management Center.

Step 4 For the **Primary Server**, enter a **Host Name/IP Address**. Then, change the **Port** from the default and then enter the **RADIUS Secret Key**.

Step 5 (Optional) Enter the **Backup Server** parameters, then provide the **RADIUS-Specific Parameters**.

a) Enter the **Timeout** in seconds before retrying the primary server, between 1 and 1024. The default is 30.

Note
The timeout range is different for the Firewall Threat Defense and the Firewall Management Center so if you share an object, be sure not to exceed the Firewall Threat Defense 's smaller timeout range (1-300 seconds). If you set the timeout to a higher value, the Firewall Threat Defense RADIUS configuration will not work.

b) Enter the **Retries** before rolling over to the backup server. The default is 3.

c) In the fields that correspond to user roles, enter the name of each user or identifying attribute-value pair that should be assigned to those roles.

Separate usernames and attribute-value pairs with commas.

Example:

If you know all users who should be Security Analysts have the value `Analyst` for their `User-Category` attribute, you can enter `User-Category=Analyst` in the **Security Analyst** field to grant that role to those users.

Example:

To grant the Administrator role to the users `jsmith` and `jdoe`, enter `jsmith, jdoe` in the **Administrator** field.

Example:

To grant the Maintenance User role to all users with a `User-Category` value of `Maintenance`, enter `User-Category=Maintenance` in the **Maintenance User** field.

d) Select the **Default User Role** for users that do not belong to any of the specified groups.

After you change a user's role, save and deploy the updated external authentication object. Remove the user from the **Users** screen. The system adds the user when they log in again.

Step 6 (Optional) **Define Custom RADIUS Attributes**.

If your RADIUS server returns values for attributes not included in the `dictionary` file in `/etc/radiusclient/`, define those attributes. You must define them if you plan to use the attributes to set roles for users.You can locate the attributes returned for a user by looking at the user's profile on your RADIUS server.

Note
Before you configure custom RADIUS attributes, obtain the Attribute ID and Type from the RADIUS server administrator. If there is a mismatch, the RADIUS packet capture displays the custom RADIUS attribute as `Unassigned Attribute`.

a) Enter an **Attribute Name**.

Enter an attribute name using alphanumeric characters. Separate words in an attribute name with dashes instead of spaces.

b) Enter the **Attribute ID** as an integer.

Enter an integer for the attribute ID. Make sure the attribute ID does not conflict with any existing IDs in the `etc/radiusclient/dictionary` file.

c) Choose the **Attribute Type** from the drop-down list.

You also specify the type of attribute: string, IP address, integer, or date.

d) Click **Add** to add the custom attribute.

When you create a RADIUS authentication object, a new dictionary file for that object is created on the device in the `/var/sf/userauth` directory. Any custom attributes you add are added to the dictionary file.

Example:

If a RADIUS server is used on a network with a Cisco router, you might want to use the `Ascend-Assign-IP-Pool` attribute to grant a specific role to all users logging in from a specific IP address

pool. `Ascend-Assign-IP-Pool` is an integer attribute that defines the address pool where the user is allowed to log in, with the integer indicating the number of the assigned IP address pool.

To declare that custom attribute, you create a custom attribute with an attribute name of `Ascend-IP-Pool-Definition`, an attribute ID of `218`, and an attribute type of `integer`.

You could then enter `Ascend-Assign-IP-Pool=2` in the **Security Analyst (Read Only)** field to grant read-only security analyst rights to all users with an `Ascend-IP-Pool-Definition` attribute value of `2`.

Step 7 (Optional) In the **CLI Access Filter** area **Administrator CLI Access User List** field, enter the usernames that should have CLI access, separated by commas.

Enter usernames that match those on the RADIUS server. Usernames must be valid Linux usernames:

- Maximum of 32 alphanumeric characters, plus period (.), hyphen (-), and underscore (_).

  (_)

- All letters must be lowercase.
- "The username cannot start with a hyphen (-), cannot be composed entirely of numbers, and cannot include an at sign (@) or a slash (/).

Leave the field blank to disable RADIUS authentication for CLI access.

Note
Users with CLI access can gain Linux shell access with the **expert** command. Linux shell users can obtain root privileges, which can present a security risk. Make sure that you restrict the list of users with CLI or Linux shell access.

Note
Remove any internal users that have the same user name as users included in the shell access filter. For the Firewall Management Center, the only internal CLI user is **admin**, so do not also create an **admin** external user.

Step 8 (Optional) Click **Test** to test Firewall Management Center connectivity to the RADIUS server.

Step 9 (Optional) You can also enter **Additional Test Parameters** to test user credentials for a user who should be able to authenticate: enter a **User Name** and **Password**, and then click **Test**. Click **Save**. For more information, refer to Enable external authentication for users on the Firewall Management Center, on page 166.

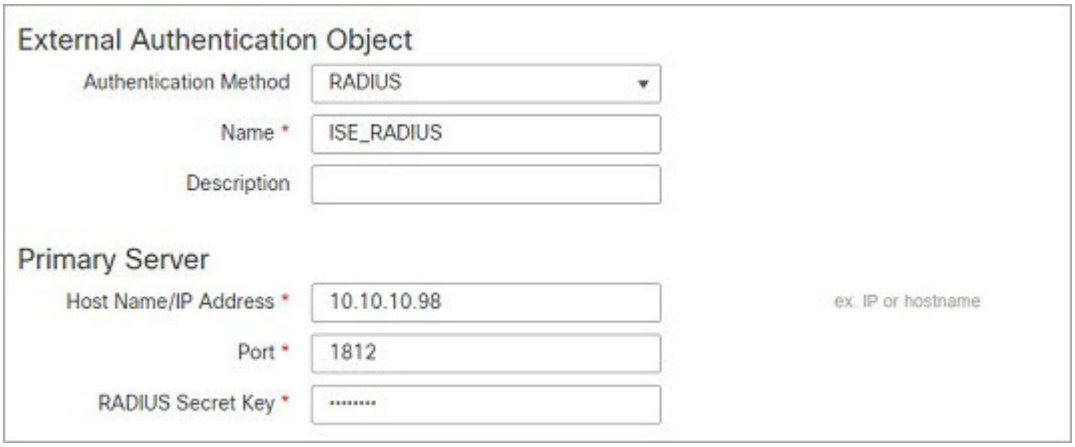
Tip
If you mistype the name or password of the test user, the test fails even if the server configuration is correct. To verify that the server configuration is correct, click **Test** without entering user information in the **Additional Test Parameters** field first. If that succeeds, supply a user name and password to test with the specific user.

Example:

To test if you can retrieve the `JSmith` user credentials at the Example company, enter `JSmith` and the correct password.

Simple user role assignments

This figure illustrates a sample RADIUS login authentication object for a server running Cisco Identity Services Engine (ISE) with an IP address of 10.10.10.98 on port 1812. No backup server is defined.



This example shows RADIUS-specific parameters, including the timeout (30 seconds) and number of failed retries before the Secure Firewall System attempts to contact the backup server, if any.

This example illustrates important aspects of RADIUS user role configuration:

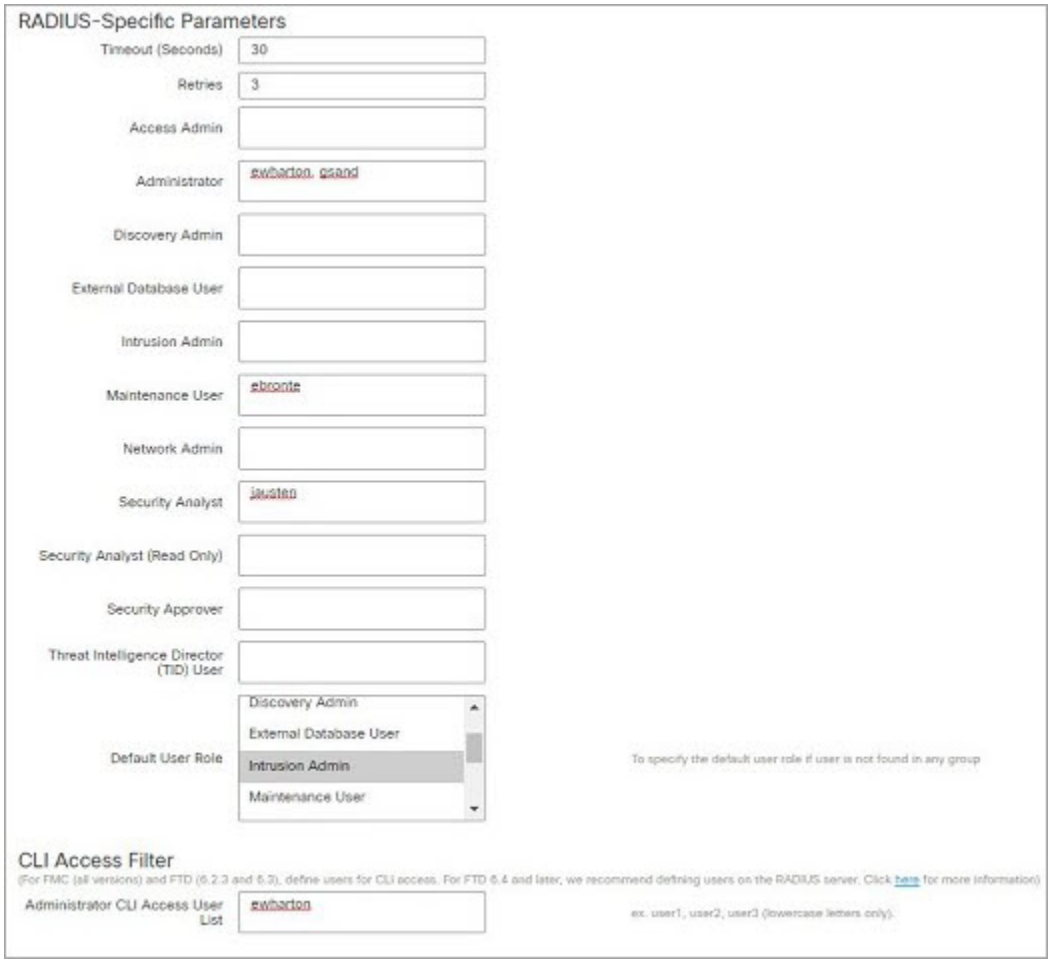
Users `ewharton` and `gsand` are granted web interface Administrative access.

The user `cbronte` is granted web interface Maintenance User access.

The user `jausten` is granted web interface Security Analyst access.

The user `ewharton` can log into the device using a CLI account.

This graphic depicts the role configuration for the example:

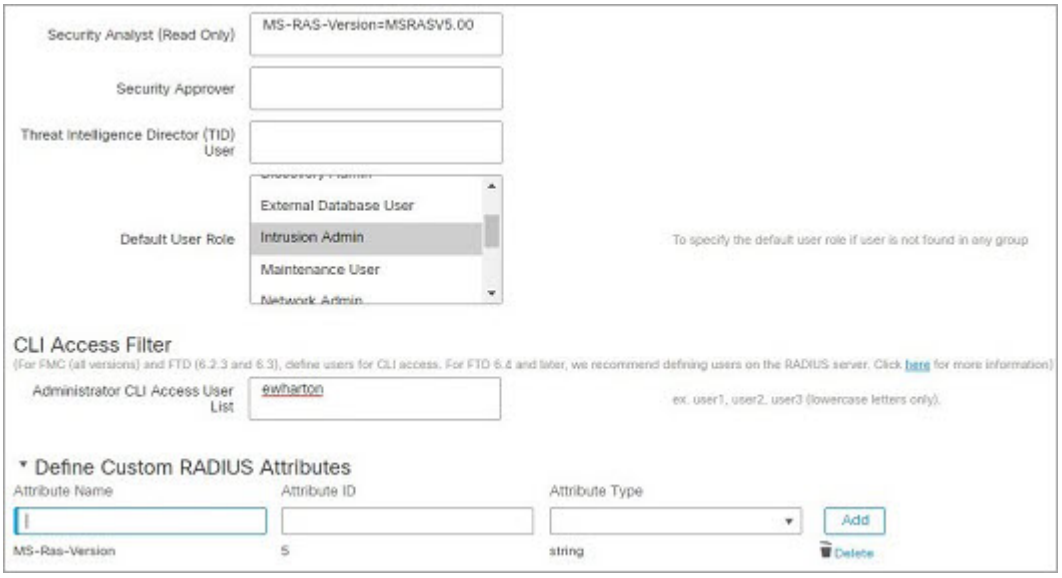


Roles for users matching an attribute-value pair

You can use an attribute-value pair to identify users who should receive a particular user role. If the attribute you use is a custom attribute, you must define the custom attribute.

This figure illustrates the role configuration and custom attribute definition in a sample RADIUS login authentication object for the same ISE server as in the previous example.

In this example, however, the `MS-RAS-Version` custom attribute is returned for one or more of the users because a Microsoft remote access server is in use. Note the `MS-RAS-Version` custom attribute is a string. In this example, all users logging in to RADIUS through a Microsoft version 5.00 remote access server should receive the Security Analyst (Read Only) role, so you enter the attribute-value pair of `MS-RAS-Version=MSRASV5.00` in the **Security Analyst (Read Only)** field.



Enable external authentication for users on the Firewall Management Center

Enable external authentication for users to allow the Firewall Management Center to verify user credentials with an LDAP or RADIUS server, enhancing security and centralized user management.

When you enable external authentication for management users, the Firewall Management Center verifies the user credentials with an LDAP or RADIUS server as specified in an External Authentication object.

Before you begin

Add one or more external authentication objects according to Add an LDAP external authentication object for the Firewall Management Center , on page 152 and Add a RADIUS external authentication object for Firewall Management Center , on page 160.

Follow these steps to enable external authentication for users on the Firewall Management Center:

Procedure

Step 1 Choose **System** (⚙) > **Users**.

Step 2 Click **External Authentication**.

Step 3 Set the default user role for external web interface users.

If a user does not have a role, they cannot perform any actions. Roles set in the external authentication object override the default user role.

a) Click the **Default User Role** value (by default, no selection).
b) In the **Default User Role Configuration** dialog box, check the role or roles that you want to use.
c) Click **Save**.

Step 4 Click the **Slider enabled** ( ) next to the each external authentication object that you want to use. If you enable more than one object, users are compared against servers in the specified order. Refer to the next step to reorder servers.

If you enable shell authentication, you must enable an external authentication object that includes a **CLI Access Filter**. Also, CLI access users can only authenticate against the server whose authentication object is highest in the list.

Step 5 (Optional) Drag and drop servers to change the order in which the system accesses them during authentication.

Step 6 From the **Shell Authentication** drop-down list, click **Enabled** if you want to allow CLI access for external users.

Note
The multidomain feature is not supported in CLI. Therefore, the **Shell Authentication** option is available only in the Global domain and not in Subdomains.

The first external authentication object name is shown next to the **Enabled** option to remind you that only the first object is used for CLI.

Step 7 Click **Save and Apply**.

When you complete this task, the system enables external authentication for users. Firewall Management Center verifies their credentials with the LDAP or RADIUS servers you configured and assigns roles as you specified.

Configure common access card authentication with LDAP

Configure CAC authentication with LDAP to allow users to log in directly without providing a separate username and password for the device.

If your organization uses Common Access Cards (CACs), you can configure LDAP authentication to authenticate Firewall Management Center users who log in to the web interface. With CAC authentication, users have the option to log in directly without providing a separate username and password for the device.

The system identifies CAC-authenticated users by their electronic data interchange personal identifier (EDIPI) numbers.

After 24 hours of inactivity, the device deletes CAC-authenticated users from the **Users** tab. The system re-adds these users after each login. Reconfigure any manual user role changes after users are re-added.

Caution Assign a default access role with the least privileges required when you configure CAC authentication with LDAP. When a user first logs in to the system with their CAC credentials, their account will be assigned this default access role.

If you do not follow the principles of least privilege while assigning the default access role, users may be assigned an unintended privilege level on subsequent logins. This could result in the users having privileges beyond their required access role.

If a user who logged in with the default access role needs a temporary privilege elevation, an administrator can assign a higher access role. The system revokes this privilege after 24 hours of inactivity, returning the user to their default access role.

If a user requires a permanent access role reassignment to a higher privilege level, such as System Admin, use the **Group Controlled Access Roles** method to grant admin access. With this method, the assigned access role remains active for more than 24 hours, and users get the correct privileges based on group assignment. For more information on configuring Group Controlled Access Roles, refer to the Add an LDAP External Authentication Object for Management Center section.

Before you begin

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

Follow these steps to configure Common Access Card authentication with LDAP:

Procedure

Step 1 Insert a CAC as your organization directs.

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

Step 3 If prompted, enter the PIN for the CAC you inserted.

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

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

You **cannot** yet log in using your CAC credentials. Choose **System** (⚙) > **Users** > **External Authentication**.

Step 6 Create an LDAP authentication object exclusively for CAC. Refer to the procedure in Add an LDAP external authentication object for the Firewall Management Center , on page 152. You must configure these items:

- **CAC** check box.
- Under **LDAP-Specific Parameters**, click **Show Advanced Options**, and then enter the **User Name Template**.
- Under **Attribute Mapping**, enter the **UI Access Attribute**.

Click **Save**.

Step 7 Enable external authentication and CAC authentication as described in Enable external authentication for users on the Firewall Management Center, on page 166. Choose **System** (⚙) > **Configuration**, and click **HTTPS Certificate**.

Step 8 Import an HTTPS server certificate if needed. Refer to the procedure outlined in Import HTTPS server certificates, on page 80.

The same certificate authority (CA) must issue the HTTPS server certificate and the user certificates on the CACs you plan to use.

Step 9 Under **HTTPS Client Certificate Settings**, choose **Enable Client Certificates**. For more information, refer to Require valid HTTPS client certificates, on page 81. Log in to the device according to Log into the Secure Firewall Management Center with CAC credentials, on page 36.

Users can now log in to the device using their CAC credentials without providing a separate username and password.

SAML Single Sign-On

SAML Single Sign-On (SSO) is a centralized authentication system that

- enables a central identity provider (IdP) to provide authentication and authorization for users logging into the Firewall Management Center and other applications within an organization,
- allows applications configured to participate in the SSO arrangement to function as federated service provider applications, and
- permits SSO users to log in once to gain access to all service provider applications that are members of the same federation.

SAML SSO features

A SAML SSO feature is an authentication and authorization mechanism that

- redirects users to an identity provider for authentication instead of entering credentials on the management center login page,
- uses the browser as an intermediary for communication between the management center and the identity provider, and
- does not require a direct network connection between the management center and the identity provider.

You can use any provider that supports the Security Assertion Markup Language (SAML) 2.0 standard for authentication and authorization with the management center. In a multi-tenant management center, you can assign SAML users to specific subdomains. You can only configure this at the global domain level.

You can configure these SSO providers in the management center web interface:

- Okta
- OneLogin
- Azure

- PingID's PingOne for Customers cloud solution
- Other

Note

- If your identity provider requires the service provider's signature on authentication requests, SSO fails because the management center cannot sign SAML authentication requests.
- The Cisco Secure Sign On SSO product does not recognize the Firewall Management Center as a pre-integrated service provider.

A user clicks the SSO link on the management center login page, is redirected to the identity provider for authentication, and upon successful authentication, is redirected back to the management center web interface and logged in.

If the identity provider requires a signed authentication request and the management center cannot provide it, SSO fails and users cannot log in using SAML single sign-on.

Using SAML single sign-on is like using a passport to access multiple countries. The identity provider acts as the passport authority, so you do not have to present separate credentials.

Guidelines for SSO configuration on Firewall Management Center

Restrictions for SSO provider configuration

The Firewall Management Center can support SSO with only one SSO provider at a time. You cannot configure the Firewall Management Center to use multiple SSO providers, such as both Okta and OneLogin, simultaneously.

SSO configuration in high availability deployments

Firewall Management CenterFirewall Management Centers in a high availability configuration can support SSO, but you must address several considerations.

- SSO configuration does not synchronize between members of a high availability pair. Configure SSO separately for each member.
- Both Firewall Management Centers in a high availability pair must use the same IdP for SSO. Configure a service provider application at the IdP for each Firewall Management Center configured for SSO.
- Before a user can use SSO to access the secondary Firewall Management Center for the first time, the user must first use SSO to log into the primary Firewall Management Center at least once.
- When configuring SSO for Firewall Management Centers in a high availability pair:
  - If you configure SSO on the primary Firewall Management Center, you do not need to configure SSO on the secondary Firewall Management Center.
  - If you configure SSO on the secondary Firewall Management Center, you are required to configure SSO on the primary Firewall Management Center as well. SSO users must log into the primary Firewall Management Center at least once before logging into the secondary Firewall Management Center.

SSO configuration in multi-tenancy environments

In a Firewall Management Center that uses multi-tenancy, the SSO configuration can be applied only at the global domain level. That configuration applies to the global domain and all subdomains.

SSO configuration permissions and limitations

- Only users with the Admin role authenticated internally or by LDAP or RADIUS can configure SSO.
- The Firewall Management Center does not support SSO initiated from the IdP.
- The Firewall Management Center does not support logging in with CAC credentials for SSO accounts.
- Do not configure SSO in deployments using CC mode.
- View SSO activity in the Firewall Management Center audit log. The Subsystem field shows whether an activity was Login or Logout.

Related Topics

High availability, on page 313
Domains, on page 243
Log into the Secure Firewall Management Center with CAC credentials, on page 36
Security certifications compliance, on page 347
Audit records, on page 425

SSO user accounts

An SSO user account is an identity management entity that

- has its username and password established at the identity provider (IdP)
- does not appear on the Firewall Management Center web interface Users page until the first login, and
- can have certain account characteristics configured from the management center web interface.

Account configuration options

You can configure users and groups directly in the identity provider, or import them from user management applications such as Active Directory, RADIUS, or LDAP.

This documentation explains how to configure the Firewall Management Center to work with the IdP to support SSO. It assumes that IdP users and groups are already established. To configure an IdP to support users and groups from other user management applications, consult the IdP vendor documentation.

Most account characteristics for SSO users, including the user name and password, are established at the IdP.

SSO accounts do not appear on the Firewall Management Center web interface Users page until those accounts log in the first time.

Note The system requires that user names for SSO accounts as well as the NameID attribute the IdP sends to the Firewall Management Center during the SAML login process must be both be valid email addresses. Many IdP's automatically use the username of the user trying to logon as the NameID attribute, but you should confirm this is the case for your IdP. Keep this in mind when configuring a service provider application at your IdP and creating IdP user accounts that are to be granted SSO access to the Firewall Management Center.

This account characteristic for SSO users can be configured from the Firewall Management Center web interface under **System** (⚙) > **Users** and click the **Edit User**:

- Real Name
- Exempt from Browser Session Timeout

SSO user account example

For example, an SSO user account is created at the IdP and appears in the management center web interface only after the user logs in for the first time. The Real Name and Exempt from Browser Session Timeout options can be configured in the Edit User dialog.

Non-SSO user account

A local user account created directly in the management center web interface is not an SSO user account and does not rely on the IdP for establishment.

User role mappings for SSO users

A user role mapping is an access control configuration that

- associates an IdP role with one or more Firewall Management Center roles, allowing access to one or more domains using the **Group Member Attribute Value**
- coordinates configuration settings between the Firewall Management Center and the SSO IdP application, and
- assigns user roles to users or groups defined at the IdP application, based on role attribute values and expressions.

Role attribute configuration for SSO user mappings

In user role mapping, the IdP sets a role attribute for the Firewall Management Center service provider application. Each user or group with access to that Firewall Management Center is configured with a string or expression for the role attribute. Attribute value requirements depend on the IdP you use.

At the Firewall Management Center, the name of the role attribute and the list of expressions assigned to Firewall Management Center user roles are part of the SSO configuration. When a user logs in to the Firewall Management Center using SSO, the Firewall Management Center compares the value of the user's role attribute or the group's attribute (depending on configuration) to the expressions for each Firewall Management Center user role. If the values match, you receive all roles that correspond to the matching expressions.

Note You can configure Firewall Management Center roles to be mapped based on individual user permissions or based on group permissions, but a single Firewall Management Center application cannot support role mapping for both groups and individual users.

Default role assignment for SSO users

By default, all users who are given SSO access to a Firewall Management Center are assigned the Security Analyst (Read Only) role. You can change this default, as well as override it for specific SSO users or groups using user role mapping.

Enable SSO at the Firewall Management Center

Enable single sign-on (SSO) at the Firewall Management Center to allow users to authenticate through an external SSO provider.

Perform this task after configuring a service provider application for your SSO provider and assigning users or groups to it. This process is intended for organizations that want centralized authentication and improved access management.

Before you begin

In the SAML SSO management application, configure a service provider application for the Firewall Management Center and assign users or groups to it.

- To configure a Firewall Management Center service provider application for Okta, refer to Configure the Firewall Management Center service provider application for Okta, on page 175.
- To configure a Firewall Management Center service provider application for OneLogin, refer to Configure a Firewall Management Center service provider application for OneLogin, on page 188.
- To configure a Firewall Management Center service provider application for Azure, refer to Configure the Firewall Management Center service provider application for Azure, on page 202.
- To configure a Firewall Management Center service provider application for PingID's PingOne for Customers cloud solution, refer to Configure a service provider application for PingID PingOne for Customers, on page 215.
- To configure a Firewall Management Center service provider application for any SAML 2.0-compliant SSO provider, refer to Configure Firewall Management Center service provider application for any SAML 2.0-compliant SSO provider, on page 220.

Follow these steps to enable SSO at the Firewall Management Center:

Procedure

Step 1 Choose **System** (⚙) > **Users** > **Single Sign-On**.

Step 2 Click the **Single Sign-On (SSO) Configuration** slider to enable SSO.

Step 3 Click the **Configure SSO** button.

Step 4 At the **Select Firewall Management Center SAML Provider** dialog box, click the radio button for the SSO IdP of your choice and click **Next**.

After completing these steps, SSO is enabled at the Firewall Management Center. Users can authenticate using the configured SSO provider, streamlining access and improving security.

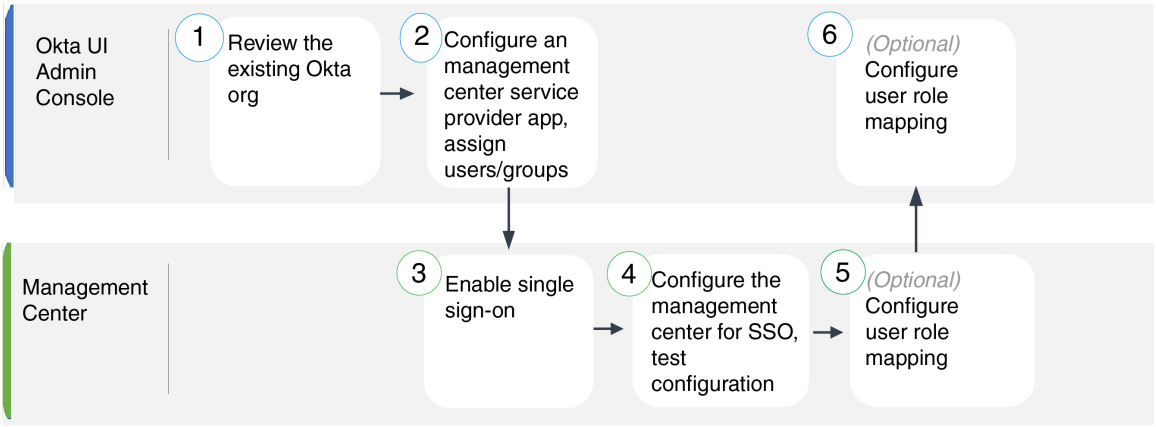
What to do next

Proceed with the instructions appropriate to your choice of SSO provider:

- Configure the Firewall Management Center for Okta SSO. Refer to Configure the Firewall Management Center for Okta SSO, on page 177.
- Configure the Firewall Management Center for SSO using PingID's PingOne for Customers cloud solution. Refer to Configure the Firewall Management Center for SSO with the PingID PingOne for Customers, on page 217.
- Configure the Firewall Management Center for Azure SSO. Refer to Configure the Firewall Management Center for Azure SSO, on page 203.
- Configure the Firewall Management Center for OneLogin SSO. Refer to Configure the Firewall Management Center for OneLogin SSO, on page 190.
- Configure the Firewall Management Center for SSO using any SAML 2.0-compliant provider. Refer to Configure the Firewall Management Center for SSO using any SAML 2.0-compliant SSO provider, on page 222.

Configure SSO with Okta

Refer to these tasks to configure SSO using Okta:



1	Okta UI Admin Console	Review the Okta Org, on page 175
2	Okta UI Admin Console	Configure the Firewall Management Center service provider application for Okta, on page 175

3	Firewall Management Center	Enable SSO at the Firewall Management Center, on page 173
4	Firewall Management Center	Configure the Firewall Management Center for Okta SSO, on page 177
5	Firewall Management Center	Configure User Role Mapping for Okta in the Firewall Management Center, on page 178
6	Okta UI Admin Console	Okta IdP user role mappings, on page 179

Review the Okta Org

In Okta, the entity that encompasses all the federated devices and applications that a user can access with the same SSO account is called an *org*. Before adding the Firewall Management Center to an Okta org, be familiar with its configuration; consider the following questions:

- How many users will have access to the Firewall Management Center?
- Are users within the Okta org members of groups?
- Are user and group definitions native to Okta or imported from a user management application such as Active Directory, RADIUS, or LDAP?
- Do you need to add more users or groups to the Okta org to support SSO on the Firewall Management Center?
- What kind of user role assignments do you want to make? (If you choose not to assign user roles, the Firewall Management Center automatically assigns a configurable default user role to all SSO users.)
- How must users and groups within the Okta org be organized to support the required user role mapping?

Keep in mind that you can configure Firewall Management Center roles to be mapped based on individual user permissions or based on group permissions, but a single Firewall Management Center application cannot support role mapping for both groups and individual users.

This documentation assumes you are already familiar with the Okta Classic UI Admin Console, and have an account that can perform configuration functions requiring Super Admin permissions. If you need more information, see Okta's documentation available online.

Configure the Firewall Management Center service provider application for Okta

Configure the Firewall Management Center service provider application in Okta to enable SAML-based single sign-on (SSO) for users and groups. This task allows you to integrate Okta with your Firewall Management Center for centralized authentication and access management.

Use these instructions at the Okta Classic UI Admin Console to create a Firewall Management Center service provider application within Okta and assign users or groups to that application. You should be familiar with SAML SSO concepts and the Okta admin console. This documentation does not describe all the Okta functions you need to establish a fully functional SSO org. For example, to create users and groups or to import user and group definitions from another user management application, refer to the Okta documentation.

Note If you plan to assign user groups to the Firewall Management Center application, do not also assign users within those groups as individuals.

Note The Firewall Management Center cannot support role mapping using multiple SSO attributes. You must select either user role mapping or group role mapping and configure a single attribute to convey user role information from OneLogin to the Firewall Management Center.

Assign users or groups to the application as appropriate for your organization.

Before you begin

- Familiarize yourself with the SSO federation and its user and groups; refer to Review the Okta Org, on page 175.
- Create user accounts and groups, or both, in your Okta org if necessary.

Note The system requires that user names for SSO accounts as well as the NameID attribute the IdP sends to the Firewall Management Center during the SAML login process must be both be valid email addresses. Many IdP's automatically use the username of the user trying to logon as the NameID attribute, but you should confirm this is the case for your IdP. Keep this in mind when configuring a service provider application at your IdP and creating IdP user accounts that are to be granted SSO access to the Firewall Management Center.

- Confirm the login URL for the target Firewall Management Center (`https://ipaddress_or_hostname`).

Note If your Firewall Management Center web interface can be reached with multiple URLs (for instance, a fully-qualified domain name as well as an IP address), SSO users must consistently access the Firewall Management Center using the login URL that you configure in this task.

Configure the Firewall Management Center service provider application for Okta by completing these steps:

Procedure

Step 1 From the Okta Classic UI Admin Console, create a service provider application for the Firewall Management Center. Configure the Firewall Management Center application with these selections:

- Select `Web` for the **Platform**.
- Select `SAML 2.0` for the **Sign on method**.

- Provide a **Single sign on URL**.

  This is the Firewall Management Center URL to which the browser sends information on behalf of the IdP.

  Append the string `SAML/acs` to the Firewall Management Center login URL. For example, use `https://ExampleFMC/SAML/acs`.

- Enable **Use this for Recipient URL and Destination URL**.
- Enter an **Audience URI (SP Entity ID)**.

  This is a globally unique name for the service provider (the Firewall Management Center), often formatted as a URL.

  Append the string `/SAML/metadata` to the Firewall Management Center login URL. For example, use `https://ExampleFMC/SAML/metadata`.

- For **Name ID Format** choose `Unspecified`.

Step 2 (Optional if you are assigning groups to the application.) Assign individual Okta users to the Firewall Management Center application. (If you plan to assign groups to the Firewall Management Center application, do not assign users that are members of those groups as individuals.)

Step 3 (Optional if you are assigning individual users to the application.) Assign Okta groups to the Firewall Management Center application.

Step 4 (Optional) To make SSO setup at the Firewall Management Center easier, you can download the SAML XML metadata file for the Firewall Management Center service provider application from Okta to your local computer.

After completing this task, your Firewall Management Center will be integrated with Okta for SAML SSO, allowing assigned users and groups to authenticate using Okta credentials.

What to do next

Enable single sign-on; refer to Enable SSO at the Firewall Management Center, on page 173.

Configure the Firewall Management Center for Okta SSO

Configure Okta SSO on the Firewall Management Center to allow users to authenticate using Okta credentials. This integration centralizes authentication and enhances security for your organization.

Use these instructions at the Firewall Management Center web interface.

This task is relevant when you want to enable single sign-on for users accessing the management center through Okta.

- Create a Firewall Management Center service provider application at the Okta Classic UI Admin Console; refer to Configure the Firewall Management Center service provider application for Okta, on page 175.
- Enable single sign-on; refer to Enable SSO at the Firewall Management Center, on page 173.

Before you begin

Before you begin, ensure you have completed the required setup in Okta and enabled single sign-on on the management center.

Follow these steps to configure Okta SSO on the management center.

Procedure

Step 1 (This step continues directly from Enable SSO at the Firewall Management Center, on page 173.) At the **Configure Okta Metadata** dialog box, you have two choices:

- To enter the SSO configuration information manually:
  a. Click the **Manual Configuration** radio button.
  b. Enter the values from the Okta SSO Service Provider application. Retrieve the values from the Okta Classic UI Admin Console.
     - **Identity Provider Single Sign-On (SSO) URL**
     - **Identity Provider Issuer**
     - **X.509 Certificate**
- If you have saved the XML metadata file generated by Okta on your local computer, (step 4 in Configure the Firewall Management Center service provider application for Okta, on page 175), upload the file to the Firewall Management Center.
  a. Click the **Upload XML File** radio button.
  b. Follow the on-screen instructions to navigate to and choose the XML metadata file on your local computer.

Step 2 Click **Next**.

Step 3 At the **Verify Metadata** dialog, review the configuration parameters and click **Save**.

Step 4 Click **Test Configuration**. If an error message appears, review the SSO configuration for the Firewall Management Center and the Okta service provider application configuration. Correct any errors, and try again.

Step 5 When a message confirms a successful configuration test, click **Apply**.

After you complete these steps, Okta SSO is enabled on the management center. Users can now authenticate using their Okta credentials.

What to do next

You may optionally configure user role mapping for SSO users; refer to Configure User Role Mapping for Okta in the Firewall Management Center, on page 178. If you choose not to configure role mapping, all SSO users who log into the Firewall Management Center are assigned the user role set in step 4 of Configure User Role Mapping for Okta in the Firewall Management Center, on page 178.

Configure User Role Mapping for Okta in the Firewall Management Center

The fields to configure for user role mapping in the Firewall Management Center web interface are the same regardless of your choice of SSO provider. However, the values you configure must take into account how the SAML SSO provider you use implements user role mapping.

Before you begin

- Review the Okta user group mapping information; refer to Review the Okta Org, on page 175.
- Configure the Firewall Management Center as an SSO service provider application; refer to Configure the Firewall Management Center service provider application for Okta, on page 175.
- Enable and configure single sign-on at the Firewall Management Center; refer to Enable SSO at the Firewall Management Center, on page 173, and Configure the Firewall Management Center for Okta SSO, on page 177.

Procedure

Step 1 Choose **System** (⚙) > **Users** > **Single Sign-On**.

Step 2 Expand **Advanced Configuration (Role Mapping)**.

Step 3 From the **Default User Role** drop-down list, choose a default Firewall Management Center user role to assign users.

Step 4 In the **Group Member Attribute** field, enter an attribute configured in Okta for Firewall Management Center user role mapping for users or groups. (Refer to Step 1 of Configure a User Attribute for Role Mapping at the Okta IdP, on page 180 or Step 1 of Configure a group attribute for role mapping at the Okta IdP, on page 181.)

Step 5 Next to each Firewall Management Center user role you wish to assign to SSO users, enter a regular expression. (The Firewall Management Center uses a restricted version of Google's RE2 regular expression standard supported by Golang and Perl.) The Firewall Management Center compares these values against the user role mapping attribute value the IdP sends to the Firewall Management Center with SSO user information. The Firewall Management Center grants users a union of all the roles for which a match is found.

Step 6 Click **Test Configuration**. If the system displays an error message, review the SSO configuration for the Firewall Management Center as well as the identity service provider application configuration, correct any errors, and try again.

Step 7 When the system reports a successful configuration test, click **Apply**.

What to do next

Configure user role mapping at the service provider application; refer to Okta IdP user role mappings, on page 179.

Okta IdP user role mappings

An Okta IdP user role mapping is an SSO configuration method that

- assigns user roles based on individual user or group permissions in the Okta Classic UI Admin Console
- compares user or group attribute values from Okta with regular expressions assigned to each Firewall Management Center user role, and
- grants all matching roles to the user, or a default role if no match is found.

User role mapping configuration details

You can configure SSO user role mapping at the Okta Classic UI Admin Console based on individual user permissions or group permissions.

You can set up this mapping in two ways.

- To map based on individual user permissions, refer to Configure a User Attribute for Role Mapping at the Okta IdP, on page 180.
- To map based on group permissions, refer to Configure a group attribute for role mapping at the Okta IdP, on page 181.

When an SSO user logs in to the Firewall Management Center, Okta presents a user or group role attribute value configured at the Okta IdP, to the Firewall Management Center. The Firewall Management Center compares that attribute value with the regular expressions assigned to each Firewall Management Center user role in the SSO configuration, and grants the user all the roles for which a match is found.. If no match is found, the Firewall Management Center grants the user a configurable default user role. The expression you assign to each Firewall Management Center user role must comply with the restricted version of Google's RE2 regular expression standard supported by Golang and Perl. The Firewall Management Center treats the attribute value received from Okta as a regular expression using that same standard for purposes of comparison with the Firewall Management Center user role expressions.

Note A single Firewall Management Center cannot support role mapping for both groups and individual users; you must choose one mapping method for the Firewall Management Center service provider application and use it consistently. Furthermore, the Firewall Management Center can support group role mapping using only one group attribute statement per Firewall Management Center service provider application configured in Okta. Generally group-based role mapping is more efficient for a Firewall Management Center with many users. You should take into account user and group definitions established throughout your Okta org.

Configure a User Attribute for Role Mapping at the Okta IdP

Use these instructions at the Okta Classic UI Admin Console to add a custom role mapping attribute to the Okta default user profile.

Okta service provider applications may use one of two types of user profiles:

- Okta user profiles, which can be extended with any custom attribute.
- App user profiles, which can be extended only with attributes from a predefined list that Okta generates by querying a third-party application or directory (such as Active directory, LDAP, or Radius) for supported attributes.

You may use either type of user profile in your Okta org; consult Okta documentation for information on how to configure them. Whichever type of user profile you use, to support user role mapping with the Firewall Management Center you must configure a custom attribute in the profile to convey each user's role mapping expression to the Firewall Management Center.

This documentation describes role mapping using Okta user profiles; mapping with App profiles requires familiarity with the third-party user management application in use at your organization to set up custom attributes. See the Okta documentation for details.

Before you begin

- Configure a Firewall Management Center service provider application at the Okta IdP as described in Configure the Firewall Management Center service provider application for Okta, on page 175.
- Configure SSO user role mapping at the Firewall Management Center as described in Configure User Role Mapping for Okta in the Firewall Management Center, on page 178.

Procedure

Step 1 Add a new attribute to the default Okta user profile:

- For **Data type** choose `string`.
- Provide the **Variable name** the Okta IdP will send to the Firewall Management Center, containing an expression to match for user role mapping. This variable name must match the string you entered at the Firewall Management Center SSO configuration for **Group Member Attribute**. (See Step 5 in Configure User Role Mapping for Okta in the Firewall Management Center, on page 178.)

Step 2 For each user assigned to the Firewall Management Center service provider application using this profile, assign a value to the user role attribute you have just created.

Use an expression to represent the role or roles the Firewall Management Center will assign to the user. The Firewall Management Center compares this string against the expressions you assigned to each Firewall Management Center user role in Step 6 of Configure User Role Mapping for Okta in the Firewall Management Center, on page 178. (For purposes of comparison with the Firewall Management Center user role expressions, the Firewall Management Center treats the attribute value received from Okta as an expression complying with the restricted version of Google's RE2 regular expression standard supported by Golang and Perl.)

Configure a group attribute for role mapping at the Okta IdP

Use these instructions at the Okta Admin Console to add a custom role mapping group attribute to the Firewall Management Center service provider application. The Firewall Management Center can support group role mapping using only one group attribute statement per Okta Firewall Management Center service provider application.

Okta service provider applications may use one of two types of groups:

- Okta groups, which can be extended with any custom attribute.
- Application groups, which can be extended only with attributes from a predefined list that Okta generates by querying a third-party application or directory (such as Active directory, LDAP, or Radius) for supported attributes.

You may use either type of group in your Okta org; consult Okta documentation for information on how to configure them. Whichever type of group you use, to support user role mapping with the Firewall Management Center you must configure a custom attribute for the group to convey its role mapping expression to the Firewall Management Center.

This documentation describes role mapping using Okta groups; mapping with application groups requires familiarity with the third-party user management application in use at your organization to set up custom attributes. Refer to the Okta documentation for details.

Before you begin

- Configure a Firewall Management Center service provider application at the Okta IdP; refer to Configure the Firewall Management Center service provider application for Okta, on page 175.
- Configure user role mapping at the Firewall Management Center; Configure User Role Mapping for Okta in the Firewall Management Center, on page 178.

Procedure

Create a new SAML group attribute for the Firewall Management Center service provider application:

- For **Name**, use the same string you entered at the Firewall Management Center SSO configuration for **Group Member Attribute**. (Refer to Step 4 in Configure User Role Mapping for Okta in the Firewall Management Center, on page 178.)
- For **Filter**, specify an expression to represent the role or roles the Firewall Management Center will assign to the members of the group. Okta compares this value against the names of the groups of which a user is a member, and sends the Firewall Management Center the group names that match. The Firewall Management Center in turn compares those group names against the regular expressions you assigned to each Firewall Management Center user role in Step 5 of Configure User Role Mapping for Okta in the Firewall Management Center, on page 178.

Okta user role mapping examples

As this examples demonstrate, the SSO configurations at the Firewall Management Center to support user role mapping are the same for both individual users and for groups. The difference lies in the settings at the Firewall Management Center service provider application in Okta.

Note You can configure Firewall Management Center roles to be mapped based on individual user permissions or based on group permissions, but a single Firewall Management Center application cannot support role mapping for both groups and individual users. Furthermore, the Firewall Management Center can support group role mapping using only one group attribute statement per Firewall Management Center service provider application configured in Okta.

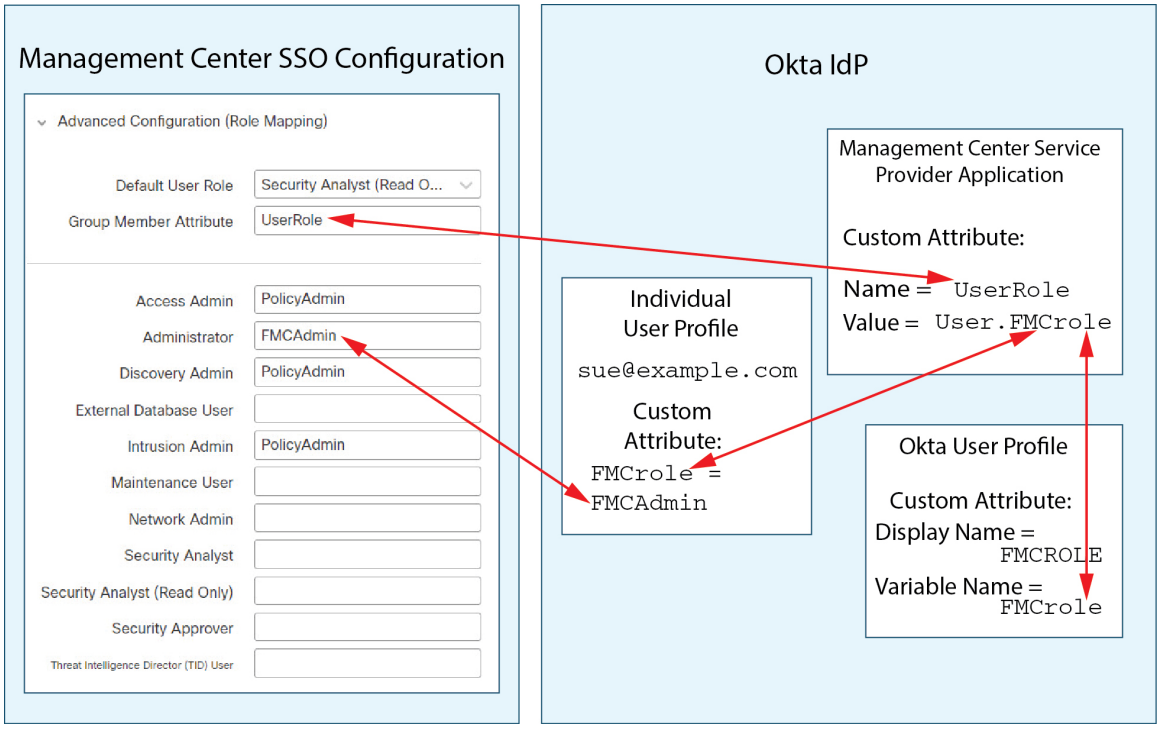
Okta role mapping example for individual user accounts

In role mapping for individual users, the Okta Firewall Management Center service application has a custom attribute whose name matches the name of the Group Member Attribute on the Firewall Management Center. (In this example, `UserRole`). The user profile in Okta also has a custom attribute (in this example, a variable named `FMCrole`.) The definition for the application custom attribute `UserRole` establishes that when Okta passes user role mapping information to the Firewall Management Center, it will use the custom attribute value assigned for the user in question.

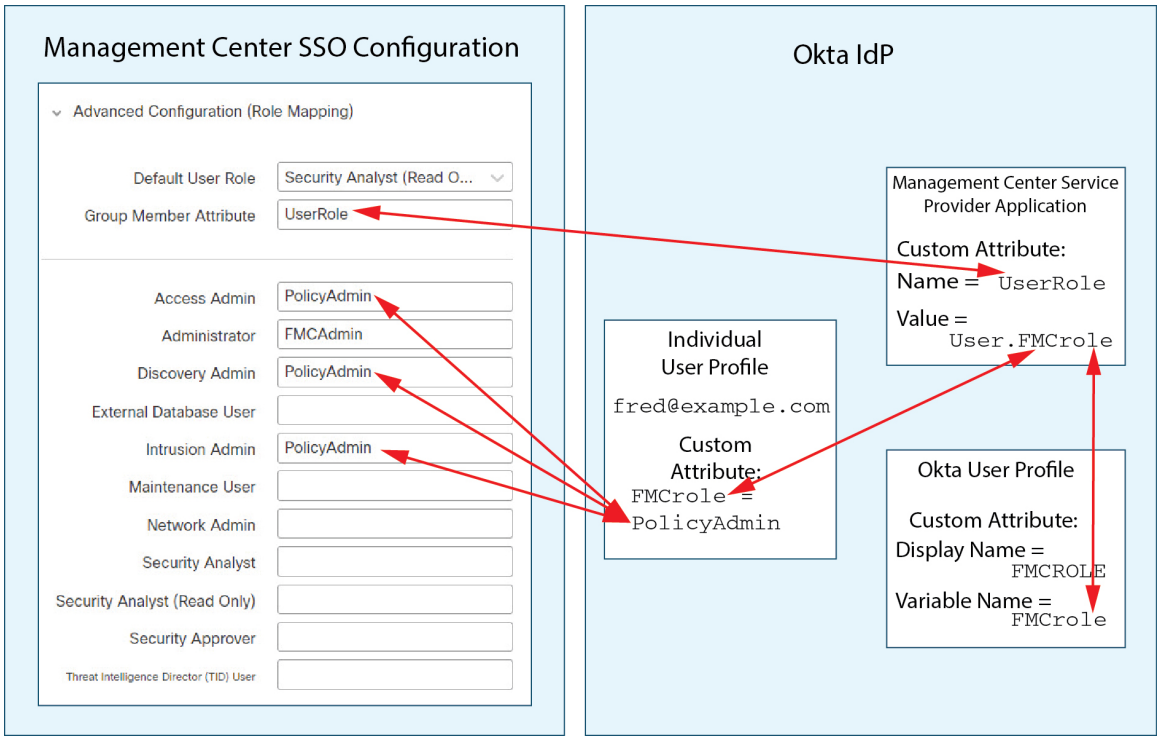
These diagrams illustrate how the relevant fields and values in the Firewall Management Center and Okta configurations correspond to each other in user role mapping for individual accounts. Each diagram uses the same SSO configurations at the Firewall Management Center and at the Okta UI Admin Console, but the

configuration for each user at the Okta UI Admin Console differs to assign each user different roles at the Firewall Management Center.

- In this diagram sue@example.com uses the `FMCrole` value `FMCAdmin` and the Firewall Management Center assigns her the Administrator role.



- In this diagram fred@example.com uses the `FMCrole` value `PolicyAdmin`, and the Firewall Management Center assigns him the roles Access Admin, Discovery Admin, and Intrusion Admin.



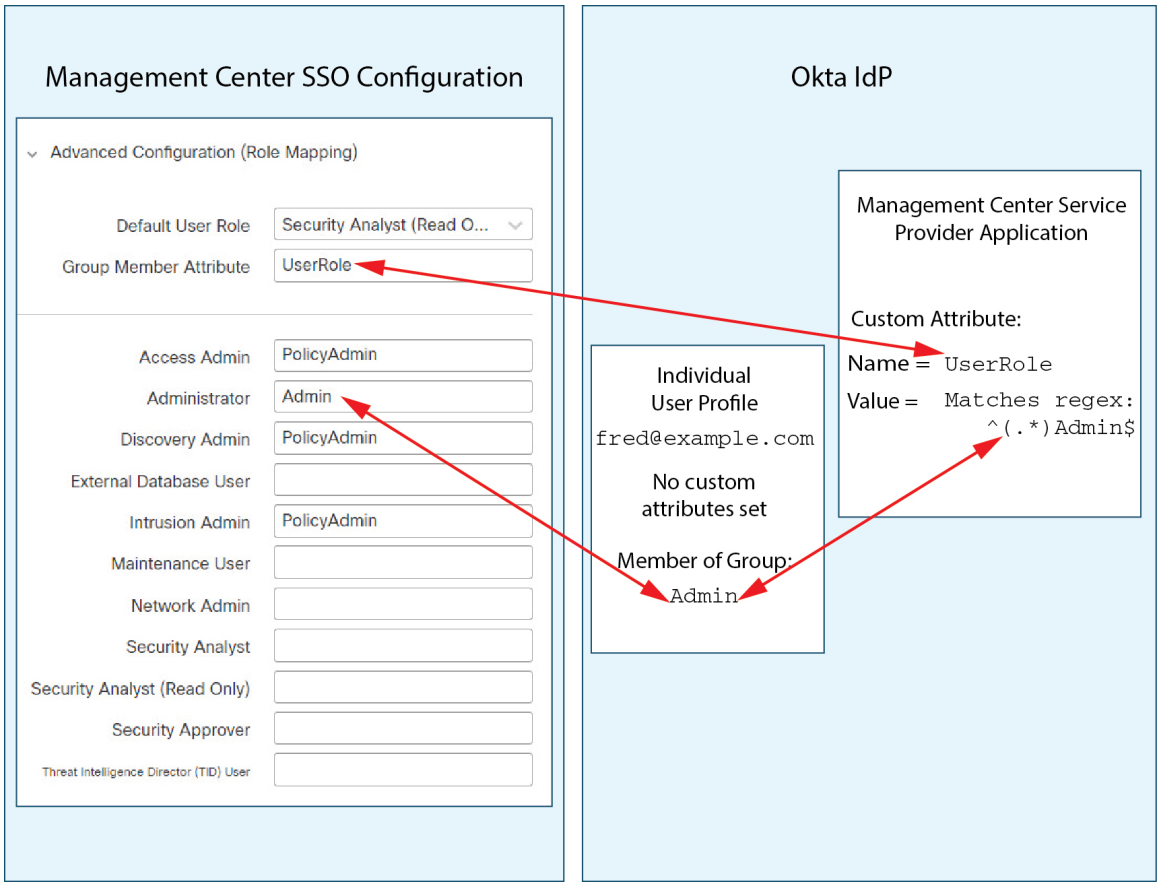
- Other users assigned to the Okta service application for this Firewall Management Center are assigned the default user role Security Analyst (Read Only) for one of these reasons:
  - They have no value assigned to the `FMCrole` variable in their Okta user profile.
  - The value assigned to the `FMCrole` variable in their Okta user profile does not match any expression configured for a user role in the SSO configuration at the Firewall Management Center.

Okta role mapping example for groups

In role mapping for groups, the Okta Firewall Management Center service application has a custom group attribute whose name matches the name of the **Group Member Attribute** on the Firewall Management Center (in this example, `UserRole`). When Okta processes a request for Firewall Management Center SSO login, it compares the user's group membership against the expression assigned to the Firewall Management Center service application group attribute (in this case `^(.*)Admin$`). Okta sends to the Firewall Management Center the user's group membership(s) that match the group attribute. The Firewall Management Center compares the group names it receives against the regular expressions it has configured for each user role, and assigns user roles accordingly.

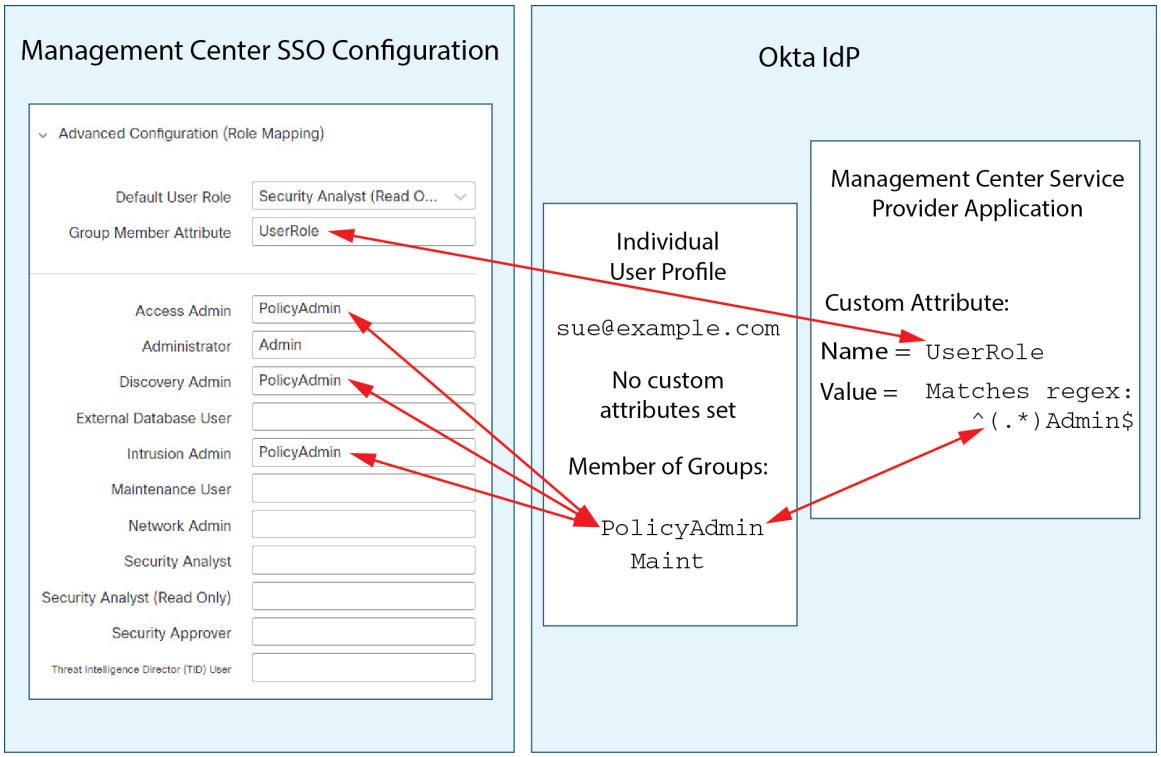
These diagrams illustrate how the relevant fields and values in the Firewall Management Center and Okta configurations correspond to each other in user role mapping for groups. Each diagram uses the same SSO configurations at the Firewall Management Center and at the Okta UI Admin Console, but the configuration for each user at the Okta UI Admin Console differs to assign each user different roles at the Firewall Management Center.

- In this diagram fred@example.com is a member of the Okta IdP group `Admin` , which matches the expression `^(.*)Admin$`. Okta sends the Firewall Management Center Fred's `Admin` group membership, and the Firewall Management Center assigns him the Administrator role.

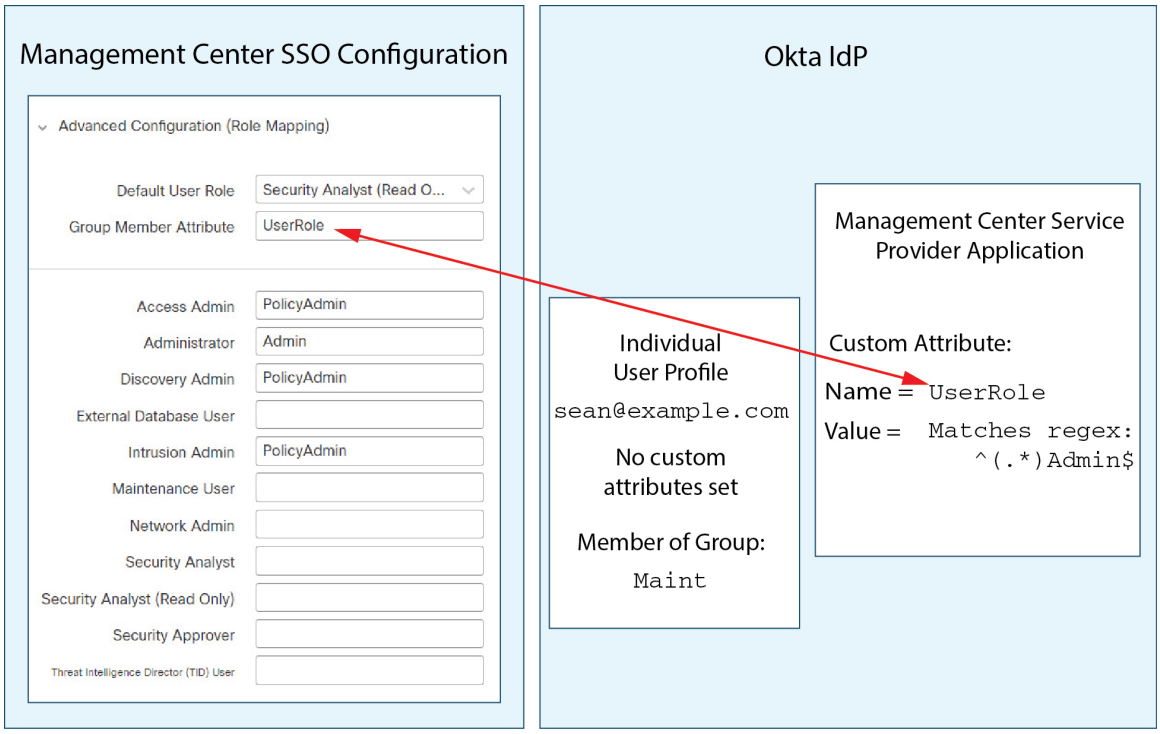


- In this diagram sue@example.com is a member of the Okta IdP group `PolicyAdmin` , which matches the expression `^(.*)Admin$`. Okta sends the Firewall Management Center Sue's `PolicyAdmin` group membership, and the Firewall Management Center assigns her the roles Access Admin, Discovery Admin, and Intrusion Admin.

  Sue is also a member of the Okta group `Maint` , but because this group name does not match the expression assigned to the group membership attribute in the Okta Firewall Management Center service application, Okta does not send information about Sue's `Maint` group membership to the Firewall Management Center , and her membership in the `Maint` group plays no part in the roles the Firewall Management Center assigns to her.



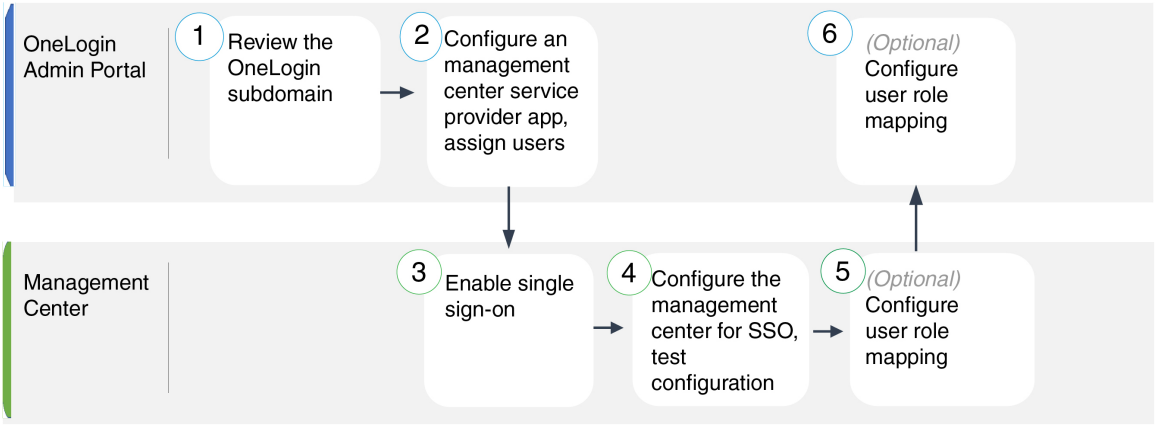
- In this diagram sean@example.com is a member of the Okta IdP group `Maint`. This group name does not match the expression `^(.*)Admin$` , so, when sean@example.com logs into the Firewall Management Center , Okta does not send information about Sean's `Maint` group membership to the Firewall Management Center and Sean is assigned the default user role (Security Analyst (Read Only)) rather than the Maintenance User role.



These diagrams illustrate the importance of advance planning when establishing a role mapping strategy. In this example, any Okta user with access to this Firewall Management Center who is a member of only the `Maint` group can be assigned only the default user role. The Firewall Management Center supports using only one custom group attribute in its Okta Service Application configuration. The expression you assign to that attribute and the group names you establish to match against it must be carefully crafted. You can add more flexibility to role mapping by using regular expressions in the user role assignment strings in the Firewall Management Center SSO configuration. (The expression you assign to each Firewall Management Center user role must comply with the restricted version of Google's RE2 regular expression standard supported by Golang and Perl.)

Configure SSO with OneLogin

Refer to these tasks to configure SSO using OneLogin:



1	Firewall Management Center	Review the OneLogin subdomain, on page 188
2	Firewall Management Center	Configure a Firewall Management Center service provider application for OneLogin, on page 188
3	OneLogin Admin Portal	Enable SSO at the Firewall Management Center, on page 173
4	OneLogin Admin Portal	Configure the Firewall Management Center for OneLogin SSO, on page 190
5	OneLogin Admin Portal	Configure User Role Mapping for OneLogin in the Firewall Management Center , on page 191
6	Firewall Management Center	Configure User Role Mapping at the OneLogin IdP, on page 192

Review the OneLogin subdomain

In OneLogin, the entity that encompasses all the federated devices and applications that a user can access with the same SSO account is called a subdomain. Before adding the Firewall Management Center to a OneLogin subdomain, be familiar with its configuration; consider these questions:

- How many users will have access to the Firewall Management Center ?
- Are users within the OneLogin subdomain members of groups?
- Are users and groups from a third-party directory such as Active Directory, Google Apps, or LDAP synchronized with the OneLogin subdomain?
- Do you need to add more users or groups to the OneLogin subdomain to support SSO on the Firewall Management Center?
- What kind of Firewall Management Center user role assignments do you want to make? (If you choose not to assign user roles, the Firewall Management Center automatically assigns a configurable default user role to all SSO users.)
- How must users and groups within the OneLogin subdomain be organized to support the required user role mapping?

Keep in mind that you can configure Firewall Management Center roles to be mapped based on individual users or based on groups, but a single Firewall Management Center application cannot support role mapping for both groups and individual users.

This documentation assumes you are already familiar with the OneLogin Admin Portal, and have an account with Super User privilege. To configure user role mapping, you will also need a subscription to the OneLogin Unlimited plan, which supports Custom User Fields. If you need more information, refer to the OneLogin documentation available online.

Configure a Firewall Management Center service provider application for OneLogin

This task guides you to configure a Firewall Management Center service provider application in OneLogin, enabling SAML SSO integration for your organization. Establish secure SSO access for users through OneLogin.

Use these instructions at the OneLogin Admin Portal to create a Firewall Management Center service provider application within OneLogin and assign users or groups to that application. You should be familiar with SAML SSO concepts and the OneLogin Admin Portal. This documentation does not describe all the OneLogin functions required to establish a fully functional SSO organization. For example, refer to the OneLogin documentation to learn how to create users and groups, or to import user and group definitions from another user management application.

- If you plan to assign user groups to the Firewall Management Center application, do not also assign users within those groups as individuals.
- The Firewall Management Center cannot support role mapping using multiple SSO attributes; you must select either user role mapping or group role mapping and configure a single attribute to convey user role information from OneLogin to the Firewall Management Center.

Before you begin

- Familiarize yourself with the OneLogin subdomain and its users and groups; refer to Review the OneLogin subdomain, on page 188.
- Create user accounts in your OneLogin subdomain if necessary.

Note The system requires that user names for SSO accounts as well as the NameID attribute the IdP sends to the Firewall Management Center during the SAML login process must be both be valid email addresses. Many IdP's automatically use the username of the user trying to logon as the NameID attribute, but you should confirm this is the case for your IdP. Keep this in mind when configuring a service provider application at your IdP and creating IdP user accounts that are to be granted SSO access to the Firewall Management Center.

- Confirm the login URL for the target Firewall Management Center (for example, `https://ipaddress_or_hostname/`).

Note If your Firewall Management Centerweb interface can be reached with multiple URLs, such as a fully qualified domain name or an IP address, SSO users must consistently access the Firewall Management Center using the login URL that you configure in this task.

Follow these steps to configure a Firewall Management Center service provider application for OneLogin:

Procedure

Step 1 Create the Firewall Management Center service provider application using the **SAML Test Connector (Advanced)** as its basis.

Step 2 Configure the application with these settings:

- For the **Audience (Entity ID)**, append the string `/SAML/metadata` to the Firewall Management Center login URL. For example: `https://ExampleFMC/SAML/metadata`.

- For **Recipient**, append the string `/SAML/ACS` to the Firewall Management Center login URL. For example: `https://ExampleFMC/SAML/ACS`.
- For **ACS (Consumer) URL Validator**, enter an expression that OneLogin uses to confirm it is using the correct Firewall Management Center URL. You can create a simple validator by using the ACS URL and altering it as follows:
  - Append a `^` to the beginning of the ACS URL.
  - Append a `$` to the end of the ACS URL.
  - Insert a `\` preceding every `/` and `?` within the ACS URL.

For example, for the ACS URL `https://ExampleFMC/SAML/ACS`, an appropriate URL validator would be `^https:\/\/ExampleFMC\/SAML\/ACS$`.

- For **ACS (Consumer) URL**, append the string `/SAML/ACS` to the Firewall Management Center login URL. For example: `https://ExampleFMC/SAML/ACS`.
- For **Login URL**, append the string `/SAML/ACS` to the Firewall Management Center login URL. For example: `https://ExampleFMC/SAML/ACS`.
- For the **SAML Initiator**, choose `Service Provider`.

Step 3 Assign OneLogin users to the Firewall Management Center service provider application.

Step 4 (Optional) To make SSO setup at the Firewall Management Center easier, you can download the SAML XML metadata for the Firewall Management Center service provider application from OneLogin to your local computer.

After completing this task, your Firewall Management Center will be integrated with OneLogin for SAML-based SSO, allowing assigned users to authenticate securely using their OneLogin credentials.

What to do next

Enable single sign-on; refer to Enable SSO at the Firewall Management Center, on page 173.

Configure the Firewall Management Center for OneLogin SSO

Configure the Firewall Management Center for OneLogin SSO to integrate single sign-on authentication, allowing users to access the system securely using their OneLogin credentials.

Use these instructions at the Firewall Management Center web interface.

This task is relevant when you want to enable SSO for your organization using OneLogin as the identity provider for the Firewall Management Center. Ensure you have administrative access to both the OneLogin Admin Portal and the Firewall Management Center.

Before you begin

- Create a Firewall Management Center service provider application at the OneLogin Admin Portal; refer to Configure a Firewall Management Center service provider application for OneLogin, on page 188.
- Enable single sign-on; refer to Enable SSO at the Firewall Management Center, on page 173.

Follow these steps to configure the Firewall Management Center for OneLogin SSO:

Procedure

Step 1 (This step continues directly from Enable SSO at the Firewall Management Center, on page 173.) At the **Configure OneLogin Metadata** dialog, you have two choices:

- To enter the SSO configuration information manually:
  a. Click the **Manual Configuration** radio button.
  b. Enter these SSO configuration values from the OneLogin service provide application:
     - **Identity Provider Single Sign-On URL**: Enter the **SAML 2.0 Endpoint (HTTP)** from OneLogin.
     - **Identity Provider Issuer**: Enter the **Issuer URL** from OneLogin.
     - **X.509 Certificate**: Enter the **X.509 Certificate** from OneLogin.
- If you have saved the XML metadata file generated by OneLogin to your local computer (step 4 in Configure a Firewall Management Center service provider application for OneLogin, on page 188), upload the file to the Firewall Management Center:
  a. Click the **Upload XML File** radio button.
  b. Follow the on-screen instructions to navigate to and choose the XML metadata file on your local computer.

Step 2 Click **Next**.

Step 3 At the **Verify Metadata** dialog, review the configuration parameters and click **Save**.

Step 4 Click **Test Configuration**. If the system displays an error message, review the SSO configuration for the Firewall Management Center and the OneLogin service provider application configuration. Correct any errors, then try again.

Step 5 When the system reports a successful configuration test, click **Apply**.

After completing this task, the Firewall Management Center is integrated with OneLogin SSO. Users can authenticate using their OneLogin credentials for secure access.

What to do next

You may optionally configure user role mapping for SSO users; refer to Configure User Role Mapping for OneLogin in the Firewall Management Center , on page 191. If you choose not to configure role mapping, all SSO users who log in to the Firewall Management Center are assigned the user role during user role mapping configuration (step 4 of Configure User Role Mapping for OneLogin in the Firewall Management Center , on page 191).

Configure User Role Mapping for OneLogin in the Firewall Management Center

The fields to configure for user role mapping at the Firewall Management Center web interface are the same regardless of your choice of SSO provider. But the values you configure must take into account how the SAML SSO provider you use implements user role mapping.

Before you begin

- Review the OneLogin users and groups, refer to Review the OneLogin subdomain, on page 188.
- Configure an SSO service provider application for the Firewall Management Center; refer to Configure a Firewall Management Center service provider application for OneLogin, on page 188.
- Enable and configure single sign-on at the Firewall Management Center; refer to Enable SSO at the Firewall Management Center, on page 173, and Configure a Firewall Management Center service provider application for OneLogin, on page 188.

Procedure

Step 1 Choose **System** (⚙) > **Users** > **Single Sign-On**.

Step 2 Expand **Advanced Configuration (Role Mapping)**.

Step 3 From the **Default User Role** drop-down list, choose a default Firewall Management Center user role to assign users.

Step 4 In the **Group Member Attribute** field, enter an attribute configured in OneLogin for Firewall Management Center user role mapping for users or groups. Refer to Step 1 of Configure User Role Mapping for Individual Users at the OneLogin IdP, on page 193 or Step 1 of Configure User Role Mapping for Groups at the OneLogin IdP, on page 194.

Step 5 Next to each Firewall Management Center user roll you wish to assign to SSO users, enter a regular expression. The Firewall Management Center compares these values against the user role mapping attribute the IdP sends to the Firewall Management Center with SSO user information. The Firewall Management Center grants users a union of all the roles for which a match is found.

Step 6 Click **Test Configuration**. If the system displays an error message, review the SSO configuration for the Firewall Management Center as well as the identity service provider application configuration, correct any errors, and try again.

Step 7 When the system reports a successful configuration test, click **Apply**.

What to do next

Configure user role mapping at the service provider application; refer to Configure User Role Mapping at the OneLogin IdP, on page 192.

Configure User Role Mapping at the OneLogin IdP

You can configure SSO user role mapping at the Onelogin Admin portal based on individual permissions or based on group permissions.

- To map based on individual user permissions, see Configure User Role Mapping for Individual Users at the OneLogin IdP, on page 193.
- To map based on group permissions, see Configure User Role Mapping for Groups at the OneLogin IdP, on page 194.

When an SSO user logs into the Firewall Management Center, OneLogin presents to the Firewall Management Center a user or group role attribute value that gets its value from a custom user field configured at the

OneLogin IdP. The Firewall Management Center compares that attribute value against the regular expressions assigned to each Firewall Management Center user role in the SSO configuration, and grants the user all the roles for which a match is found.. . (If no match is found, the Firewall Management Center grants the user a configurable default user role.) The expression you assign to each Firewall Management Center user role must comply with the restricted version of Google's RE2 regular expression standard supported by Golang and Perl. The Firewall Management Center treats the attribute value received from Okta as a regular expression using that same standard for purposes of comparison with the Firewall Management Center user role expressions.

Note A single Firewall Management Center cannot support role mapping for both groups and individual users; you must choose one mapping method for the Firewall Management Center service provider application and use it consistently. The Firewall Management Center can support role mapping using only one custom user field configured in OneLogin. Generally group-based role mapping is more efficient for a Firewall Management Center with many users. You should take into account user and group definitions established throughout your OneLogin subdomain.

Configure User Role Mapping for Individual Users at the OneLogin IdP

Use the OneLogin Admin Portal to create a custom parameter for the Firewall Management Center service provider application and a custom user field. These provide the means for OneLogin to pass user role information to the Firewall Management Center during the SSO login process.

Before you begin

- Review the OneLogin subdomain and its users and groups; see Review the OneLogin subdomain, on page 188.
- Create and configure a Firewall Management Center service provider application in OneLogin; see Configure a Firewall Management Center service provider application for OneLogin, on page 188.
- Configure SSO user role mapping as described in Configure User Role Mapping for OneLogin in the Firewall Management Center , on page 191.

Procedure

Step 1 Create a custom parameter for the Firewall Management Center service provider application.

- For the **Field Name**, use the same name you used for the **Group Member Attribute** in the Firewall Management Center SSO configuration. (See Step 4 in Configure User Role Mapping for OneLogin in the Firewall Management Center , on page 191.)
- For the **Value**, provide a mnemonic name such as `FMCUserRole`. This must match the name of the customer user field you will configure in Step 2 of this procedure.

Step 2 Create a custom user field to contain user role information for each OneLogin user with access the Firewall Management Center.

- For the field **Name**, provide a mnemonic name such as `FMCUserRole`. This must match the value provided for the application custom parameter described in Step 1 of this procedure.

- For the **Short name**, provide an abbreviated alternate name for the field. (This is used for OneLogin programmatic interfaces.)

Step 3 For each user with access to the Firewall Management Center service provider application, assign a value to the custom user field you created in Step 2 of this procedure.

When a user logs into the Firewall Management Center using SSO, the value you assign to this field for that user is the value the Firewall Management Center compares against the expressions you assigned to Firewall Management Center user roles in the SSO configuration. (See Step 5 in Configure User Role Mapping for OneLogin in the Firewall Management Center , on page 191.)

What to do next

- Test your role mapping scheme by logging into the Firewall Management Center using SSO from various accounts and confirming that users are assigned Firewall Management Center user roles as you expect.

Configure User Role Mapping for Groups at the OneLogin IdP

Use the OneLogin Admin Portal to create a custom parameter for the Firewall Management Center service provider application and a custom user field. Assign OneLogin users to groups. Then create one or more mappings between the custom user field and the user group so OneLogin assigns a value to the custom user field based on the user's group membership. These provide the means for OneLogin to pass group-based user role information to the Firewall Management Center during the SSO login process.

OneLogin service provider applications may use one of two types of groups:

- Groups native to OneLogin.
- Groups synchronized from third-party applications such as Active Directory, Google Apps, or LDAP.

You may user either type of group for Firewall Management Center group role mapping. This documentation describes role mapping using OneLogin groups; using third-party application groups requires familiarity with the third-party user management application in use at your organization. See the OneLogin documentation for details.

Before you begin

- Review the OneLogin subdomain and its users and groups; see Review the OneLogin subdomain, on page 188.
- Create and configure a Firewall Management Center service provider application in OneLogin; see Configure a Firewall Management Center service provider application for OneLogin, on page 188.
- Configure SSO user role mapping as described in Configure User Role Mapping for OneLogin in the Firewall Management Center , on page 191.

Procedure

Step 1 Create a custom parameter for the Firewall Management Center service provider application.

- For the **Field Name**, use the same name you used for the **Group Member Attribute** in the Firewall Management Center SSO configuration. (See Step 4 in Configure User Role Mapping for OneLogin in the Firewall Management Center , on page 191.)
- For the **Value**, provide a mnemonic name such as `FMCUserRole`. This must match the name of the customer user field you will configure in Step 2 of this procedure.

Step 2 Create a custom user field to contain user role information for each OneLogin user with access the Firewall Management Center.

- For the field **Name**, provide a mnemonic name such as `FMCUserRole`. This must match the value provided for the application custom parameter described in Step 1 of this procedure.
- For the **Short name**, provide an abbreviated alternate name for the field. (This is used for OneLogin programmatic interfaces.)

Step 3 Create one or more user field mappings to assign group-based values to the custom user field you created in Step 2 of this procedure. Create as many mappings as you need to assign the correct Firewall Management Center user role to each OneLogin user group.

- Create one or more **Conditions** for the mapping, comparing the user **Group** field against group names.
- If you create multiple **Conditions**, choose whether a user's group must match `any` or `all` of the conditions for the mapping to take place.
- Create an **Action** for the mapping, to assign a value to the custom user field you created in Step 2 of this procedure. Provide the field **Name**, and the string that OneLogin assigns to this custom user field for all users that meet the **Conditions** you specified.

  The Firewall Management Center compares this string against the expressions you assign to each Firewall Management Center user role in Step 5 of Configure User Role Mapping for OneLogin in the Firewall Management Center , on page 191.
- **Reapply All Mappings** when you have completed your changes.

What to do next

- Test your role mapping scheme by logging into the Firewall Management Center using SSO from various accounts and confirming that users are assigned Firewall Management Center user roles as you expect.

OneLogin User Role Mapping Examples

As the following examples demonstrate, the SSO configurations at the Firewall Management Center to support user role mapping are the same for both individual users and for groups. The difference lies in the settings at the Firewall Management Center service provider application in OneLogin.

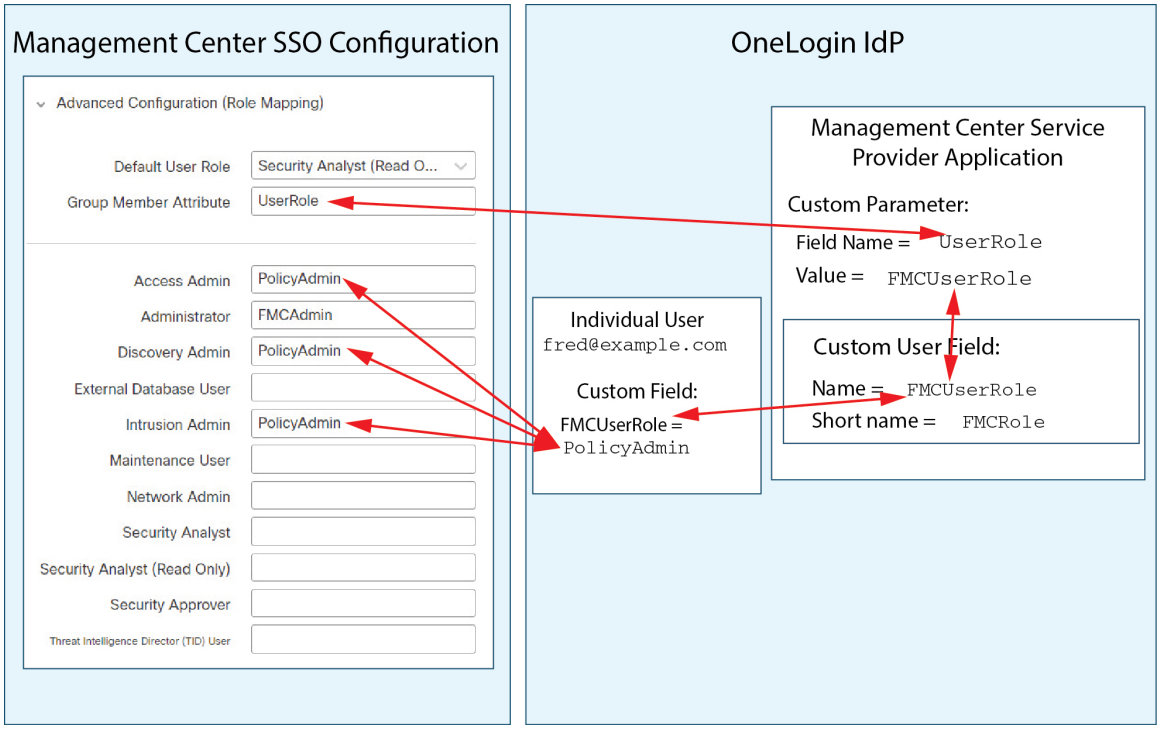
Note A single Firewall Management Center cannot support role mapping for both groups and individual users; you must choose one mapping method for the Firewall Management Center service provider application and use it consistently. The Firewall Management Center can support role mapping using only one custom user field configured in OneLogin. Generally group-based role mapping is more efficient for a Firewall Management Center with many users. You should take into account user and group definitions established throughout your OneLogin subdomain.

OneLogin Role Mapping Example for Individual User Accounts

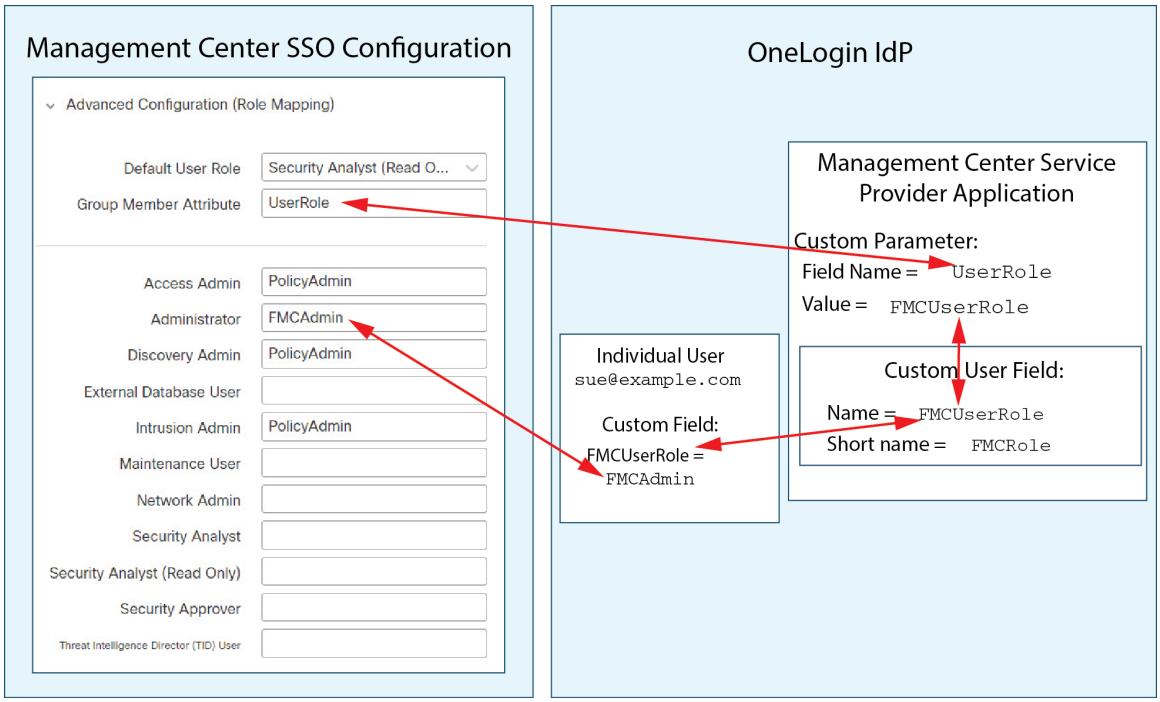
In role mapping for individual users, the OneLogin Firewall Management Center service application has a custom parameter whose name matches the name of the Group Member attribute on the Firewall Management Center (in this example, `UserRole`). OneLogin also has a custom user field defined (in this example, `FMCUserRole`). The definition for the application custom parameter `UserRole` establishes that when OneLogin passes user role mapping information to the Firewall Management Center, it will use the value of the custom user field `FMCUserRole` for the user in question.

The following diagrams illustrate how the relevant fields and values in the Firewall Management Center and OneLogin configurations correspond to each other in user role mapping for individual accounts. Each diagram uses the same SSO configurations at the Firewall Management Center and at the OneLogin Admin portal, but the configuration for each user at the OneLogin Admin portal differs to assign each user different roles at the Firewall Management Center.

- In this diagram fred@example.com uses the `FMCUserRole` value `PolicyAdmin` and the Firewall Management Center assigns him the roles Access Admin, Discovery Admin, and Intrusion Admin.



- In this diagram sue@example.com uses the `FMCUserRole` value `FMCAdmin`, and the Firewall Management Center assigns her the Administrator role.



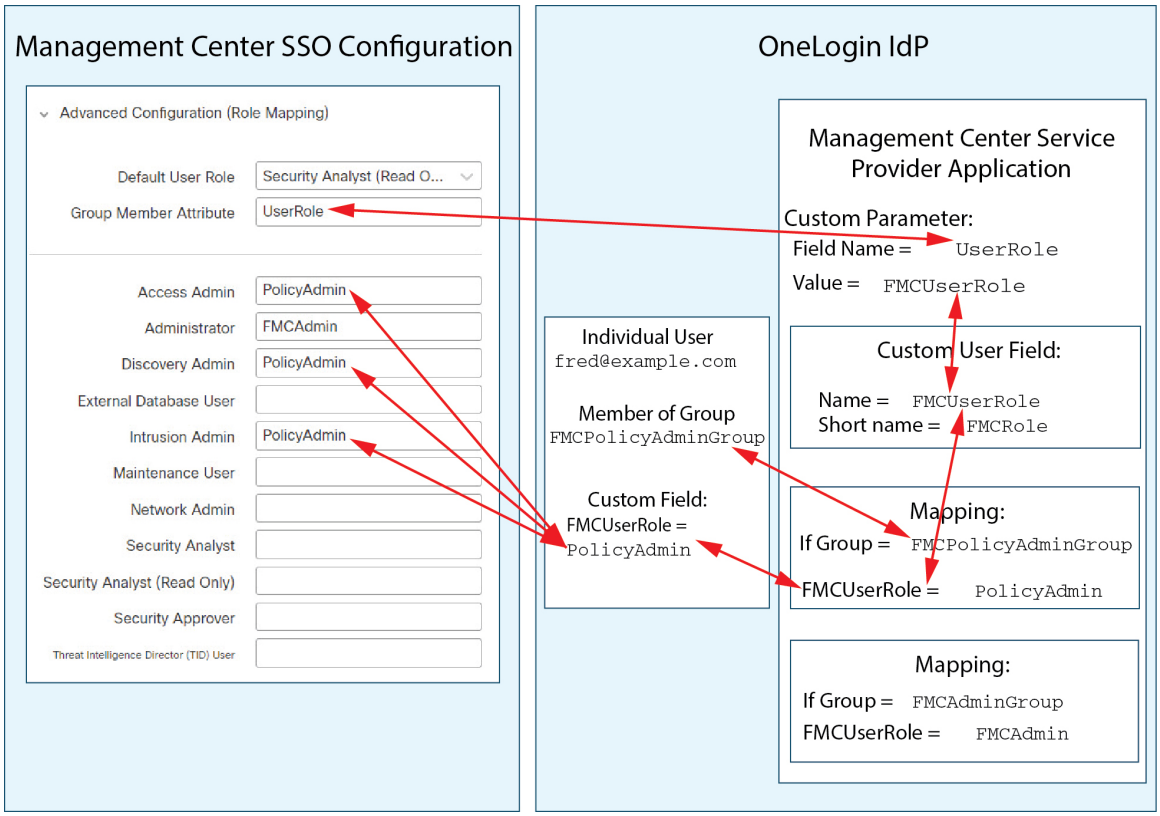
- Other users assigned to the OneLogin service application for this Firewall Management Center are assigned the default user role Security Analyst (Read Only) for one of the following reasons:
  - They have no value assigned to the `FMCUserRole` custom user field.
  - The value assigned to the `FMCUserRole` custom user field does not match any expression configured for a user role in the SSO configuration at the Firewall Management Center.

OneLogin Role Mapping Example for Groups

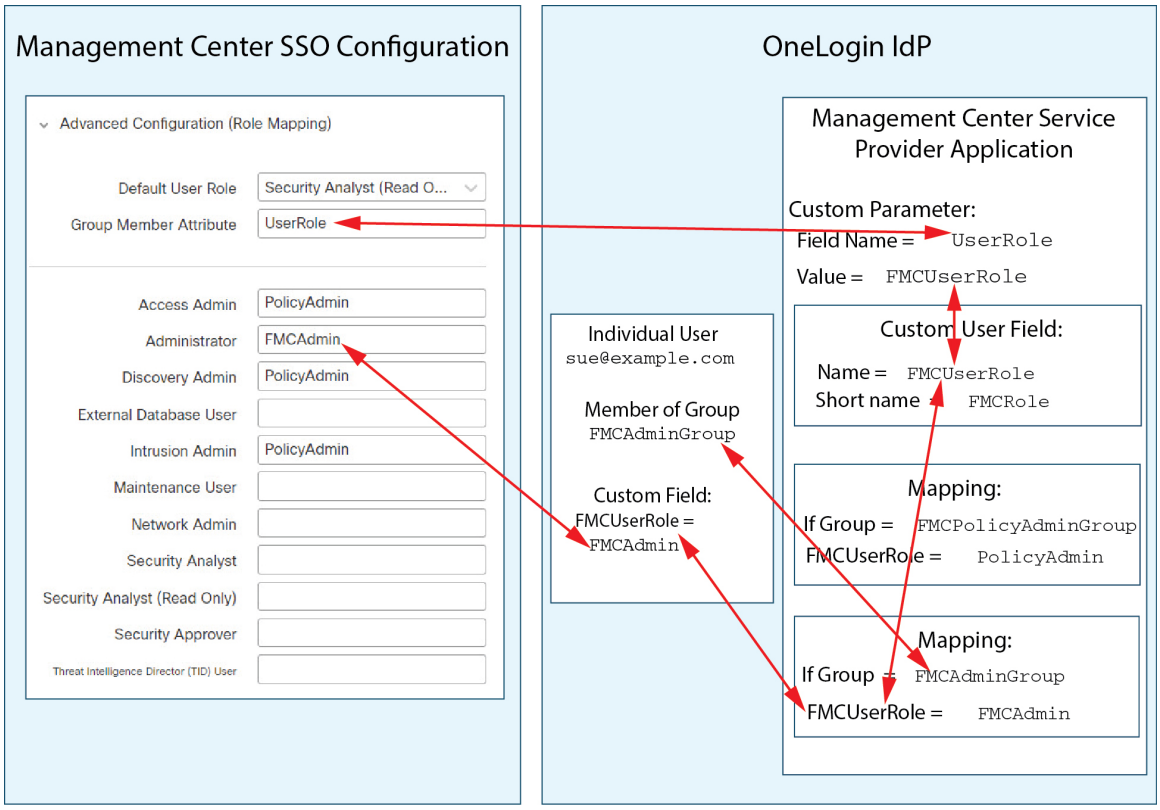
In role mapping for groups, the OneLogin Firewall Management Center service application has a has a custom parameter whose name matches the name of the Group Member attribute on the Firewall Management Center (in this example, `UserRole`). OneLogin also has a custom user field defined (in this example, `FMCUserRole`). The definition for the application custom parameter `UserRole` establishes that when OneLogin passes user role mapping information to the Firewall Management Center, it will use the value of the custom user field `FMCUserRole` for the user in question. To support user group mapping, you must establish a mapping within OneLogin to assign a value for each user's `FMCUserRole` field based on that user's OneLogin group membership.

The following diagrams illustrate how the relevant fields and values in the Firewall Management Center and OneLogin configurations correspond to each other in user role mapping for groups. Each diagram uses the same SSO configurations at the Firewall Management Center and at the OneLogin Admin portal, but the configuration for each user at the OneLogin Admin portal differs to assign each user different roles at the Firewall Management Center.

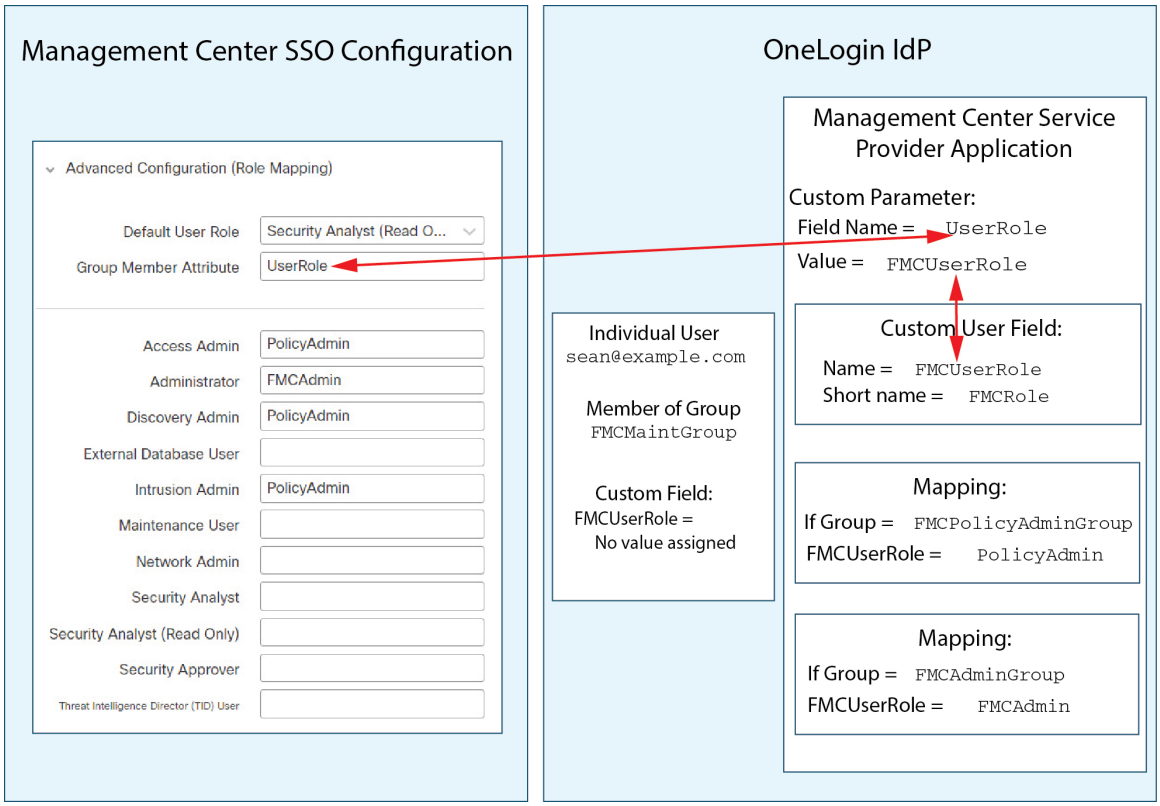
- In this diagram fred@example.com is a member of the OneLogin IdP group `FMCPolicyAdminGroup`. A OneLogin mapping assigns the value `PolicyAdmin` to the custom user field `FMCUserRole` for members of the `FMCPolicyAdminGroup`. The Firewall Management Center assigns Fred and other members of the `FMCPolicyAdminGroup` the roles Access Admin, Discovery Admin, and Intrusion Admin.



- In this diagram sue@example.com is a member of the OneLogin IdP group `FMCAdminGroup`. A OneLogin mapping assigns the value FMCAdmin to the custom user field `FMCUserRole` for members of the `FMCAdminGroup`. The Firewall Management Center assigns Sue and other members of the `FMCAdminGroup` the Administrator role.

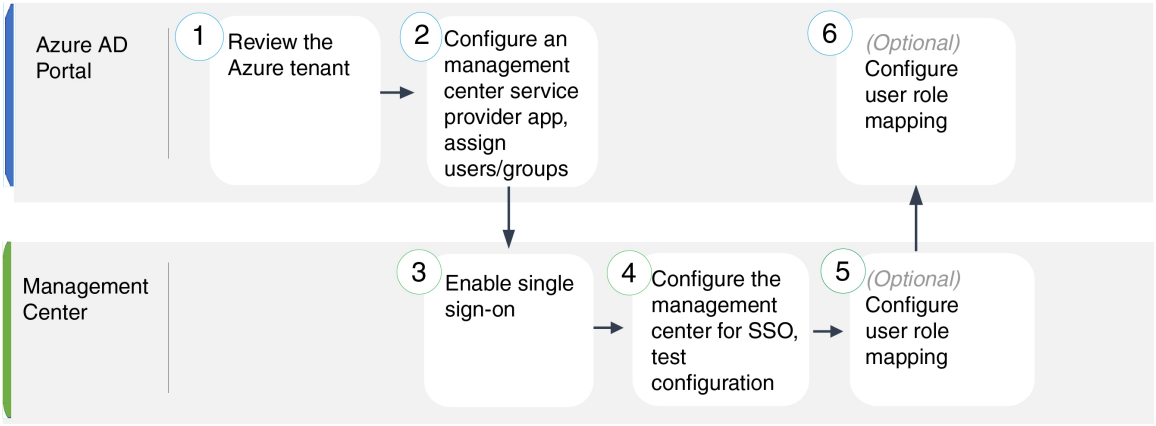


- In this diagram sean@example.com is a member of the Idp group `FMCMaintGroup`. There is no OneLogin mapping associated with this group, so OneLogin does not assign a value to the custom user field `FMCUserRole` for Sean. The Firewall Management Center assigns Sean the default user role (Security Analyst (Read Only)) rather than the Maintenance User role.



Configure SSO with Azure AD

This reference provides the sequence of tasks required to configure Single Sign-On (SSO) for your Cisco environment using Azure Active Directory (Azure AD).



Step	Location	Task
1	Azure AD Portal	Azure tenants, on page 201

Step	Location	Task
2	Azure AD Portal	Configure the Firewall Management Center service provider application for Azure, on page 202
3	Firewall Management Center	Enable SSO at the Firewall Management Center, on page 173
4	Firewall Management Center	Configure the Firewall Management Center for Azure SSO, on page 203
5	Firewall Management Center	Configure user role mapping for Azure in the Firewall Management Center, on page 205
6	Azure AD Portal	Azure IdP user role mappings, on page 206

Azure tenants

An Azure tenant is a cloud-based identity and access management entity that

- encompasses all federated devices a user can access with the same SSO account
- organizes users and groups within Azure Active Directory, and
- supports integration with applications such as the Firewall Management Center.

Azure tenant organization considerations

Consider these questions before adding the Firewall Management Center to an Azure tenant:

- How many users will have access to the Firewall Management Center?
- Are users within the Azure tenant members of groups?
- Are users and groups from another directory product?
- Do you need to add more users or groups to the Azure tenant to support SSO on the Firewall Management Center?
- What kind of Firewall Management Center user role assignments do you want to make? If you choose not to assign user roles, the Firewall Management Center automatically assigns a configurable default user role to all SSO users.
- How must users and groups within the Azure tenant be organized to support the required user role mapping?
- You can configure Firewall Management Center roles to be mapped based on individual users or based on groups, but a single Firewall Management Center application cannot support role mapping for both groups and individual users.

Before you begin, ensure you are familiar with the Azure Active Directory Portal and that you have an account with application admin privileges for the Azure AD tenant. The Firewall Management Center supports Azure SSO only with tenant-specific single sign-on and single sign-out endpoints. Ensure you have an Azure AD Premium P1 or higher license and Global Administrator permissions. For more information, refer to Azure documentation.

Configure the Firewall Management Center service provider application for Azure

This task enables you to configure a Firewall Management Center service provider application in Azure Active Directory for SAML-based single sign-on (SSO).

Use the Azure Active Directory Portal to create a Firewall Management Center service provider application within your Azure Active Directory tenant and establish basic configuration settings.

- If you plan to assign user groups to the Firewall Management Center application, do not also assign users within those groups as individuals.
- The Firewall Management Center cannot support role mapping using multiple SSO attributes. You must select either user role mapping or group role mapping. Then, configure a single attribute to convey user role information from OneLogin to the Firewall Management Center.

Before you begin

- Familiarize yourself with your Azure tenant and its users and groups; refer to Azure tenants, on page 201.
- Create user accounts or groups, or both, in your Azure tenant if necessary.

Note The system requires that user names for SSO accounts as well as the NameID attribute the IdP sends to the Firewall Management Center during the SAML login process must be both be valid email addresses. Many IdP's automatically use the username of the user trying to logon as the NameID attribute, but you should confirm this is the case for your IdP. Keep this in mind when configuring a service provider application at your IdP and creating IdP user accounts that are to be granted SSO access to the Firewall Management Center.

- Confirm the login URL for the target Firewall Management Center (`https://ipaddress_or_hostname`).

Note If your Firewall Management Center web interface can be reached with multiple URLs (for example, a fully-qualified domain name as well as an IP address), SSO users must consistently access the Firewall Management Center using the login URL that you configure in this task.

Follow these steps to configure the Firewall Management Center service provider application for Azure.

Procedure

Step 1 Create the Firewall Management Center service provider application using the Azure AD SAML Toolkit as its basis.

Step 2 Configure the application with these settings for **Basic SAML Configuration**:

- For the **Identifier (Entity ID)** append the string `/SAML/metadata` to the Firewall Management Center login URL. For example: `https://ExampleFMC/SAML/metadata`.

- For the **Reply URL (Assertion Consumer Service URL)** append the string `/SAML/acs` to the Firewall Management Center login URL. For example: `https://ExampleFMC/SAML/acs`.
- For the **Sign on URL** append the string `/SAML/acs` to the Firewall Management Center login URL. For example: `https://ExampleFMC/SAML/acs`.

Step 3 Edit the **Unique User Identifier Name (Name ID)** claim for the application to force the username for sign-on at the Firewall Management Center to be the email address associated with the user account:

- For **Source**, choose `Attribute`.
- For **Source attribute**, choose `user.mail`.

Step 4 Generate a certificate to secure SSO on the Firewall Management Center. Use these options for the certificate:

- Select Sign SAML Response and Assertion for the Signing Option.
- Select SHA-256 for the Signing Algorithm.

Step 5 Download the Base-64 version of the certificate to your local computer. You will need it when you configure Azure SSO at the Firewall Management Center web interface.

Step 6 In the SAML-based Sign-on information for the application, note these values:

- **Login URL**
- **Azure AD Identifier**

Use these values when you configure Azure SSO at the Firewall Management Center web interface.

Step 7 (Optional) To make SSO setup at the Firewall Management Center easier, you can download the SAML XML metadata file for the Firewall Management Center service provider application (called the **Federation Metadata XML** in the Azure Portal) to your local computer.

Step 8 Assign existing Azure users and groups to the Firewall Management Center service application.

- If you plan to assign user groups to the Firewall Management Center Application, do not also assign users within those groups as individuals.
- If you plan to configure user role mapping, you can configure roles to be mapped based on individual user permissions or based on group permissions, but a single Firewall Management Center application cannot support role mapping for both groups and individual users.

After you complete this task, your Azure AD tenant has a configured service provider application for your Firewall Management Center, which enables SAML-based SSO integration and secure authentication for assigned users and groups.

Configure the Firewall Management Center for Azure SSO

Configure the Firewall Management Center for Azure SSO to integrate with Azure Active Directory and enable single sign-on authentication for users.

Use these instructions at the Firewall Management Center web interface. Perform this configuration when you want to enable SSO for your management center using Azure AD as the identity provider. This is typically done by administrators responsible for secure access management.

Before you begin

- Create a Firewall Management Center service provider application at the Azure AD Portal; refer to Configure the Firewall Management Center service provider application for Azure, on page 202.
- Enable single sign-on; refer to Enable SSO at the Firewall Management Center, on page 173.

Follow these steps to configure the Firewall Management Center for Azure SSO:

Procedure

Step 1 This step continues directly from Enable SSO at the Firewall Management Center, on page 173. At the **Configure Azure Metadata** dialog, you have two choices:

- To enter the SSO configuration information manually:
  a. Click the **Manual Configuration** radio button.
  b. Enter the values you retrieved from the Azure SSO Service Provider application:
     - For **Identity Provider Single Sign-On URL** enter the **Login URL** you noted in Step 6 of Configure the Firewall Management Center service provider application for Azure, on page 202.
     - For **Identity Provider Issuer** enter the **Azure AD Identifier** you noted in Step 6 of Configure the Firewall Management Center service provider application for Azure, on page 202.
     - For the **X.509 Certificate**, use the certificate you downloaded from Azure in Step 5 of Configure the Firewall Management Center service provider application for Azure, on page 202. (Use a text editor to open the certificate file, copy the contents, and paste it into the **X.509 Certificate** field.)
- If you have the XML metadata file generated by Azure on your local computer (step 7 of Configure the Firewall Management Center service provider application for Azure, on page 202), upload the file to the Firewall Management Center:
  a. Click the **Upload XML File** radio button.
  b. Follow the on-screen instructions to navigate to and choose the XML metadata file on your local computer.

Step 2 Click **Next**.

Step 3 At the **Verify Metadata** dialog, review the configuration parameters and click **Save**.

Step 4 Click **Test Configuration**. If the system displays an error message, review the SSO configuration for the Firewall Management Center and the Azure service provider application. Correct any errors and try again.

Step 5 Click **Apply** after the system reports a successful configuration test.

After completing these steps, the Firewall Management Center is configured for Azure SSO. Users can now authenticate using their Azure AD credentials for secure access.

What to do next

You may optionally configure role mapping for SSO users; refer to Configure user role mapping for Azure in the Firewall Management Center, on page 205. If you do not configure role mapping, all SSO users who log in to the Firewall Management Center receive the default user role specified in step 4 of Configure user role mapping for Azure in the Firewall Management Center, on page 205.

Configure user role mapping for Azure in the Firewall Management Center

Configure user role mapping for Azure in the management center to assign roles to users authenticated through SSO. This configuration enables granular access control based on Azure user or group attributes. You can map Azure users and groups to management center roles to ensure secure access management

The fields to configure for user role mapping at the Firewall Management Center web interface are the same regardless of your choice of SSO provider. However, the values you configure must consider the implementation of user role mapping by your SAML SSO provider.

This task is relevant when integrating Azure SSO with the management center to ensure users are assigned the correct roles based on their Azure attributes. Apply this configuration when onboarding new Azure SSO users or updating role assignments.

Before you begin

- Review the existing Azure users and groups; refer to Azure tenants, on page 201.
- Configure an SSO service provider application for the Firewall Management Center; refer to Configure the Firewall Management Center service provider application for Azure, on page 202.
- Enable and configure single sign-on at the Firewall Management Center; refer to Enable SSO at the Firewall Management Center, on page 173, and Configure the Firewall Management Center for Azure SSO, on page 203.

Follow these steps to configure user role mapping for Azure in the management center:

Procedure

Step 1 Choose **System** (⚙) > **Users** > **Single Sign-On**.

Step 2 Expand **Advanced Configuration (Role Mapping)**.

Step 3 From the **Default User Role** drop-down list, choose a default Firewall Management Center user role to assign users.

Step 4 In the **Group Member Attribute** field, enter an attribute configured in Azure for Firewall Management Center user role mapping for users or groups. Refer to step 1 of Configure User Role Mapping for Individual Users at the Azure IdP, on page 206 or step 1 of Configure User Role Mapping for Groups at the Azure IdP, on page 208.

Step 5 Next to each Firewall Management Center user role you want to assign to SSO users, enter a regular expression. The Firewall Management Center compares these values against the user role mapping attribute the IdP sends to the Firewall Management Center with SSO user information. The Firewall Management Center grants users a union of all the roles for which a match is found.

Step 6 Click **Test Configuration**. If the system displays an error message, review the SSO configuration for the Firewall Management Center as well as the identity service provider application configuration, correct any errors, and try again.

Step 7 When the system reports a successful configuration test, click **Apply**.

After completing this task, Azure SSO users are mapped to the appropriate management center roles based on their Azure attributes, ensuring correct access permissions are applied automatically.

What to do next

Configure user role mapping at the service provider application; refer to Azure IdP user role mappings, on page 206.

Azure IdP user role mappings

An Azure IdP user role mapping is a configuration method that assigns user or group role attribute values from application roles in the Azure AD portal. This mapping also determines user access by comparing attribute values with regular expressions assigned to each Firewall Management Center user role in the SSO configuration.

User role mapping configuration options

You can configure SSO user role mapping at the Azure AD portal based on individual user permissions or group permissions.

Choose one of these mapping methods:

- To map based on individual user permissions, refer to Configure User Role Mapping for Individual Users at the Azure IdP.
- To map based on group permissions, refer to Configure User Role Mapping for Groups at the Azure IdP.

During SSO login, Azure presents a user or group role attribute value from an application role configured at the Azure AD portal to the Firewall Management Center. The Firewall Management Center compares that attribute value with the regular expressions assigned to each Firewall Management Center user role in the SSO configuration and grants the user all the roles for which a match is found. If no match is found, the Firewall Management Center grants the user a configurable default user role. The regular expression assigned to each Firewall Management Center user role must comply with the restricted version of Google's RE2 standard, which is supported by Golang and Perl. The Firewall Management Center treats the attribute value received from Okta as a regular expression using that standard when comparing it against Firewall Management Center user role expressions.



Note A single Firewall Management Center cannot support role mapping for both groups and individual users; you must choose one mapping method for the Firewall Management Center service provider application and use it consistently. The Firewall Management Center can support role mapping using one claim configured in Azure. Generally, group-based role mapping is more efficient for a Firewall Management Center that has many users. Consider user and group definitions established in your Azure tenant when making this decision.

Configure User Role Mapping for Individual Users at the Azure IdP

To establish role mapping for individual users of the Firewall Management Center service application in Azure, use the Azure AD Portal to add a claim to the application, add roles to the application's registration manifest, and assign roles to users.

Before you begin

- Review the Azure tenant; see Azure tenants, on page 201.
- Create and configure a Firewall Management Center service provider application in Azure; see Configure the Firewall Management Center service provider application for Azure, on page 202.
- Configure SSO user role mapping as described in Configure user role mapping for Azure in the Firewall Management Center, on page 205.

Procedure

Step 1 Add a user claim to the SSO configuration for the Firewall Management Center service application with the following characteristics:

- **Name**: Use the same string you entered for the **Group Member Attribute** in the Firewall Management Center SSO configuration. (See Step 5 in Configure user role mapping for Azure in the Firewall Management Center, on page 205.)
- **Name identifier format**: Choose `Persistent`.
- **Source**: Choose `Attribute`.
- **Source attribute**: Choose `user.assignedroles`.

Step 2 Edit the manifest for the Firewall Management Center service application (in JSON format) and add application roles to represent Firewall Management Center user roles you wish to assign to SSO users. The simplest approach is to copy an existing application role definition and change the following properties:

- `displayName`: The name for the role that will appear in the AD Azure Portal.
- `description`: A brief description of the role.
- `Id`: An alphanumeric string that must be unique among ID properties within the manifest.
- `value`: A string to represent one or more Firewall Management Center user roles. (Note: Azure does not permit spaces in this string.)

Step 3 For each user assigned to the Firewall Management Center Service application, assign one of the application roles you have added to the manifest for that application. When a user logs in to the Firewall Management Center using SSO, the application role you assign to that user is the value Azure sends to the Firewall Management Center in the claim for the service application. The Firewall Management Center compares the claim against the expressions you assigned to Firewall Management Center user roles in the SSO configuration (See Step 6 of Configure user role mapping for Azure in the Firewall Management Center, on page 205.), and assigns the user all the Firewall Management Center user roles for which there is a match.

What to do next

- Test your role mapping scheme by logging into the Firewall Management Center using SSO from various accounts and confirming that users are assigned Firewall Management Center user roles as you expect.

Configure User Role Mapping for Groups at the Azure IdP

To establish role mapping for user groups for the Firewall Management Center service application in Azure, use the Azure AD Portal to add a claim to the application, add roles to the application's registration manifest, and assign roles to groups.

Before you begin

- Review the Azure tenant; see Azure tenants, on page 201.
- Create and configure a Firewall Management Center service provider application in Azure; see Configure the Firewall Management Center service provider application for Azure, on page 202.
- Configure SSO user role mapping as described in Configure user role mapping for Azure in the Firewall Management Center, on page 205.

Procedure

Step 1 Add a user claim to the SSO configuration for the Firewall Management Center service application with the following characteristics:

- **Name**: Use the same string you entered for the **Group Member Attribute** in the Firewall Management Center SSO configuration. (See Step 5 in Configure user role mapping for Azure in the Firewall Management Center, on page 205.)
- **Name identifier format**: Choose `Persistent`.
- **Source**: Choose `Attribute`.
- **Source attribute**: Choose `user.assignedroles`.

Step 2 Edit the manifest for the Firewall Management Center service application (in JSON format) and add application roles to represent Firewall Management Center user roles you wish to assign to SSO users. The simplest approach is to copy an existing application role definition and change the following properties:

- `displayName`: The name for the role that will appear in the Ad Azure Portal.
- `description`: A brief description of the role.
- `Id`: An alphanumeric string that must be unique among id properties within the manifest.
- `value`: A string to represent one or more Firewall Management Center user roles. (Azure does not permit spaces in this string.)

Step 3 For each group assigned to the Firewall Management Center Service application, assign one of the application roles you have added to the manifest for that application. When a user logs in to the Firewall Management Center using SSO, the application role you assign to that user's group is the value Azure sends to the Firewall Management Center in the claim for the service application. The Firewall Management Center compares the claim against the expressions you assigned to Firewall Management Center user roles in the SSO configuration (see Step 6 of Configure user role mapping for Azure in the Firewall Management Center, on page 205), and assigns the user all the Firewall Management Center user roles for which there is a match.

What to do next

Test your role mapping scheme by logging into the Firewall Management Center using SSO from various accounts and confirming that users are assigned Firewall Management Center user roles as you expect.

Azure User Role Mapping Examples

As the following examples demonstrate, the SSO configurations at the Firewall Management Center to support user role mapping are the same for both individual users and for groups. The difference lies in the settings at the Firewall Management Center service provider application in Azure.



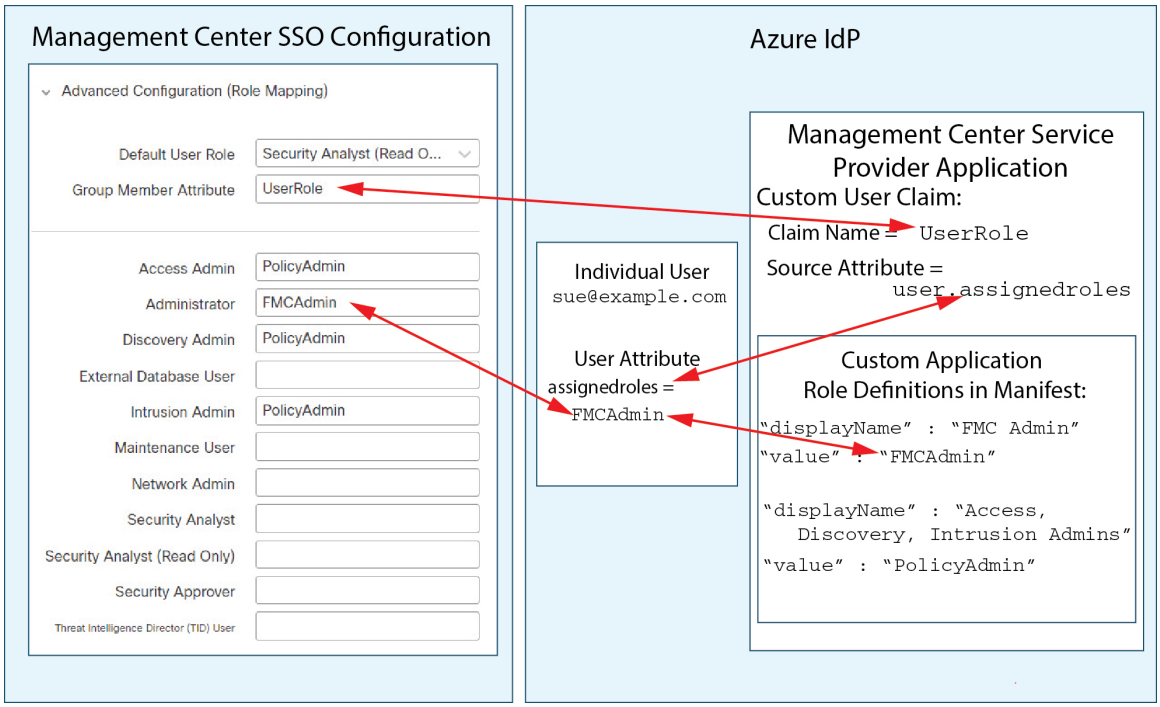
Note You can configure Firewall Management Center roles to be mapped based on individual permissions or based on group permissions, but a single Firewall Management Center application cannot support role mapping for both groups and individual users. The Firewall Management Center can support role mapping using only one claim configured in Azure.

Azure Role Mapping Example for Individual User Accounts

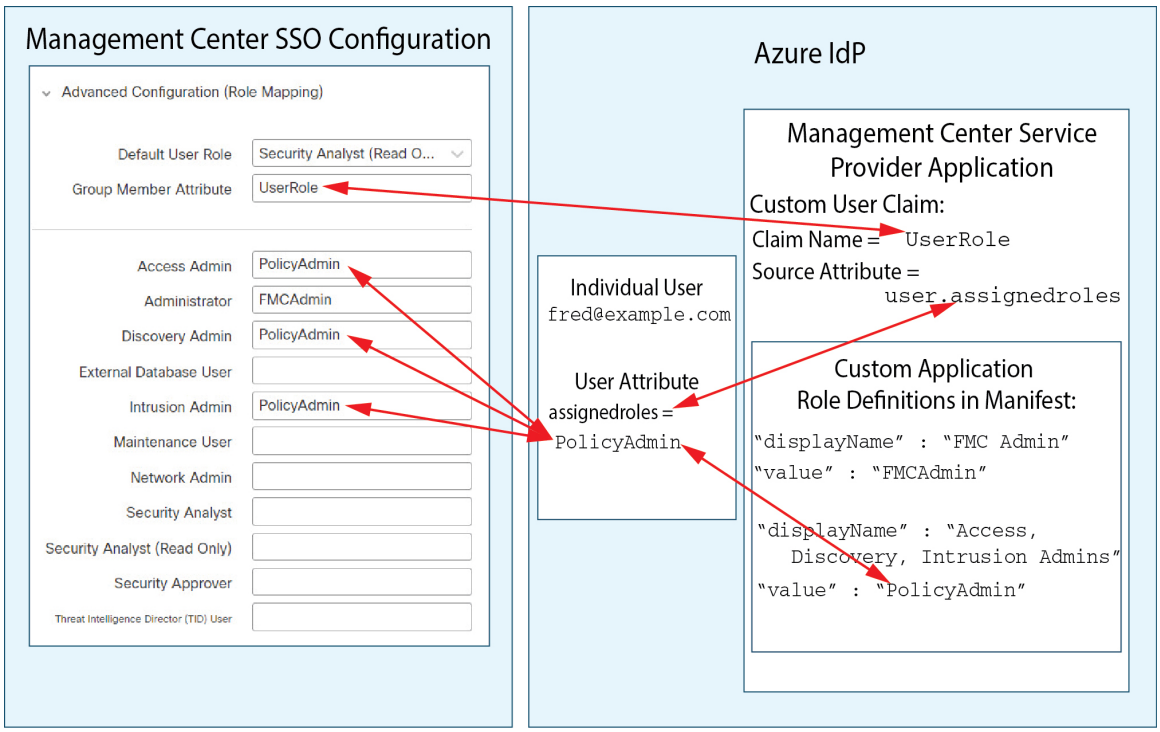
In role mapping for individual users, the Azure Firewall Management Center service application has custom roles defined within its manifest. (In this case, FMCAdmin and PolicyAdmin.) These roles can be assigned to users; Azure stores role assignments for each user in that user's assignedroles attribute. The application also has a custom user claim defined, and this claim is configured to get its value from the assigned user role for a user logging into the Firewall Management Center using SSO. Azure passes the claim value to the Firewall Management Center during the SSO login process, and the Firewall Management Center compares the claim value against strings assigned to each Firewall Management Center user role in the Firewall Management Center SSO configuration.

The following diagrams illustrate how the relevant fields and values in the Firewall Management Center and Azure configurations correspond to each other in user role mapping for individual accounts. Each diagram uses the same SSO configurations at the Firewall Management Center and at the Azure AD portal, but the configuration for each user at the Azure AD portal differs to assign each user different roles at the Firewall Management Center.

- In this diagram sue@ example.com uses the assignedroles attribute value `FMCAdmin`, and the Firewall Management Center assigns her the Firewall Management Center Administrator role.



- In this diagram fred @ example .com uses the assignedroles attribute value `PolicyAdmin`, and the Firewall Management Center assigns him the roles Access Admin, Discovery Admin, and Intrusion Admin.



- Other users assigned to the Azure service application for this Firewall Management Center are assigned the default user role Security Analyst (Read Only) for one of the following reasons:

- They have no value assigned to their assignedroles attribute.
- The value assigned to their assignedroles attribute does not match any expression configured for a user role in the SSO configuration at the Firewall Management Center.

Azure role mapping examples for groups

An Azure role mapping example for groups is a user access configuration that

- assigns custom roles defined in the Azure Firewall Management Center service application to groups
- passes role assignments for each group to group members' assigned roles attribute, and
- uses a custom user claim to map the assigned user role during SSO login, enabling the Firewall Management Center to compare claim values against configured user roles.

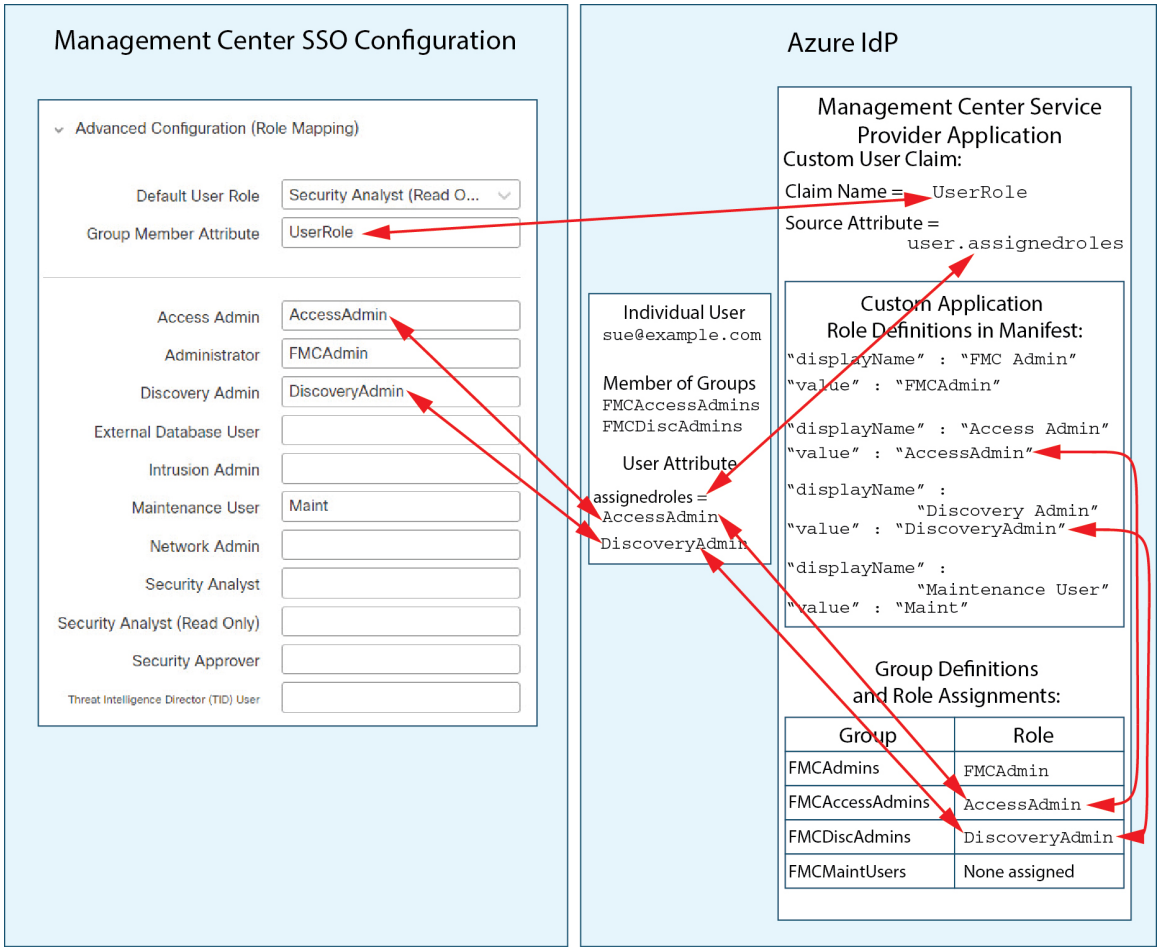
Role mapping process in Azure and Firewall Management Center

Describes how fields and values in the Firewall Management Center and Azure configurations correspond in user role mapping for groups.

Each diagram uses the same SSO configurations at the Firewall Management Center and Azure AD portal. However, user configurations at the Azure AD portal differ, which results in different roles being assigned at the Firewall Management Center.

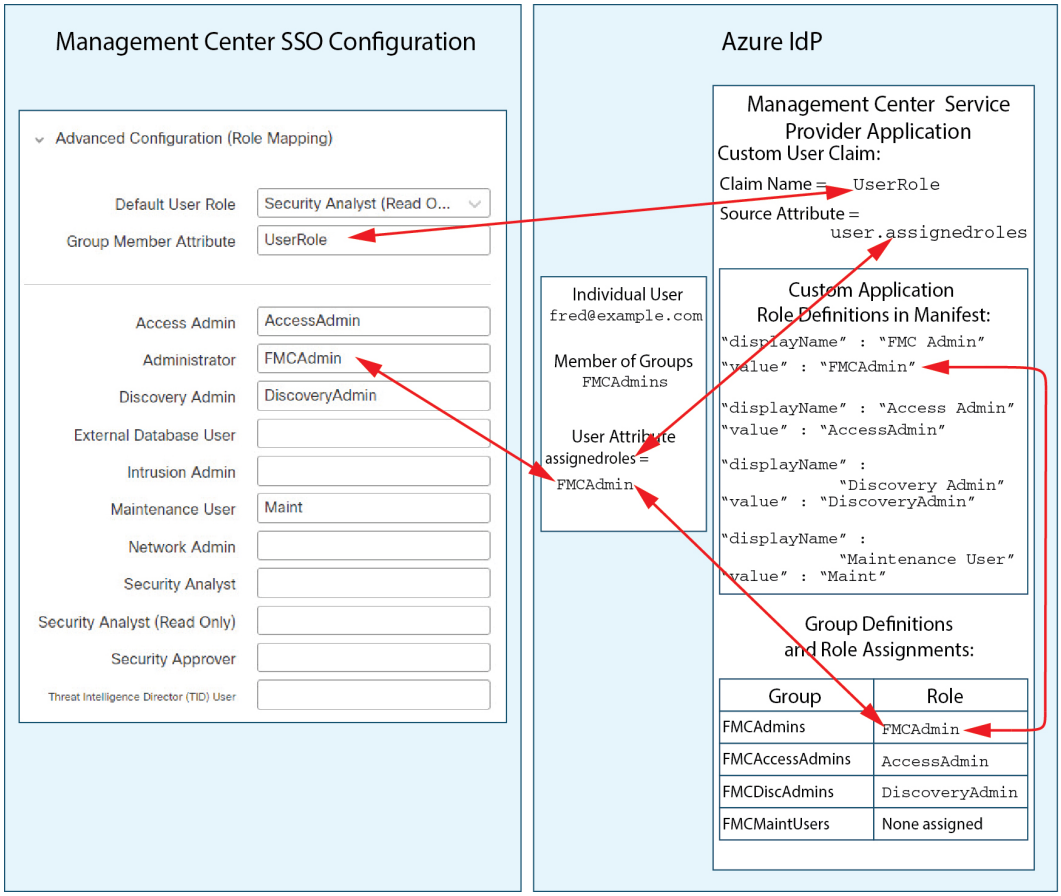
- In this diagram, sue@example.com is a member of the groups `FMCAccessAdmins` and `FMCDiscoveryAdmins`. She inherits the custom roles `AccessAdmin` and `DiscoveryAdmin`. When Sue logs into the Firewall Management Center using SSO, the Firewall Management Center assigns her the roles Access Admin and Discovery Admin.

Figure 23: Role mapping for sue@example.com



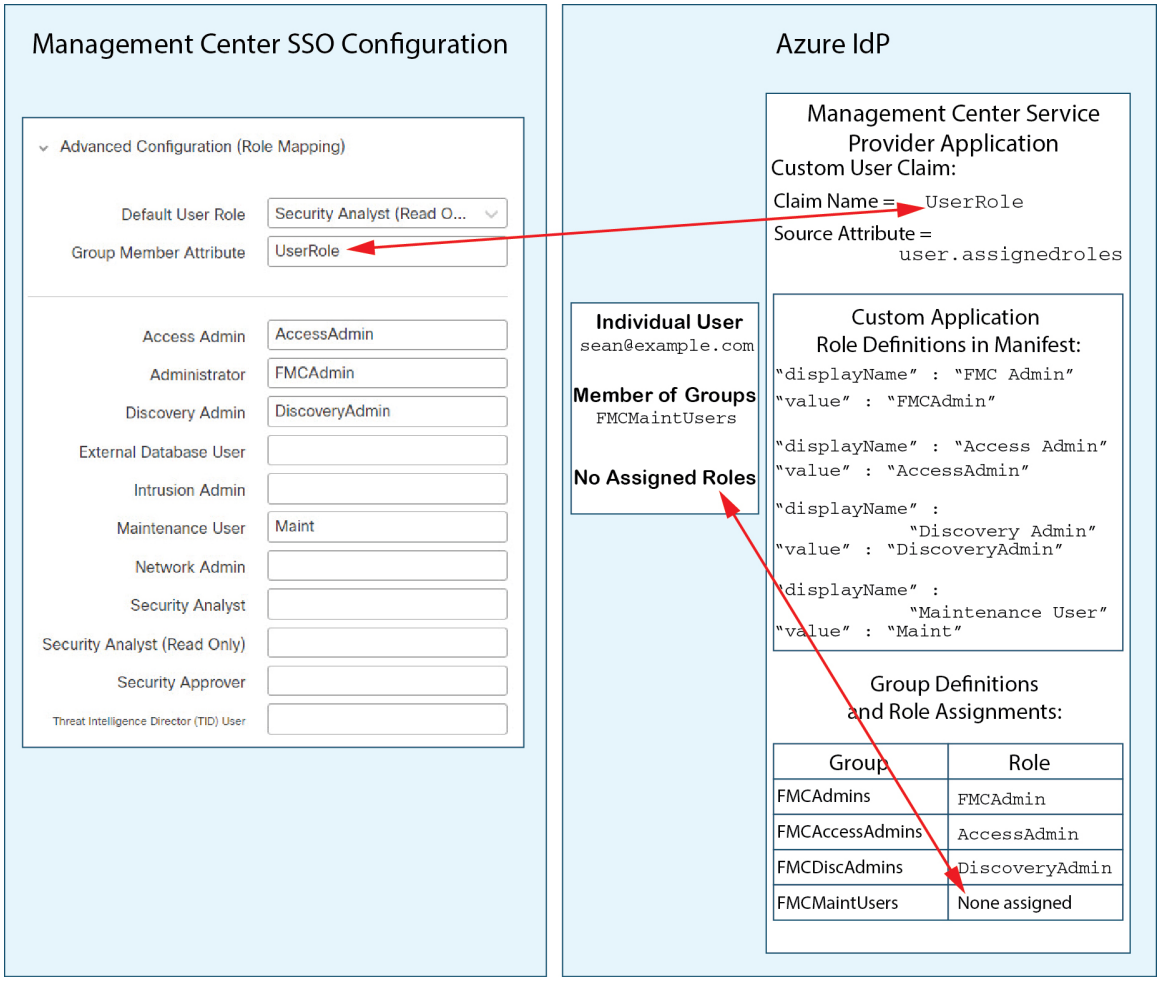
- In this diagram, fred@example.com is a member of the `FMCAdmins` group, from which he inherits the custom role `FMCAdmin`. When Fred logs into the Firewall Management Center using SSO, the Firewall Management Center assigns him the Administrator role.

Figure 24: Role mapping for fred@example.com



- In this diagram, sean@example.com is a member of the `FMCMaintUsers` group, but because no custom role has been assigned to `FMCMaintUsers` within the Azure Firewall Management Center service provider application, Sean has no roles assigned to him. When he logs into the Firewall Management Center using SSO, the Firewall Management Center assigns him the default role Security Analyst (Read Only).

Figure 25: Role mapping for sean@example.com

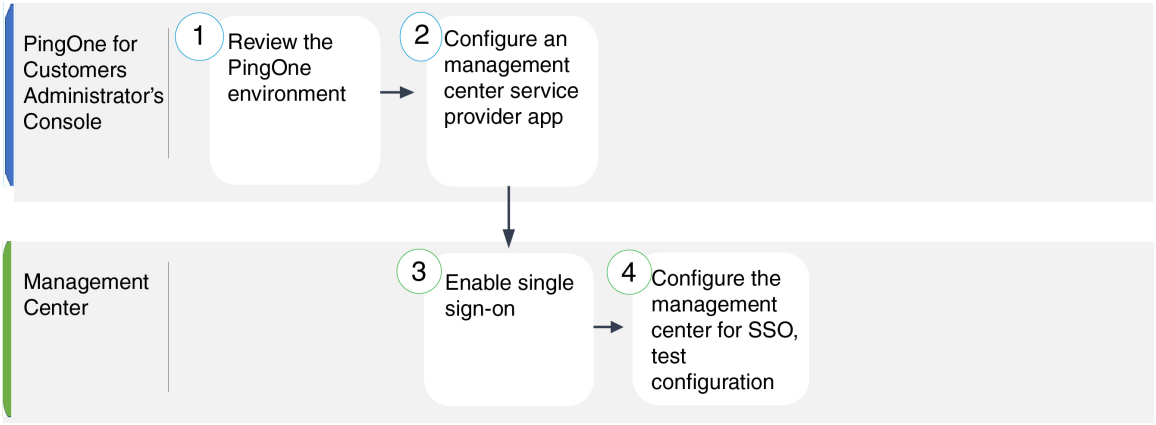


Azure role mapping user scenarios

Examples include users who inherit custom roles from their group memberships and are assigned those roles during SSO login, as illustrated in the diagrams for sue@example.com, fred@example.com, and sean@example.com.

Configure Single Sign-On with PingID

See the following tasks to configure SSO using PingID's PingOne for Customers product:



1	PingOne for Customers Administrator's Console	Review the PingID PingOne for Customers Environment, on page 215.
2	PingOne for Customers Administrator's Console	Configure a service provider application for PingID PingOne for Customers, on page 215.
3	Firewall Management Center	Enable SSO at the Firewall Management Center, on page 173.
4	Firewall Management Center	Configure the Firewall Management Center for SSO with the PingID PingOne for Customers, on page 217.

Review the PingID PingOne for Customers Environment

PingOne for Customers is PingID's cloud-hosted identity-as-a-service (IDaaS) product. In PingOne for Customers, the entity that encompasses all the federated devices that a user can access with the same SSO account is called an environment. Before adding the Firewall Management Center to a PingOne environment, be familiar with its organization; consider the following questions:

- How many users will have access to the Firewall Management Center?
- Do you need to add more users to support SSO access to the Firewall Management Center?

This documentation assumes you are already familiar with the PingOne for Customers Administrator Console and have an account with the Organization Admin role.

Configure a service provider application for PingID PingOne for Customers

Configure a service provider application for PingID PingOne for Customers to establish SAML-based single sign-on (SSO) integration with the Secure Firewall Management Center. This task enables secure authentication and centralized user access management for your environment.

Use the PingOne for Customers Administrator Console to create a Firewall Management Center service provider application within your PingOne for Customers environment and establish basic configuration

settings. This documentation does not describe all functions required to configure an SSO environment. To create users, refer to the PingOne for Customers documentation.

Before you begin

- Familiarize yourself with your PingOne for Customers environment and its users.
- Create additional users if necessary.



Note The system requires that user names for SSO accounts as well as the NameID attribute the IdP sends to the Firewall Management Center during the SAML login process must be both be valid email addresses. Many IdP's automatically use the username of the user trying to logon as the NameID attribute, but you should confirm this is the case for your IdP. Keep this in mind when configuring a service provider application at your IdP and creating IdP user accounts that are to be granted SSO access to the Firewall Management Center.

- Confirm the login URL for the target Firewall Management Center (`https://`*`ipaddress_or_hostname`*)



Note If your Firewall Management Center web interface can be reached with multiple URLs (for instance, a fully-qualified domain name as well as an IP address), SSO users must consistently access the Firewall Management Center using the login URL that you configure in this task.

Follow these steps to configure a service provider application for PingID PingOne for Customers:

Procedure

Step 1 Use the PingOne for Customer Administrator Console to create the application in your environment using these settings:

- Choose the **Web App** application type.
- Choose the **SAML** connection type.

Step 2 Configure the application with these settings for the SAML Connection:

- For the **ACS URL**, append `/sam/ACS` to the Firewall Management Center login URL. For example: `https://ExampleFMC/SAML/ACS`.
- For the **Signing Certificate**, choose Sign Assertion & Response.
- For the **Signing Algorithm**, choose RSA_SHA256.
- For the **Entity ID**, append `/SAML/metadata` to the Firewall Management Center login URL. For example: `https://ExampleFMC/SAML/metadata`.
- For the **SLO Binding**, select HTTP POST.

- For the **Assertion Validity Duration**, enter 300.

Step 3 In the SAMLConnection information for the application, note these values:

- **Single Sign-On Service**
- **Issuer ID**

You will need these values when you configure SSO using PingID's PingOne for Customers product at the Firewall Management Center web interface.

Step 4 For **SAML ATTRIBUTES**, make these selections for a single required attribute:

- **PINGONE USER ATTRIBUTE:** `Email Address`
- **APPLICATION ATTRIBUTE:** `saml_subject`

Step 5 Download the signing certificate in X509 PEM (`.crt`) format and save it to your local computer.

Step 6 (Optional) to make SSO setup at the Firewall Management Center easier, you can download the SAML XML metadata file for the Firewall Management Center service provider application to your local computer.

Step 7 Enable the application.

After you finish this task, your PingOne for Customers environment will have a configured service provider application for the Firewall Management Center, enabling SAML-based SSO integration.

What to do next

Enable single sign-on. Refer to Enable SSO at the Firewall Management Center, on page 173.

Configure the Firewall Management Center for SSO with the PingID PingOne for Customers

Configure the Firewall Management Center to integrate with PingID PingOne for Customers for single sign-on (SSO) authentication. This task establishes SSO connectivity and user role mapping between your system and PingID PingOne for Customers.

Use this task when you need to enable SSO for your Firewall Management Center using PingID PingOne for Customers as the identity provider.

This configuration is relevant for organizations seeking centralized authentication and streamlined user access management. Ensure you have administrative access to both the Firewall Management Center

and the PingOne for Customers Administrator Console.

Before you begin

- Create a Firewall Management Center service provider application at the PingOne for Customers Administrator Console; refer to Configure a service provider application for PingID PingOne for Customers, on page 215.
- Enable single sign-on; refer to Enable SSO at the Firewall Management Center, on page 173.

Follow these steps to configure SSO integration between the Firewall Management Center and PingID PingOne for Customers:

Procedure

Step 1 (This step continues directly from Enable SSO at the Firewall Management Center, on page 173.) At the **Configure PingID Metadata** dialog, you have two choices:

- To enter the SSO configuration information manually:
  a. Click the **Manual Configuration** radio button.
  b. Enter the values you retrieved from the PingOne for Customers Administrator Console:
     - For **Identity Provider Single Sign-On URL** enter the **Single Signon Service** you noted in Step 3 of Configure a service provider application for PingID PingOne for Customers, on page 215.
     - For **Identity Provider Issuer** enter the **Issuer ID** you noted in Step 3 of Configure a service provider application for PingID PingOne for Customers, on page 215.
     - For the **X.509 Certificate**, use the certificate you downloaded from PingOne for Customers in Step 5 of Configure a service provider application for PingID PingOne for Customers, on page 215. (Use a text editor to open the certificate file, copy the contents, and paste it into the **X.509 Certificate** field.)
- If you saved the XML metadata file generated by PingOne for Customers to your local computer (Step 6 of Configure a service provider application for PingID PingOne for Customers, on page 215), click the **Upload XML File** radio button to upload the file to the Firewall Management Center.

Click **Next**.

Step 2 At the **Verify Metadata** dialog, review the configuration parameters and click **Save**.

Step 3 Expand **Advanced Configuration (Role Mapping)**.

Step 4 From the **Default User Role** drop-down list, choose a default Firewall Management Center user role to assign users.

Step 5 In the **Group Member Attribute** field, enter an attribute configured in PingID PingOne for Firewall Management Center user role mapping for users or groups.

Step 6 Next to each Firewall Management Center user role you wish to assign to SSO users, enter a regular expression. The Firewall Management Center compares these values against the user role mapping attribute the IdP sends to the Firewall Management Center with SSO user information. The Firewall Management Center grants users a union of all the roles for which a match is found.

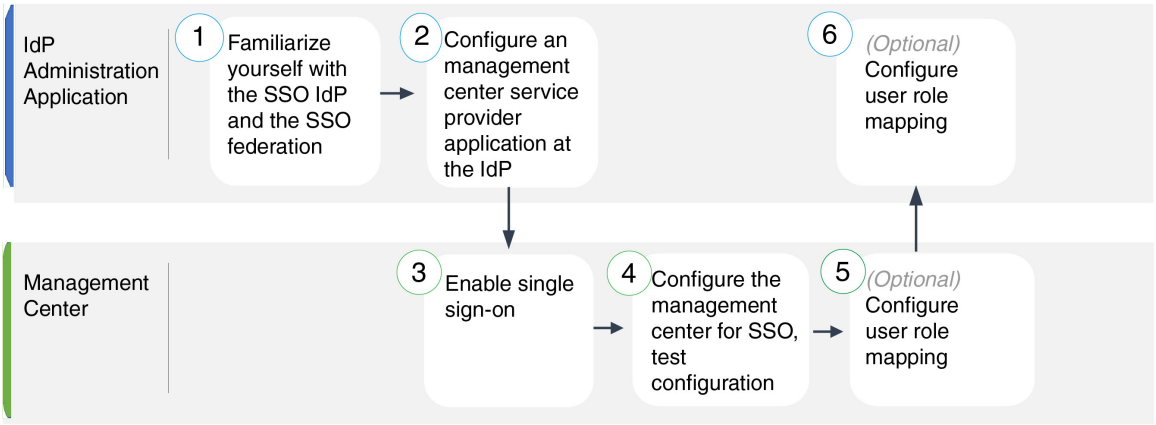
Step 7 Click **Test Configuration**. If the System displays an error message, review the SSO configuration for the Firewall Management Center as well as the PingOne for Customers service provider application, correct any errors, and try again.

Step 8 When the system reports a successful configuration test, click **Apply**.

After completing this task, the Firewall Management Center will be integrated with PingID PingOne for Customers for SSO. Users will be able to authenticate using their PingID credentials and be assigned roles based on your configuration.

Configure Single Sign-On with Any SAML 2.0-Compliant SSO Provider

The Firewall Management Center supports single sign-on with any SSO identity provider (IdP) compliant with the SAML 2.0 SSO protocol. Generic instructions to use a wide range of SSO providers must address the tasks to be performed at a high level; establishing SSO using a provider not specifically addressed in this documentation requires that you be proficient with the IdP of your choice. These tasks help you determine the steps to configure the Firewall Management Center for single sign-on using any SAML 2.0-compliant SSO provider:



1	IdP Administration Application	Familiarize Yourself with the SSO Identity Provider and the SSO Federation, on page 220.
2	IdP Administration Application	Configure Firewall Management Center service provider application for any SAML 2.0-compliant SSO provider, on page 220.
3	Firewall Management Center	Enable SSO at the Firewall Management Center, on page 173.
4	Firewall Management Center	Configure the Firewall Management Center for SSO using any SAML 2.0-compliant SSO provider, on page 222.
5	Firewall Management Center	Configure User Role Mapping in the Firewall Management Center for SAML 2.0-Compliant SSO Providers, on page 224.
6	IdP Administration Application	Configure Firewall Management Center User Role Mapping at the IdP for SAML 2.0-Compliant SSO Providers, on page 225.

Familiarize Yourself with the SSO Identity Provider and the SSO Federation

Read the IdP vendor documentation with the following considerations in mind:

- Does the SSO provider require that users subscribe to or register with any services before using the IdP?
- What terminology does the SSO provider use for common SSO concepts? For instance, to refer to a group of federated service provider applications, Okta uses "org" where Azure uses "tenant."
- Does the SSO provider support SSO exclusively, or a suite of functions—for instance, multifactor authentication or domain management? (This can affect configuration of some elements shared between features—especially users and groups.)
- What permissions does an IdP user account need to configure SSO?
- What configurations does the SSO provider require you to establish for a service provider application? For instance, Okta automatically generates an X509 Certificate to secure its communications with the Firewall Management Center, while Azure requires that you generate that certificate using the Azure portal interface.
- How are users and groups created and configured? How are users assigned to groups? How are users and groups granted access to service provider applications?
- Does the SSO provider require that at least one user be assigned to a service provider application before the SSO connection can be tested?
- Does the SSO provider support user groups? How are user and group attributes configured? How can you map attributes to Firewall Management Center user roles in the SSO configuration?
- Do you need to add more users or groups to the federation to support SSO on the Firewall Management Center?
- Are users within the federation members of groups?
- Are user and group definition native to the IdP or imported from a user management application such as Active Directory, RADIUS, or LDAP?
- What kind of user role assignments do you want to make? (If you choose not to assign user roles, the Firewall Management Center automatically assigns the user a configurable default user role to all SSO users.)
- How must users and groups within the federation be organized to support your plan for user role mapping?

Configure Firewall Management Center service provider application for any SAML 2.0-compliant SSO provider

This task enables you to configure a service provider application at your identity provider (IDP) for SAML 2.0 SSO integration with the Firewall Management Center.

- Establish secure SSO authentication for users and groups.
- Ensure the IDP and service provider exchange the required SAML attributes and certificates.

Generally, SSO providers require you to configure a service provider application at the IDP for each federated application. All IdPs that support SAML 2.0 SSO need the same configuration information for service provider applications. Some IdPs automatically generate some configuration settings for you, while others require you to configure all settings yourself.

- If you plan to assign user groups to the Firewall Management Center application, do not also assign users within those groups as individuals.
- The Firewall Management Center cannot support role mapping using multiple SSO attributes. You must select either user role mapping or group role mapping, and configure a single attribute to convey user role information from the IDP to the Firewall Management Center.

Before you begin

- Familiarize yourself with the SSO federation, its users, and its groups. Refer to Familiarize Yourself with the SSO Identity Provider and the SSO Federation, on page 220.
- Confirm your IDP account has the necessary permissions to perform this task.
- Create user accounts or groups, or both, in your SSO federation if necessary.

Note The system requires that user names for SSO accounts as well as the NameID attribute the IdP sends to the Firewall Management Center during the SAML login process must be both be valid email addresses. Many IdP's automatically use the username of the user trying to logon as the NameID attribute, but you should confirm this is the case for your IdP. Keep this in mind when configuring a service provider application at your IdP and creating IdP user accounts that are to be granted SSO access to the Firewall Management Center.

- Confirm the login URL for the target Firewall Management Center (for example, `https://ipaddress_or_hostname`)

Note If your Firewall Management Center web interface can be reached with multiple URLs (for example, a fully qualified domain name as well as an IP address), SSO users must consistently access the Firewall Management Center using the login URL that you configure in this task.

Follow these steps to configure a service provider application for SAML 2.0 SSO integration:

Procedure

Step 1 Create a new service provider application at the IDP.

Step 2 Configure values required by the IDP. Be sure to include the fields listed below, required to support SAML 2.0 SSO functionality with the Firewall Management Center. (Because different SSO service providers use different terminology for SAML concepts, this list provides alternate names for these fields to help you find the right settings in the IDP application.):

- Service Provider Entity ID, Service Provider Identifier, Audience URI: A globally unique name for the service provider (the Firewall Management Center), formatted as a URL. To create this, append the string `/SAML/metadata` to the Firewall Management Center login URL, such as `https://ExampleFMC/SAML/metadata`.

- Single Sign on URL, Recipient URL, Assertion Consumer Service URL: The service provider (Firewall Management Center) address to which the browser sends information on behalf of the IDP. To create this, append the string `SAML/acs` to the Firewall Management Center login URL, such as `https://ExampleFMC/SAML/acs`.
- X.509 Certificate: Certificate to secure communications between the Firewall Management Center and the IDP. Some IdPs may automatically generate the certificate, and some may require that you explicitly generate it using the IDP interface.

Step 3 (Optional if you are assigning groups to the application) Assign individual users to the Firewall Management Center application. (If you plan to assign groups to the Firewall Management Center application, do not assign members of those groups as individuals.)

Step 4 (Optional if you are assigning individual users to the application.) Assign user groups to the Firewall Management Center application.

Step 5 (Optional) Some IdPs provide the ability to generate a SAML XML metadata file containing the information you have configured in this task formatted to comply with SAML 2.0 standards. If your IDP provides this ability, you can download the file to your local computer to ease the SSO configuration process at the Firewall Management Center.

After completing this task, your SAML 2.0-compliant SSO provider will be able to authenticate users or groups for access to the Firewall Management Center using single sign-on. The IDP and service provider will exchange the required SAML attributes and certificates for secure authentication.

What to do next

Enable single sign-on. Refer to Enable SSO at the Firewall Management Center, on page 173.

Configure the Firewall Management Center for SSO using any SAML 2.0-compliant SSO provider

Configure the Firewall Management Center to use single sign-on (SSO) with any SAML 2.0-compliant SSO provider. This task allows you to centralize authentication and streamline user access. This configuration helps integrate your security management with your organization's identity provider (IdP).

Use these instructions at the Firewall Management Center web interface. To configure the Firewall Management Center for SSO using any SAML 2.0-compliant SSO provider, you need information from the IdP.

This task is relevant when you want to enable SSO for centralized user authentication and access control. Ensure you have access to the IdP and the necessary SSO configuration details before starting.

Before you begin

- Review the organization of your SSO federation, and its users and groups.
- Configure a Firewall Management Center service provider application at the IdP; see Configure the Firewall Management Center for SSO using any SAML 2.0-compliant SSO provider, on page 222.
- Gather the following SSO configuration information for the service provider application from the IdP. Because different SSO service providers use different terminology for SAML concepts, this list provides alternate names for these fields to help you find the right values in the IdP application:
  - Identity Provider Single Sign-On URL, Login URL: The IdP URL where the browser sends information on behalf of the Firewall Management Center.

- Identity Provider Issuer, Identity Provider Issuer URL, Issuer URL: A globally unique name for the IdP, often formatted as a URL.
- An X.509 digital certificate to secure communications between the Firewall Management Center and the IdP.

- Enable single sign-on; see Enable SSO at the Firewall Management Center, on page 173.

Follow these steps to configure the Firewall Management Center for SSO using any SAML 2.0-compliant SSO provider:

Procedure

Step 1 (This step continues directly from Enable SSO at the Firewall Management Center, on page 173.) At the **Configure SAML Metadata** dialog, you have two choices:

- To enter the SSO configuration information manually:
  a. Click the **Manual Configuration** radio button.
  b. Enter the following values previously obtained from the SSO Service Provider application:
     - **Identity Provider Single Sign-On URL**
     - **Identity Provider Issuer**
     - **X.509 Certificate**

- If you saved the XML metadata file generated at the IdP (Step 5 in Configure Firewall Management Center service provider application for any SAML 2.0-compliant SSO provider, on page 220), you can upload the file to the Firewall Management Center:
  a. Click the **Upload XML File** radio button.
  b. Follow the on-screen instructions to navigate to and choose the XML metadata file on your local computer.

Step 2 Click **Next**.

Step 3 At the **Verify Metadata** dialog, review the configuration parameters and click **Save**.

Step 4 Click **Test Configuration**. If the system displays an error message, review the SSO configuration for the Firewall Management Center as well as the service provider application configuration at the IdP, correct any errors, and try again.

Step 5 When the system reports a successful configuration test, click **Apply**.

After completing this task, the Firewall Management Center is configured for SSO using your SAML 2.0-compliant SSO provider. Users can now authenticate using their IdP credentials.

What to do next

You may optionally configure user role mapping for SSO users; see Configure User Role Mapping in the Firewall Management Center for SAML 2.0-Compliant SSO Providers, on page 224. If you choose not to configure role mapping, by default all SSO users that log into the Firewall Management Center are assigned the default user role you configure in Step 4 of Configure User Role Mapping in the Firewall Management Center for SAML 2.0-Compliant SSO Providers, on page 224.

Configure User Role Mapping in the Firewall Management Center for SAML 2.0-Compliant SSO Providers

To implement SAML SSO user role mapping you must establish coordinating configurations at the IdP and at the Firewall Management Center.

- At the IdP, establish user or group attributes to convey user role information and assign values to them; the IdP sends these to the Firewall Management Center once it has authenticated and authorized an SSO user.
- At the Firewall Management Center, associate values with each of the Firewall Management Center user roles you want to assign to users.

When the IdP sends the user or group attribute associated with an authorized user, the Firewall Management Center, the Firewall Management Center compares the attribute value with the values associated with each Firewall Management Center user role, assigns the user all the roles that match, . The Firewall Management Center performs this comparison by treating both values as regular expressions that comply with the restricted version of Google's RE2 regular expression standard supported by Golang and Perl.

The fields to configure for user role mapping at the Firewall Management Center web interface are the same regardless of your choice of SSO provider. But the values you configure must take into account how the SAML SSO provider you use implements user role mapping. Your IdP may enforce syntactical limitations on user or group attributes; if so, you must devise a user role mapping scheme using role names and regular expressions compatible with those requirements.

Before you begin

- Configure an SSO service provider application for the Firewall Management Center; see Configure Firewall Management Center service provider application for any SAML 2.0-compliant SSO provider, on page 220.
- Enable and configure single sign-on at the Firewall Management Center, see Enable SSO at the Firewall Management Center, on page 173, and Configure the Firewall Management Center for SSO using any SAML 2.0-compliant SSO provider, on page 222.

Procedure

Step 1 Choose **System** (⚙) > **Users** > **Single Sign-On**.

Step 2 Expand **Advanced Configuration (Role Mapping)**.

Step 3 Select a Firewall Management Center user role to assign users as a default value from the **Default User Role** drop-down.

Step 4 Enter a **Group Member Attribute**. This string must match an attribute name configured at the IdP Firewall Management Center service provider application for user role mapping using either users or groups. (See Step 1 of Configure Firewall Management Center User Role Mapping at the IdP for SAML 2.0-Compliant SSO Providers, on page 225.)

Step 5 Next to each Firewall Management Center user roll you wish to assign to SSO users, enter a regular expression. The Firewall Management Center compares these values against the user role mapping attribute the IdP sends to the Firewall Management Center with SSO user information. The Firewall Management Center grants users a union of all the roles for which a match is found.

Step 6 Click **Test Configuration**. If the system displays an error message, review the SSO configuration for the Firewall Management Center as well as the identity service provider application configuration, correct any errors, and try again.

Step 7 When the system reports a successful configuration test, click **Apply**.

What to do next

Configure user role mapping at the service provider application; see Configure Firewall Management Center User Role Mapping at the IdP for SAML 2.0-Compliant SSO Providers, on page 225.

Configure Firewall Management Center User Role Mapping at the IdP for SAML 2.0-Compliant SSO Providers

The detailed steps for configuring user role mapping are different for each IdP. You must determine how to create a custom user or group attribute for the service provider application, and assign values to the attribute for each user or group at the IdP to convey user or group privileges to the Firewall Management Center. Keep in mind the following:

- If your IdP imports user or group profiles from a third-party user management application (such as Active directory, LDAP, or Radius), this may affect how you can use attributes for role mapping.
- Take into account user and group role definitions throughout your SSO federation.
- The Firewall Management Center cannot support role mapping using multiple SSO attributes; you must select either user role mapping or group role mapping and configure a single attribute to convey user role information from the IdP to the Firewall Management Center.
- Group role mapping is generally more efficient for a Firewall Management Center with many users.
- If you assign user groups to Firewall Management Center applications, do not also assign users within those groups as individuals.
- For the purpose of determining a match with Firewall Management Center user roles, the Firewall Management Center treats user and group role attribute values received from the IdP as regular expressions complying with the restricted version of Google's RE2 regular expression standard supported by Golang and Perl. Your IdP may enforce certain syntactical limitations on user or group attributes. if so, you must devise a user role mapping scheme using role names and regular expressions compatible with those requirements.

Before you begin

- Confirm your IdP account has the necessary permissions to perform this task.

- Configure a Firewall Management Center service provider application at the IdP (see Configure Firewall Management Center service provider application for any SAML 2.0-compliant SSO provider, on page 220).

Procedure

Step 1 At the IdP, create or designate an attribute to be sent to the Firewall Management Center to contain role mapping information for each user sign-in. This may be a user attribute, a group attribute, or a different attribute that obtains its value from a source such as user or group definitions maintained by the IdP or a third party user management application.

Step 2 Configure how the attribute gets its value. Coordinate the possible values with the values associated with the user roles in the Firewall Management Center SSO configuration.

Custom user roles for the web interface

A custom user role is a user access configuration that

- assigns specific permissions to each user account
- enables management of user roles for web interface access, and
- allows configuration of custom roles beyond default user roles.

Web interface user role reference information

Each user account must be defined with a user role. For default user roles, refer to User roles, on page 140.

Create Custom User Roles

Custom user roles can have any set of menu-based and system permissions, and may be completely original, copied from a predefined or another custom user role, or imported from another Firewall Management Center.

Note (Requires Version 7.2.6+) Although you can enable access to content updates without product upgrades, we recommend against the reverse: product without content. That is, if you enable **Product Upgrades** in a custom user role, also enable **Content Updates**. Otherwise, you could have trouble manually uploading upgrade packages and upgrading older ASA FirePOWER and NGIPSv devices.

Procedure

Step 1 Choose **System** (⚙) > **Users**.

Step 2 Click **User Roles**.

Step 3 Add a new user role with one of the following methods:

- Click **Create User Role**.
- Click the **Copy** ( ) next to the user role you want to copy.
- Import a custom user role from another Firewall Management Center:
  - **a.** On the other Firewall Management Center, click the **Export**( ) to save the role to your computer.
  - **b.** On the new Firewall Management Center, choose **System** ( ) > **Tools** > **Import/Export**.
  - **c.** Click **Upload Package**, then follow the instructions to import the saved user role to the new Firewall Management Center.

Step 4 Enter a **Name** for the new user role. User role names are case sensitive.

Step 5 (Optional) Add a **Description**. Limit the description length to 128 characters.

Step 6 Choose **Menu-Based Permissions** for the new role.

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

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

You can apply restrictive searches to a custom user role. These searches constrain the data a user can see in the tables on the pages available under the Analysis menu. You can configure a restrictive search by first creating a private saved search and selecting it from the **Restrictive Search** drop-down menu under the appropriate menu-based permission.

Step 7 (Optional) Check the **External Database Access (Read Only)** check box to set database access permissions for the new role.

This option provides read-only access to the database using an application that supports JDBC SSL connections. For the third-party application to authenticate to the Firewall Management Center, you must enable database access in the system settings.

Step 8 (Optional) To set escalation permissions for the new user role, see Enable user role escalation, on page 229.

Step 9 Click **Save**.
The custom role is saved. If the system determines it is a read-only role, it labels the role with '(Read Only)'. This is relevant when configuring the number of concurrent sessions for read-only vs read-write users. You cannot make a role read-only by adding '(Read Only)' to the role name. For more information on concurrent session limits, see Configure user settings, on page 128.

Example

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

The following table lists custom roles that you could create and user permissions granted for each example. The table lists the privileges required for each custom role. In this example, Policy Approvers

can view (but not modify) access control and intrusion policies. They can also deploy configuration changes to devices.

Table 9: Sample Access Control Custom Roles

Menu-Based Permission	Example Roles		
	Access Control Editor	**Intrusion & Network Analysis Editor**	**Policy Approver**
Access Control	yes	no	yes
Access Control Policy	yes	no	yes
Modify Access Control Policy	no	no	no
Intrusion Policy	no	yes	yes
Modify Intrusion Policy	no	yes	no
Deploy Configuration to Devices	no	no	yes

Deactivate User Roles

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

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

Procedure

Step 1 Choose **System** (⚙) > **Users**.

Step 2 Click **User Roles**.

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

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

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

Enable user role escalation

Enable user role escalation to allow custom user roles to temporarily gain the privileges of another targeted user role, in addition to their base role. This feature helps substitute one user for another during an absence or track advanced privilege usage.

- This feature is not supported for default user roles.

You can configure user role escalation so that users can use their own passwords or the password of another user that you specify. Managing a single escalation password for all applicable users is also possible.

- For example, a user with limited privileges can escalate to the Administrator role to perform administrative actions.

Before you begin

No explicit prerequisites are required for this task.

Follow these steps to enable user role escalation:

Procedure

Step 1 Set the escalation target role, on page 229. Only one user role at a time can be the escalation target role.

Step 2 Configure a custom user role for escalation, on page 230.

Step 3 (For the logged in user) Escalate your user role, on page 231.

After completing these steps, custom user roles can temporarily gain the privileges of a targeted user role, allowing for flexible privilege management and tracking of advanced user privileges.

Set the escalation target role

Set a user role as the escalation target role to allow custom roles to escalate permissions during login sessions.

- Configure a predefined or custom user role to act as the escalation target.
- Enable system-wide escalation for user roles as needed.

You can assign any user role, predefined or custom, to act as the system-wide escalation target role.

This role is used when a custom role escalates, if permitted. Only one user role at a time can be the escalation target role. Each escalation lasts for the duration of a login session and is recorded in the audit log.

- Escalation is effective for the current login session and tracked in the audit log.

Before you begin

No explicit prerequisites are required for this task.

- Ensure you have access to the Secure Firewall Management Center and appropriate permissions to configure user roles.

Follow these steps to set the escalation target role:

Procedure

Step 1 Choose **System** (⚙) > **Users**.

Step 2 Click **User Roles**.

Step 3 Click **Configure Permission Escalation**.

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

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

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

After completing this task, the selected user role becomes the escalation target. Users in escalated sessions will have the permissions of the new escalation target role, and all escalations are logged in the audit log.

What to do next

No post-requisites are required for this task.

- Review the audit log to verify escalation events as needed.

Configure a custom user role for escalation

Configure a custom user role to enable escalation, so users can perform privileged actions when required.

- Allow users to escalate privileges using a designated password.
- Support efficient management of escalating users and passwords.

Users must belong to a custom user role with escalation enabled to perform privileged actions. This procedure describes how to enable escalation for a custom user role.

Consider your organization's needs when configuring the escalation password. You may choose another user whose password serves as the escalation password. Changing or deactivating that user affects all escalating users who require that password. Using an externally-authenticated user allows centralized management.

- Deactivating the user whose password is used for escalation makes escalation impossible for users with the role that requires it.
- This feature can be used to quickly remove escalation powers if necessary.

Before you begin

Set a target user role according to Set the escalation target role, on page 229.

- Ensure you have access to create or modify custom user roles.

Follow these steps to configure a custom user role for escalation:

Procedure

Step 1 Begin configuring your custom user role as described in Create Custom User Roles, on page 226.

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

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

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

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

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

Step 4 Click **Save**.

After completing this task, users in the custom role will be able to escalate privileges as configured. Escalation is managed according to the password option you selected.

What to do next

Review user access and escalation settings to ensure compliance with organizational policies.

- Monitor changes to escalation passwords and user roles regularly.

Escalate your user role

Escalate your user role to gain the permissions of a target role in addition to your current role. This is useful when you need temporary access to additional features or administrative capabilities.

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

Before you begin

Your user role must have escalation permissions enabled by your administrator.

Follow these steps to escalate your user role:

Procedure

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

If you do not see this option, your administrator did not enable escalation for your user role.

Step 2 Enter the authentication password.

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

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

After completing these steps, you will have the permissions of the escalation target role for the duration of your login session. Logging out will return you to your base role privileges.

Troubleshoot LDAP authentication connections

This reference explains how to troubleshoot LDAP authentication connection issues, identify common failure causes, and apply solutions to restore communication between Cisco devices and LDAP servers.

Configure User Preferences

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

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

Change your password

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

When password strength checking is enabled, passwords must comply with the strong password requirements described in Guidelines and limitations for user accounts for Firewall Management Center, on page 146.

If you are an LDAP or a RADIUS user, you cannot change your password through the web interface. Ensure you are not an LDAP or RADIUS user if you intend to change your password through the web interface.

Before you begin

Follow these steps to change your password:

Procedure

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

Step 2 Click **Change Password**.

Step 3 Optionally, check the **Show password** check box. This action displays the password while you use the dialog.

Step 4 Enter your **Current Password**.

Step 5 You have two options:

- Enter your new password for **New Password** and **Confirm Password**.

- Click **Generate Password** to let the system create a compliant password for you. Generated passwords are non-mnemonic. Record your password carefully if you select this option.

Step 6 Click **Apply**.

Your password is updated, and you can use the new password to access your account. If password requirements are not met, the system prompts you to enter a compliant password.

What to do next

Log in with your new password to verify the change was successful.

Change an expired password

Change your password when it has expired to regain access to your account and maintain security compliance.

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

Before you begin

Follow these steps to change an expired password:

Procedure

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

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

  Tip
  When password strength checking is enabled, passwords must comply with the strong password requirements described in Guidelines and limitations for user accounts for Firewall Management Center, on page 146.

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

After completing this task, your password is updated and you regain access to your account. If you choose to skip, you may be prompted again until you change your password.

Change the web interface appearance

Change the web interface appearance by choosing from available themes or enable new design options for the light theme variant.

You can change the way the web interface appears. Theme options and design toggles may vary based on your system version and available features.

Before you begin

Follow these steps to change the web interface appearance:

Procedure

From the drop-down list under your user name, choose a theme:

- **Light**
- **Dusk**
- **Classic** (the look and feel before Version 6.6)

The web interface appearance changes according to your selected theme or design preference, providing a personalized and comfortable user experience.

Specify your home page

You can set a custom home page within the web interface for your appliance. This helps personalize your experience and improve workflow efficiency..

You can specify the page within the web interface to use as your home page for the appliance. The default home page is the default dashboard (**Overview** > **Dashboards heading** > **Dashboard**), except for user accounts with no dashboard access, such as External Database users. (Refer to Specify your default dashboard, on page 240 to set the default dashboard.)

In a multidomain deployment, the home page selection applies to all domains accessible to your user account. Certain pages are only available in the Global domain.

Procedure

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

Step 2 Click **Home Page**.

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

The options in the drop-down list are based on the access privileges for your user account. For more information, refer to User roles, on page 140.

Step 4 Click **Save**.

Your selected page becomes the new home page and will be displayed when you log in to the appliance web interface.

Configure event view settings

Configure event view settings to customize how events are displayed and managed on the Firewall Management Center.

Use the Event View Settings page to configure characteristics of event views on the Firewall Management Center. Some configurations are available only for specific user roles. If you have the External Database User role, you can view event view settings, but changing those settings does not affect system behavior.

Procedure

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

Step 2 Click **Event View Settings**.

Step 3 In the **Event Preferences** section, configure the basic characteristics of event views ; refer to Configure event view preferences, on page 235.

Step 4 In the **File Preferences** section, configure file download preferences ; refer to Configure file download preferences, on page 236.

Step 5 In the **Default Time Windows** section, configure the default time window or windows ; refer to Default time windows, on page 237.

Step 6 In the **Default Workflow** sections, configure default workflows ; refer to Default workflows, on page 239.

Step 7 Click **Save**.

The system saves your event view settings and uses them to display and manage events according to your preferences.

Configure event view preferences

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

The Event Preferences section contains these fields.

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

  If this setting is enabled and you click **Delete All** in an event view, you must confirm deletion of all the events that meet the current constraints. This confirmation includes events not shown on the current page. The appliance then deletes the events from the database.

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

  An event view may display slowly if it contains many IP addresses and this option is enabled. To enable this setting, use the management interfaces configuration to establish a DNS server in the system settings.

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

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

Configure file download preferences

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

If your appliance does not support downloading captured files, these options are disabled.

These fields appear in the File Preferences section:

- The **Confirm 'Download File' Actions** check box determines whether a File Download pop-up window appears when you download a file. The pop-up displays a warning and prompts you to continue or cancel.

Caution Cisco recommends that you avoid downloading malware because it can cause harmful effects. Always use caution when downloading files. Files may contain malware, so secure your download destination by taking precautions before downloading files.

You can disable this option whenever you download a file.

- When you download a captured file, the system creates a password-protected .zip archive containing the file. The **Zip File Password** field defines the password you want to use to restrict access to the .zip file.

  If you leave this field blank, the system creates archive files without passwords.
- The **Show Zip File Password** check box toggles displaying plain text or obfuscated characters in the **Zip File Password** field. When this field is cleared, the **Zip File Password** displays obfuscated characters.

Default time windows

A default time window is a time constraint setting that

- controls the default behavior of event views by limiting events to a specified time range
- allows users to manually change the time window for individual event views during event analysis, and
- resets to the defaults configured on the Event View Settings page at the start of each session.

User access and event types for default time windows

User-role access to the Default Time Windows section varies:

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

Time window settings are valid only for the current session. When you log out and log back in, time windows reset to the defaults configured on the Event View Settings page.

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

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

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

Time window settings do not affect event views that display hosts, host attributes, applications, clients, vulnerabilities, user identity, or compliance allow list violations.

Configuration options for time windows:

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

There are three types of time window:

- *static*, which displays all the events generated from a specific start time to a specific end time.
- *expanding*, which displays all the events generated from a specific start time to the present; as time moves forward, the time window expands and new events are added to the event view.
- *sliding*, which displays all the events generated from a specific start time (for example, one day ago) to the present; as time moves forward, the time window “slides” so that you see only the events for the range you configured (in this example, for the last day).

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

Time window settings options in the **Time Window Settings** drop-down list:

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

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

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

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

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

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

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

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

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

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

  For **expanding** time windows, disable the **Use End Time** check box. You see all events generated from midnight Sunday to the present. When you change event views, the time window expands to the current time. If you analyze events for more than one week before logging out, this time window can extend beyond one week.

Default time window configuration example

For example, an administrator can configure a sliding time window to display events from the last hour, or set a static time window for the current day to review events that occurred during a specific period.

Non-constrained event views

Event views that display hosts, host attributes, applications, clients, vulnerabilities, user identity, or compliance allow list violations are not affected by time window settings.

Time window analogy

A default time window is like a filter that limits the events you see to a specific period, similar to viewing only the most recent messages in an email inbox.

Default workflows

A default workflow is a predefined workflow that

- is configured by the appliance for each event type
- determines which workflow is displayed when viewing events of a specific type, and
- can be changed based on user role permissions.

Default workflow configuration and behavior

The appliance provides a default workflow for each event type. This workflow defines how event data is presented to you.

- Each event type has at least one predefined workflow.
- For intrusion events, you see the Events by Priority and Classification workflow by default.
- When you view intrusion events, the appliance automatically displays the default workflow.

You can change the default workflow for each event type. However, your user role determines whether you can configure these defaults. If you are an intrusion event analyst, you cannot set default discovery event workflows.

Default workflow example

For example, as a Security Analyst, you can choose among ten different intrusion event workflows, each presenting intrusion event data differently. The appliance displays the Events by Priority and Classification workflow by default when you view intrusion events.

Set your default time zone

Set your default time zone so that all times shown in the web interface for your user account reflect your local time preferences. This helps you view scheduled tasks and dashboards in your preferred time zone.

This setting determines the times displayed in the web interface for your user account, such as for task scheduling and dashboard viewing. It does not change the system time, affect other users, or modify data stored in the system, which generally uses UTC.

- The Time Zone function in User Preferences assumes that the system clock is set to UTC. Do not change the system time. Changing the system time from UTC is not supported and will require reimaging the device to recover from an unsupported state.

- This feature does not affect the time zone used for time-based policy application. Set the time zone for a device in **Devices** > **Platform Settings**.

Before you begin

Follow these steps to set your default time zone:

Procedure

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

Step 2 Click the **TIME Zone** drop-down list.

Step 3 Choose the continent or country and the state name that corresponds with the time zone you want to use.

After completing these steps, your user account will display times in the web interface according to your selected time zone. System time and other users are not affected.

Specify your default dashboard

Configure your default dashboard to personalize your workspace and enable quicker access to the most relevant widgets and data for your role.

The default dashboard appears when you choose **Overview** > **Dashboards heading** > **Dashboard**. Unless changed, the default dashboard for all users is the Summary dashboard.

In a multidomain deployment, the default dashboard you select applies to every domain your account can access. If your account frequently accesses multiple domains, note that some dashboard widgets are only available in specific domains.

Before you begin

Your user role must be Administrator, Maintenance, or Security Analyst to change the default dashboard. Ensure you have access to the User Preferences menu.

Follow these steps to specify your default dashboard:

Procedure

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

Step 2 Click **Dashboard Settings**.

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

Step 4 Click **Save**.

Your selected dashboard appears by default when you access the system and displays your preferred widgets and information.

Configure How-To settings

Configure the How To widget to enable guided walkthroughs for tasks in Secure Firewall Management Center. The widget helps users complete tasks efficiently by providing step-by-step instructions on various UI screens.

The How To widget provides walkthroughs to help you navigate and complete tasks on the Firewall Management Center. Walkthroughs guide you through each step, regardless of the UI screens involved. The **How To** widget is enabled by default.

For a list of feature walkthroughs supported in the Firewall Management Center, refer to Feature Walkthroughs Supported in Secure Firewall Management Center.

- Walkthroughs are generally available for all UI pages and are not specific to user roles. However, some menu items may not be visible if user privileges restrict access. Walkthroughs will not execute on those pages.
- This feature is not available in the Classic theme.

Before you begin

Follow these steps to configure How-To settings:

Procedure

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

Step 2 Click the **How-To Settings** tab.

Step 3 Check the **Enable How-Tos** check box to enable How-Tos.

Step 4 Click **Save**.

After completing these steps, the How To widget will be configured and enabled, allowing users to access guided walkthroughs for supported tasks in the Secure Firewall Management Center.

What to do next

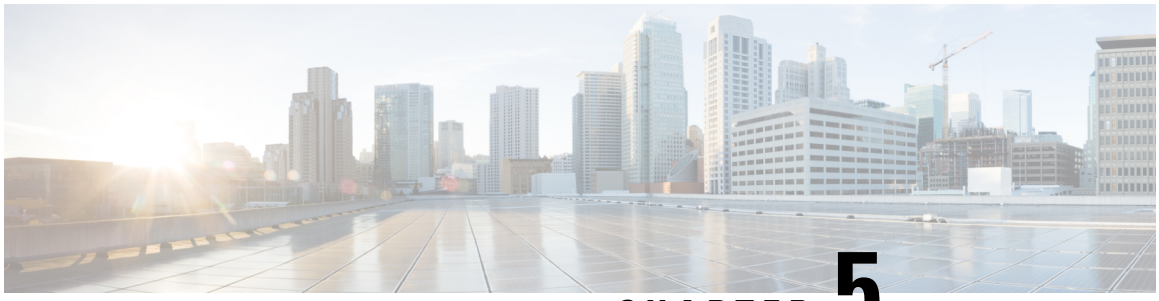
To open the How To widget, choose **Help** (?) > **On-screen Assistance** > **How-Tos**. You can search for How To walkthroughs to find tasks relevant to your needs. For more information, refer to Search for how to walkthroughs, on page 21.

History for Firewall Management Center user accounts

This reference lists the history of features, enhancements, and changes for Firewall Management Center user accounts, including version numbers and details for each update.

This table lists features, enhancements, and changes for user accounts.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Added new field for assigning Shell user name template	7.0.0	Any	The **Shell User Name Template** field provides a template for CLI access attributes for LDAP external authentication. This field uniquely identifies LDAP CLI users. New or modified screens: **System** > **Users** > **External Authentication**.
Added support for Single Sign-On with any SAML 2.0–compliant SSO provider	6.7.0	Any	Added the ability to support Single Sign-On for external users configured at any third-party SAML 2.0-compliant identity provider (IdP). This allows mapping user or group roles from the IdP to Firewall Management Center user roles. Only users with the Admin role authenticated internally or by LDAP or RADIUS can configure SSO. New or modified screens: **System** > **Users** > **Single Sign-On**.
Themes for the web interface	6.6.0	Any	You can choose the look and feel of the web interface. You can choose the Light, Dusk, or Classic theme for the web interface. New or modified screens: **User name** > **User Preference** > **General** > **UI Theme.**
Added a new field for name in user accounts	6.6.0	Any	You can use the new field to identify the user or department responsible for an internal user account. New or modified screens: **System** > **Users** > **Users > Real Name** field.
Cisco Security Manager Single Sign-on no longer supported	6.5.0	Any	Single Sign-on between the Firewall Management Center and Cisco Security Manager is no longer supported as of Firepower 6.5. New or modified screens: **System** > **Users > CSM Single Sign-on.**
Enhanced password security	6.5.0	Any	New requirements for strong passwords now appear in a single place in this chapter and are cross-referenced from other chapters. The change password interface includes new fields: **Show Password** and **Generate Password**. New or modified screens: **User name** > **User Preferences** > **General** > **Change password.**



CHAPTER 5

Domains

These topics describe how to manage multitenancy using domains:

- Multitenancy using domains, on page 243
- Requirements and prerequisites for domains, on page 246
- Manage domains, on page 247
- Create new domains, on page 247
- Move data between domains, on page 248
- Move devices between domains, on page 249
- History for domain management, on page 253

Multitenancy using domains

Multitenancy using domains is a segmentation approach that

- segments user access to managed devices, configurations, and events,
- allows creation of up to 100 subdomains under a top-level Global domain in two or three levels, and
- requires users to log into a single domain called the current domain.

Domain hierarchy and access control

The Firewall Management Center allows you to implement multitenancy using domains. When you log into the Firewall Management Center, you log into a single domain, called the *current domain*. Depending on your user account, you may be able to switch to other domains. You can create up to 100 subdomains under a top-level Global domain, in two or three levels.

When you log into the Firewall Management Center, you log into a single domain, referred to as the *current domain*. Depending on your user account, you may be able to switch to other domains.

In addition to any restrictions imposed by your user role, your current domain level can also limit your ability to modify various configurations. The Firewall Management Center limits most management tasks, like system software updates, to the Global domain.

The Firewall Management Center limits other tasks to *leaf domains*, which are domains with no subdomains. For example, you must associate each managed device with a leaf domain, and perform device management tasks from the context of that leaf domain. Note that each device can only belong to a single domain.

Each leaf domain builds its own network map, based on the discovery data collected by that leaf domain's devices. Events reported by a managed device (connection, intrusion, malware, and so on) are also associated with the device's leaf domain.

Domain deployment levels include:

- **One domain level: Global** - If you do not configure multitenancy, all devices, configurations, and events belong to the Global domain, which in this scenario is also a leaf domain. Except for domain management, the system hides domain-specific configurations and analysis options until you add subdomains.
- **Two domain levels: Global and second-level** - In a two-level multidomain deployment, the Global domain has direct descendant domains only. For example, a managed security service provider (MSSP) can use a single Firewall Management Center to manage network security for multiple customers.
- **Three domain levels: Global, second-level, and third-level** - In a three-level multidomain deployment, the Global domain has subdomains, at least one of which has its own subdomain.

Two-level domain deployment scenario

In a two-level multidomain deployment for an MSSP managing multiple customers:

- Administrators at the MSSP logging into the Global domain cannot view or edit customers' deployments. They must log into the respective second-level named subdomains to manage the customers' deployment.
- Administrators for each customer can log into second-level named subdomains to manage only the devices, configurations, and events applicable to their organizations. These local administrators cannot view or affect the deployments of other customers of the MSSP.

Three-level domain deployment scenario

In a three-level multidomain deployment where an MSSP customer wants to separately manage edge and internal network devices:

- Administrators for the customer logging into the second-level subdomain cannot view or edit the customer's edge network deployments. They must log into the respective leaf domain to manage the devices deployed on the network edge.
- Administrators for the customer's edge network can log into a third-level (leaf) domain to manage only the devices, configurations, and events applicable to devices deployed on the network edge. Similarly, administrators for the customer's internal network can log into a different third-level domain to manage internal devices, configurations, and events. Edge and internal administrators cannot view each other's deployment.



Note

- In the Firewall Management Center that uses multi-tenancy, SSO configuration can be applied only at the global domain level, and applies to the global domain and all subdomains.
- In deployments that use multitenancy with Firewall Management Center domains, passive identity sessions apply to all managed Firewall Threat Defense devices, regardless of domain.

Related Topics

SAML Single Sign-On, on page 169

Domains terminology

This documentation uses these terms when describing domains and multidomain deployments.

Domain terminology definitions:

- Global Domain: In a multidomain deployment, the top-level domain. If you do not configure multitenancy, all devices, configurations, and events belong to the Global domain. Administrators in the Global domain can manage the entire Secure Firewall System deployment.
- Subdomain: A second or third-level domain.
- Second-level domain: A child of the Global domain. Second-level domains can be leaf domains, or they can have subdomains.
- Third-level domain: A child of a second-level domain. Third-level domains are always leaf domains.
- Leaf domain: A domain with no subdomains. Each device must belong to a leaf domain.
- Descendant domain: A domain descending from the current domain in the hierarchy.
- Child domain: A domain's direct descendant.
- Ancestor domain: A domain from which the current domain descends.
- Parent domain: A domain's direct ancestor.
- Sibling domain: A domain with the same parent.
- Current domain: The domain you are logged into now. The system displays the name of the current domain before your user name at the top right of the web interface. Unless your user role is restricted, you can edit configurations in the current domain.

Domain properties

To modify a domain's properties, you must have Administrator access in that domain's parent domain.

Each domain has configurable properties that control its behavior and resource allocation:

- Name and Description: Each domain must have a unique name within its hierarchy. A description is optional.
- Parent Domain: Second- and third-level domains have a parent domain. You cannot change a domain's parent after you create the domain.
- Devices: Only leaf domains may contain devices. In other words, a domain may contain subdomains or devices, but not both. You cannot save a deployment where a non-leaf domain directly controls a device. In the domain editor, the web interface displays available and selected devices according to their current place in your domain hierarchy.
- Host Limit: The number of hosts the Firewall Management Center can monitor, and therefore store in network maps, depends on its model. In a multidomain deployment, leaf domains share the available pool of monitored hosts, but have separate network maps. To ensure that each leaf domain can populate

its network map, you can set host limits at each subdomain level. If you set a domain's host limit to `0`, the domain shares in the general pool.

Setting the host limit has a different effect at each domain level:

- Leaf — For a leaf domain, a host limit is a simple limit on the number of hosts the leaf domain can monitor.
- Second Level — For a second-level domain that manages third-level leaf domains, a host limit represents the total number of hosts that the leaf domains can monitor. The leaf domains share the pool of available hosts.
- Global — For the Global domain, the host limit is equal to the total number of hosts the Firewall Management Center can monitor. You cannot change it

The sum of subdomains' host limits can add up to more than their parent domain's host limit. For example, if the Global domain host limit is 150,000, you can configure multiple subdomains each with a host limit of 100,000. Any of those domains, but not all, can monitor 100,000 hosts.

The network discovery policy controls what happens when you detect a new host after you reach the host limit; you can drop the new host, or replace the host that has been inactive for the longest time. Because each leaf domain has its own network discovery policy, each leaf domain governs its own behavior when the system discovers a new host.

If you reduce the host limit for a domain and its network map contains more hosts than the new limit, the system deletes the hosts that have been inactive the longest.

Related Topics

Host limits
Network discovery data storage settings

Requirements and prerequisites for domains

This reference provides the requirements and prerequisites necessary for working with domains, including model support, supported domains, and required user roles.

Model support

Any.

Supported domains

Any

User roles

- Admin

Manage domains

This task allows you to modify domain properties and structure to organize your network management hierarchy effectively.

To modify a domain's properties, you must have Administrator access in that domain's parent domain.

Before you begin

Follow these steps to manage domains:

Procedure

Step 1 Choose **System** (⚙) > **Domains**.

Step 2 Manage your domains:

- Add — Click **Add Domain**, or click **Add Subdomain** next to the parent domain; see Create new domains, on page 247.
- Edit — Click **Edit** (✎) next to the domain you want to modify; see Domain properties, on page 245.
- Delete — Click **Delete** (🗑) next to the empty domain you want to delete, then confirm your choice. Move devices from domains you want to delete by editing their destination domain.

Step 3 When you are done making changes to the domain structure and all devices are associated with leaf domains, click **Save** to implement your changes.

Step 4 If prompted, make additional changes:

- If you changed a leaf domain to a parent domain, move or delete the old network map; see Move data between domains, on page 248.
- If you moved devices between domains and must assign new policies and security zones or interface groups, see Move devices between domains, on page 249.

Your domain structure is updated according to your modifications, with devices properly organized within the domain hierarchy.

What to do next

- Configure user roles and policies (access control, network discovery, and so on) for any new domains. Update device properties as needed.
- Deploy configuration changes; see the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Create new domains

Create new domains to organize devices into hierarchical domain structures that support policy management and device configuration.

You can create up to 100 subdomains under a top-level Global domain, in two or three levels.

You must assign all devices to a leaf domain before you can implement the domain configuration. When you add a subdomain to a leaf domain, the domain stops being a leaf domain and you must reassign its devices.

Procedure

Step 1 In a Global or a second-level domain, choose **System** (⚙) > **Domains**.

Step 2 Click **Add Domain**, or click **Add Subdomain** next to the parent domain.

Step 3 Enter a **Name** and **Description**.

Step 4 Choose a **Parent Domain**.

Step 5 On **Devices**, choose the **Available Devices** to add to the domain, then click **Add to Domain** or drag and drop into the list of **Selected Devices**.

Step 6 Optionally, click **Advanced** to limit the number of hosts the new domain may monitor; see Domain properties, on page 245.

Step 7 Click **Save** to return to the domain management page.

The system warns you if any devices are assigned to non-leaf domains. Click **Create New Domain** to create a new domain for those devices. Click **Keep Unassigned** if you plan to move the devices to existing domains.

Step 8 When you are done making changes to the domain structure and all devices are associated with leaf domains, click **Save** to implement your changes.

Step 9 If prompted, make additional changes:

- If you changed a leaf domain to a parent domain, move or delete the old network map; see Move data between domains, on page 248.
- If you moved devices between domains and must assign new policies and security zones or interface groups, see Move devices between domains, on page 249.

What to do next

- Configure user roles and policies (access control, network discovery, and so on) for any new domains. Update device properties as needed.
- Deploy configuration changes; see the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Move data between domains

Move network maps and associated events to appropriate domains when restructuring your domain hierarchy.

Because events and network maps are associated with leaf domains, when you change a leaf domain to a parent domain, you have two choices:

- Move the network map and associated events to a new leaf domain.

- Delete the network map but retain the events. In this case, the events remain associated with the parent domain until the system prunes events as needed or as configured. Alternatively, you can delete old events manually.

Before you begin

Implement a domain configuration where a former leaf domain is now a parent domain; see Manage domains, on page 247.

Follow these steps to move data between domains:

Procedure

Step 1 For each former leaf domain that is now a parent domain:

- Choose a new **Leaf Domain** to inherit the **Parent Domain**'s events and network map.
- Choose **None** to delete the parent domain's network map, but retain old events.

Step 2 Click **Save**.

The data is moved or retained according to your selection, and the domain restructuring is complete.

What to do next

Deploy configuration changes; see the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Move devices between domains

Move devices between domains to reorganize your network structure while preserving device configurations and ensuring proper policy assignments in the target domain.

You can move devices between domains as long as the source and the target domains are visible from the domain where you are moving the devices. Moving a device between domains can affect the configurations and policies applied to the device. The system retains the following device configurations while moving devices between domains.

- Interfaces
- Inline sets
- Routing
- DHCP
- Associated objects
- SNMP (if available)

The following changes can occur to the configuration of a device when it is moved between domains:

- If you want the system to retain the device configurations after the devices are moved to the target domain, ensure that:
  - The shared access control policies are in the Global domain. We also recommend that the other shared policies are in the Global domain.
- For VPN configurations,
  - The remote access VPN configurations and device certificates are in the global or target domain.
  - When you assign a remote access VPN policy to a device, you can move the device from one domain to another, only if the target domain is a descendant of the domain in which remote access VPN is configured.
- The network objects for SNMP are in the global domain.
- You can move the device into any child domain without deleting the enrolled certificate on the device. Specifically:
  - If the health policy applied to a moved device is inaccessible in the new domain, you can choose a new health policy.
  - If the access control policy assigned to a moved device is not valid or accessible in the new domain, choose a new policy. Every device must have an assigned access control policy.
  - If the interfaces on the moved device belong to a security zone that is inaccessible in the new domain, you can choose a new zone.
  - Interfaces are removed from:
    - Security zones that are inaccessible in the new domain and not used in an access control policy.
    - All interface groups.

If devices require a policy update but you do not need to move interfaces between zones, the system displays a message stating that zone configurations are up to date. For example, if a device's interfaces belong to a security zone configured in a common ancestor domain, you do not need to update zone configurations when you move devices from subdomain to subdomain.

Before you begin

- Create new domains. For more information, refer to Create new domains, on page 247.
- Implement a domain configuration where you moved a device from domain to domain and now must assign new policies and security zones; refer to Manage domains, on page 247.

Follow these steps to move devices between domains:

Procedure

Step 1 In the Global domain, choose **System** (⚙) > **Domains**.

Step 2 Edit the target domain to which you plan to move the devices.

Step 3 In the **Edit Domain** dialog box, do these actions:

a. Select the devices that you want to move and click **Add to Domain**.

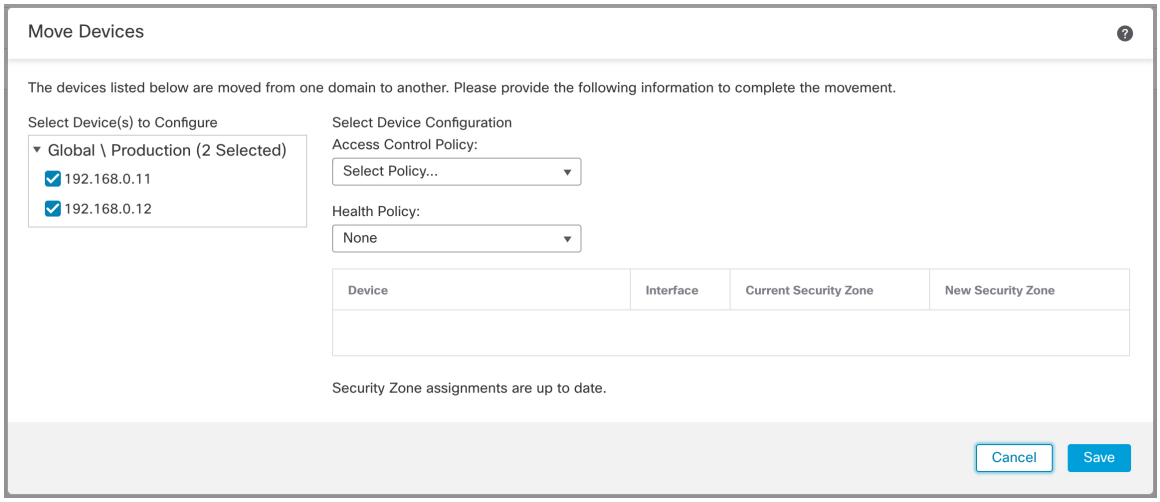
b. Click **Save**.

Step 4 On the Domains page, click **Save**.

Step 5 **(If your access control policies are not in the Global domain)** In the **Move Devices** dialog box, do these actions:

a. Under **Select Device(s) to Configure** check the device that you want to configure.

Check multiple devices to assign the same health and access control policies.



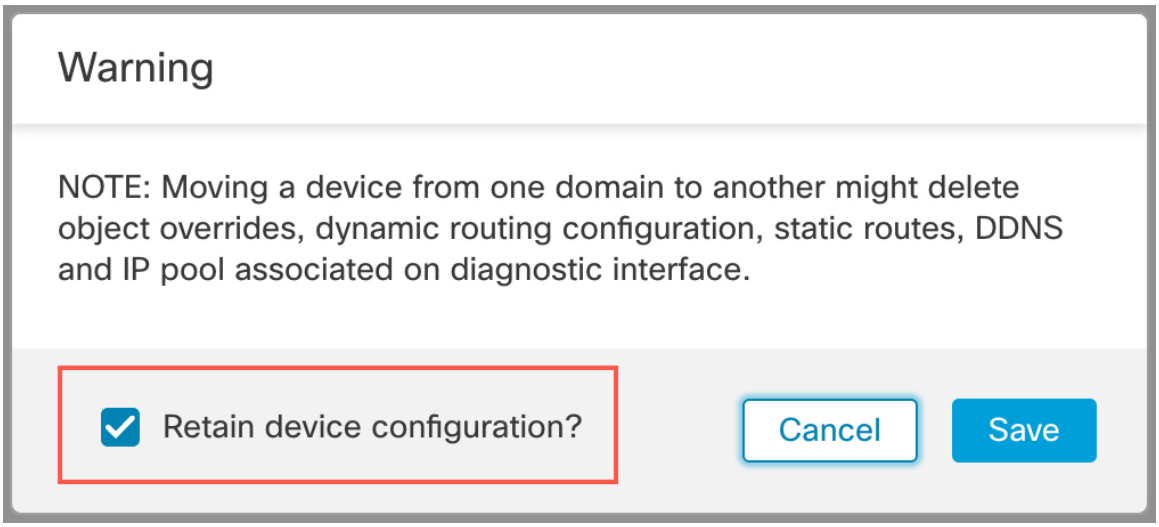
b. Choose an **Access Control Policy** to apply to the device, or choose **New Policy** to create a new policy.

c. Choose a **Health Policy** to apply to the device, or choose **None** to leave the device without a health policy.

d. If prompted to assign interfaces to new zones, choose a **New Security Zone** for each listed interface, or choose **None** to assign it later.

e. After you configure all affected devices, click **Save** to save policy and zone assignments.

Step 6 If you want to retain the device configuration after the move, check the **Retain device configuration?** check box.



If you select this option, the system retains the device configurations after the devices are moved to the target domain. If you do not select this option, you must manually update the device configurations on the moved device that were affected by the move.

This table shows how objects are handled in various scenarios.

Scenario	System Action
Objects exist in the target domain.	Reuse the objects.
Objects with the same name and value exist in the target domain.	Reuse the objects.
Objects with the same name but different values exist in the target domain.	• Network and Port—Create object overrides. • Interface Objects—Create new objects if the type is different. • Reuse all other object types depending on the name match.
Objects do not exist in the target domain.	Create new objects.

Step 7 Click **Save** to implement the domain configuration.

Step 8 After the domain configuration is complete, click **OK**.

The devices are moved to the target domain with the selected policy assignments and configuration settings.

What to do next

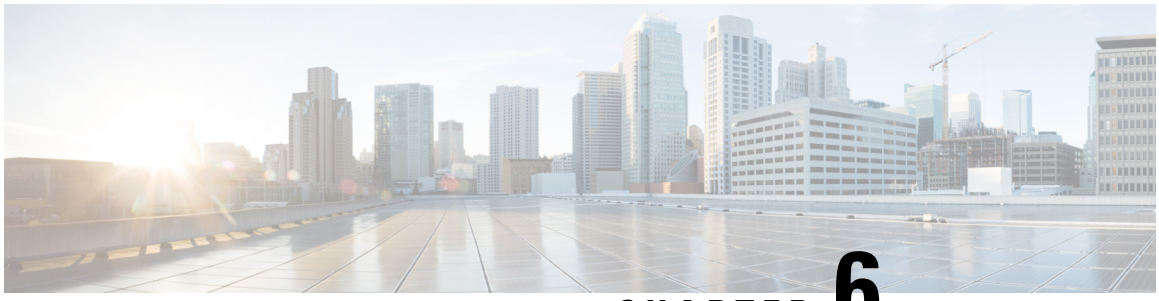
- Update other configurations on the moved device that were affected by the move.
- Deploy configuration changes; see the *Cisco Secure Firewall Management Center Device Configuration Guide*.

- You can manually restore the device configuration if the system fails to retain it after moving a device between domains. For more information, refer to *Export and Import the Device Configuration* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

History for domain management

This reference provides a chronological history of domain management features, including version information, platform compatibility, and detailed descriptions of feature enhancements.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Retain device configurations associated with site-to-site VPN	7.3	Any	While moving devices from one domain to another, you can now retain the device configurations associated with site-to-site VPN only if the site-to-site VPN is configured at the target domain.
Retain the device configuration	7.2	Any	You can now retain the device configuration while moving the device from one domain to another.
Increased maximum number of supported domains	6.5	Any	You can now add up to 100 domains. Previously, the maximum was 50 domains. Supported platforms: Secure Firewall Management Center



CHAPTER 6

Updates

- Content updates, on page 255
- Product upgrades, on page 256
- Guidelines for content updates, on page 256
- Vulnerability database (VDB) updates, on page 256
- Geolocation database (GeoDB) updates, on page 258
- Intrusion rule updates, on page 260
- Maintain your air-gapped deployment, on page 267
- History for content updates, on page 267

Content updates

Content updates keep your deployment current with changing security and contextual information. The system uses this information to evaluate network activity and respond to changes in the threat landscape. Automatic content updates are often enabled by initial configuration or when you enable the related feature.

Types of content updates

Firewall Management Center uses these types of content updates:

- Vulnerability database (VDB): A database of known vulnerabilities to which hosts may be susceptible, as well as fingerprints for operating systems, clients, and applications. The system uses the VDB to help determine whether a particular host increases your risk of compromise.

  Refer to Vulnerability database (VDB) updates, on page 256.

- Geolocation database (GeoDB): Maps IP addresses to countries and continents.

  Refer to Geolocation database (GeoDB) updates, on page 258.

- Intrusion rules (LSP/SRU): Intrusion rule updates provide new and updated intrusion and preprocessor rules and can modify rule states, categories, variables, and default policy settings.

  Refer to Intrusion rule updates, on page 260.

- Security Intelligence feeds: Collections of IP addresses, domain names, and URLs used to filter matching traffic.

  Refer to *Cisco Secure Firewall Management Center Device Configuration Guide*.

- URL categories and reputations: Classification and risk information used to control access to websites.

Refer to *Cisco Secure Firewall Management Center Device Configuration Guide*.

Product upgrades

This guide does not include procedures for upgrading the system software or firewall chassis. Instead, refer to the companion upgrade guide: https://cisco.com/go/ftd-fmc-upgrade-726.

Guidelines for content updates

These guidelines apply to all content updates. Refer to the topics for each update type for update-specific guidance.

User roles

Admin

Domains

Global, except for custom Snort 2 rule import

Internet access

Firewall Management Center must be able to access the internet to download content updates. Refer to Communication Ports and Internet Access, on page 1063.

Release information

Read any release notes or advisory text that accompanies a content update. This information describes compatibility, prerequisites, new capabilities, behavior changes, and warnings.

When to perform content updates

Follow these guidelines when performing and scheduling content updates:

- Scheduled updates: Review scheduled updates to make sure they run at the right time for your environment.
- Bandwidth: Perform updates that require large amounts of bandwidth during periods of low network use.
- Traffic inspection and flow: Perform updates that might interrupt traffic during maintenance windows.

Content updates are scheduled in Coordinated Universal Time (UTC), which remains constant year-round. Scheduled updates do not adjust automatically for local time changes, such as daylight saving time or summer time. For example, an update scheduled for 2:00 a.m. during standard time runs at 3:00 a.m. after the local clock moves forward.

Vulnerability database (VDB) updates

The VDB is a database of known vulnerabilities to which hosts may be susceptible, as well as fingerprints for operating systems, clients, and applications. The system uses the VDB to help determine whether a particular host increases your risk of compromise.

VDB versions

We periodically release VDB updates. Each update replaces the previous version.

VDB update duration

The time required to update the VDB and its associated mappings on the Firewall Management Center depends on the number of hosts in your network map. Allow approximately one minute for every 1000 hosts.

Scheduled VDB updates

Initial configuration automatically downloads and installs the latest VDB as a one-time operation. To update the VDB and deploy configurations automatically, create separate tasks. For more information, refer to Scheduling, on page 505.

Deployment after VDB updates

Deploy configuration changes after a VDB update for updated application detectors and operating system fingerprints to take effect. Updated vulnerability information takes effect without deploy.



Caution The first deploy after a VDB update usually restarts Snort because updated application detectors and operating system fingerprints require a restart. Restarting Snort briefly interrupts traffic flow and inspection on all devices, including those configured for high availability or scalability. Interface configurations determine whether traffic drops or passes without inspection during the interruption. When you deploy without restarting Snort, resource demands may result in a small number of packets dropping without inspection.

VDB release information

For VDB 343 and later, Cisco Secure Firewall Application Detectors provides searchable application detector information. The release notes describe changes in each VDB release.

VDB lite

For VDB 363 and later, the system installs a smaller VDB (also called VDB lite) on lower memory devices running Snort 2. This smaller VDB contains the same applications, but fewer detection patterns. Devices using the smaller VDB can miss some application identification versus devices using the full VDB.

Manually update the VDB

Perform an on-demand update to keep the VDB current .

Before you begin

If the Firewall Management Center cannot access the internet, download the VDB manually from https://www.cisco.com/go/firepower-software. Select or search for your model—or choose any model because all Firewall Management Centers use the same VDB—and then open the Coverage and Content Updates page.

Procedure

Step 1 Choose **System** (⚙) > **Content Updates** > **VDB Updates**.

In Versions 7.2.0–7.2.5, choose **System** (⚙) > **Updates** > **Content Updates**.

Step 2 Choose how you want to get the VDB onto the Firewall Management Center.

- Direct download: Click **Download Updates**.
- Manual upload: Click **Upload Update**, click **Choose File**, select the VDB, and click **Upload**.

Note
In Version 7.2.0–7.2.5, clicking **Download Updates** also immediately gets the latest maintenance release and the latest critical patches for your deployment.

Step 3 Install the VDB.

a) Click the **Install** icon next to the VDB update.
b) Choose the Firewall Management Center.
c) Click **Install**.

Monitor update progress in the Message Center. Do not use mapped-vulnerability features while the VDB is updating. If the Message Center shows no progress for several minutes or reports that the update failed, do not restart it. Contact Cisco TAC.

Step 4 Verify update success.

The current version is displayed on the VDB update page and on **Help** (?) > **About**.

The system uses the new vulnerability information. However, you must deploy before updated application detectors and operating system fingerprints can take effect.

What to do next

Deploy configuration changes.

Geolocation database (GeoDB) updates

The geolocation database (GeoDB) maps IP addresses to geographic locations so that you can view and filter traffic by location.

GeoDB versions

We periodically release GeoDB updates. Each update replaces the previous version.

Scheduled GeoDB updates

Initial configuration schedules weekly updates to keep geolocation information accurate. Review the schedule and adjust it as needed. Refer to Schedule GeoDB updates, on page 259.

Deployment after GeoDB updates

The Firewall Management Center automatically updates the GeoDB on managed devices, so deploy is usually unnecessary. If an update adds a new country, which is rare, deploy configuration changes as soon as possible. This allows the new country to count as part of its continent. For example, if an update adds Country to Continent, rules that filter based on "Continent" do not match traffic through Country until you deploy.

Deprecated IP package

We no longer provide the geolocation IP package, which contained contextual data associated with routable IP addresses. This saves disk space and does not affect geolocation rules or traffic handling in any way. Any contextual data is now stale, and upgrading to most later versions deletes the IP package. Options to download the IP package or view contextual data have no effect, and are removed in later versions.

Schedule GeoDB updates

Schedule recurring GeoDB updates to keep the installed geolocation information current.

Initial configuration creates a weekly schedule for GeoDB updates. Review the schedule and adjust the start time as needed.

Note The **IP Package Download** option available in some versions has no effect. We no longer provide an IP package. In later versions, this option is removed.

Procedure

Step 1 Choose **System** (⚙) > **Content Updates** > **Geolocation Updates**.

In Versions 7.2.0–7.2.5, choose **System** (⚙) > **Updates** > **Geolocation Updates**.

Step 2 Configure recurring GeoDB updates.

a) Under **Recurring Geolocation Updates**, check **Enable Recurring Weekly Updates from the Support Site**.

b) Specify the **Update Start Time**.

Step 3 Click **Save**.

Manually update the GeoDB

Perform an on-demand update to keep the GeoDB current.

Before you begin

If the Firewall Management Center cannot access the internet, download the GeoDB manually from https://www.cisco.com/go/firepower-software. Select or search for your model—or choose any model because all Firewall Management Centers use the same GeoDB—and then open the Coverage and Content Updates page.

Note The **IP Package Download** option available in some versions has no effect. We no longer provide an IP package. In later versions, this option is removed.

Procedure

Step 1 Choose **System** (⚙) > **Content Updates** > **Geolocation Updates**.

In Versions 7.2.0–7.2.5, choose **System** (⚙) > **Updates** > **Geolocation Updates**.

Step 2 Under **One-Time Geolocation Update**, choose how you want to get the GeoDB onto the Firewall Management Center.

- Direct download: Choose **Download and install geolocation update from the Support Site**.
- Manual upload: Choose **Upload and install geolocation update**, click **Choose File**, and select the GeoDB.

Step 3 Click **Import**.

Monitor update progress in the Message Center.

Step 4 Verify update success.

The current version is displayed on the GeoDB update page and on **Help** (?) > **About**.

The system uses the new geolocation information. However, you must deploy before new countries can take effect. Until you deploy, the new country does not count as part of its continent. For example, if an update adds Country to Continent, rules that filter based on "Continent" do not match traffic through Country until you deploy.

What to do next

If the update adds a new country, deploy configuration changes.

Intrusion rule updates

An intrusion rule update is a cumulative package that provides new and modified intrusion and preprocessor rules, and modifies the policies that use them.

Intrusion rule update versions

As new vulnerabilities become known, the Talos Intelligence Group releases intrusion rule updates. Keep intrusion rules up to date. Each update replaces the previous version, and you can import only an update that is newer than the currently installed version.

Contents of intrusion rule updates

An intrusion rule update can contain these types of content:

- Rules and rule states: An update can add, modify, or delete intrusion and preprocessor rules. The default state of a new rule can differ between system-provided intrusion policies. For example, a new rule can be enabled in the Security over Connectivity policy and disabled in the Connectivity over Security policy. An update can also change the default state of an existing rule.
- Rule categories: An update can contain new rule categories. Importing the update adds all new categories.
- Preprocessor and advanced settings: An update can change advanced settings in system-provided intrusion policies and preprocessor settings in system-provided network analysis policies. It can also change the default values of advanced preprocessing and performance options in access control policies.
- Variables: An update can change the default values of existing variables without overriding customized values. It also adds new variables.

Policy changes from intrusion rule updates

Intrusion rule updates can affect system-provided and custom network analysis and intrusion policies, as well as the access control policies that use them. The effects depend on the policy type.

- System-provided policies: Changes to system-provided network analysis and intrusion policies and advanced access control settings take effect after you redeploy the configuration.
- Custom policies: Every custom network analysis and intrusion policy is based directly or indirectly on a system-provided policy. For each custom policy, you can prevent rule updates from automatically applying changes from its base policy. You can then apply the base-policy changes manually on a schedule independent of rule update imports. Updates do not override customized settings.



Caution Importing an intrusion rule update discards all cached changes to network analysis and intrusion policies. Before you import an update, use the Rule Updates page to identify policies with cached changes and the users who made them.

Scheduled intrusion rule updates

Initial configuration creates a daily schedule for intrusion rule updates. Review the schedule and adjust the frequency as needed. Refer to Schedule intrusion rule updates, on page 262.

Deployment after intrusion rule updates

Deploy configuration changes after an intrusion rule update for the update's changes to take effect. You can configure an import to deploy automatically to affected devices. Automatic deployment is especially useful when you allow updates to modify system-provided base intrusion policies.

Caution Although a rule update by itself does not restart Snort when you deploy, other changes you have made may. Restarting Snort briefly interrupts traffic flow and inspection on all devices, including those configured for high availability or scalability. Interface configurations determine whether traffic drops or passes without inspection during the interruption. When you deploy without restarting Snort, resource demands may result in a small number of packets dropping without inspection.

Schedule intrusion rule updates

Schedule recurring intrusion rule updates to keep the installed intrusion rules current.

Before you begin

Make sure your process for updating intrusion rules complies with your security policies.

Initial configuration creates a daily schedule for intrusion rule updates. Review the schedule and adjust the frequency as needed. You can also configure automatic deployment after each rule update. Automatic deployment can disrupt traffic inspection and flow.

Procedure

Step 1 Choose **System** (⚙) > **Content Updates** > **Rule Updates**.

In Versions 7.2.0–7.2.5, choose **System** (⚙) > **Updates** > **Rule Updates**.

Step 2 Configure recurring rule updates.

a) Under **Recurring Rule Update Imports**, check **Enable Recurring Rule Update Imports**.
b) Specify the **Import Frequency** and start time.
c) (Optional) Check **Deploy all policies to targeted devices after rule update completes**.

Step 3 Click **Save**.

Manually update intrusion rules

Perform an on-demand update to keep intrusion rules current.

Before you begin

- Make sure your process for updating intrusion rules complies with your security policies.
- If the Firewall Management Center cannot access the internet, download the update manually from https://www.cisco.com/go/firepower-software. Select or search for your model—or choose any model because all Firewall Management Centers use the same update—and then open the Coverage and Content Updates page.

Procedure

Step 1 Choose **System** (⚙) > **Content Updates** > **Rule Updates**.

In Versions 7.2.0–7.2.5, choose **System** (⚙) > **Updates** > **Rule Updates**.

Step 2 Configure the one-time rule update.

a) Under **One-Time Rule Update/Rules Import**, choose how to update intrusion rules.

- Direct download: Choose **Download new rule update...**.
- Manual upload: Choose **Rule update or text rule file...**, click **Choose File**, and select the intrusion rule update.

b) (Optional) Check **Reapply all policies after the rule update import completes**.

Step 3 Click **Import**.

Monitor update progress in the Message Center. Even if the Message Center shows no progress for several minutes or indicates that the update has failed, do not restart the update. Instead, contact Cisco TAC.

Step 4 Verify update success.

The current version is displayed on the rule update page and on **Help** (?) > **About**.

What to do next

If you did not deploy configuration changes as part of the update, deploy them now.

Guidelines for importing custom Snort 2 rules

Use these guidelines to import custom standard text rules for use in Snort 2 intrusion policies.

File and rule syntax requirements

Import rules using a rule file that meets these requirements:

- Use a plain text file encoded in ASCII or UTF-8.
- Use only alphanumeric characters, spaces, underscores (_), periods (.), and hyphens (-) in the file name.
- To import a rule in the deleted state, precede it with one number sign (#). A rule preceded by two number signs (##) is not imported.
- Do not use escape characters in rules.
- Limit each source or destination port list to 64 characters. A longer list causes the import to fail.

For more information on writing custom rules, refer to *Custom Snort 2 Intrusion Policies for Access Control* and to the Snort manual at http://www.snort.org.

SID, GID, and revision number requirements

Follow these requirements for SIDs, GIDs, and revision numbers:

- Do not specify a GID. A standard text rule receives GID 1. In a multidomain deployment, a rule imported into a subdomain instead receives a domain-specific GID from 1000 through 2000.
- Leave the SID and revision number unspecified when importing a rule for the first time. The rule receives an available SID of 1000000 or greater and a revision number of 1.
- If you must specify a SID for a new rule, use an unassigned value from 1000000 through 2147483647 (the 32-bit limit).
- When importing an updated or previously deleted rule, include its assigned SID and a revision number greater than its current revision number.

To determine the current revision number, edit the imported rule. Deleting a custom rule moves it to the deleted rule category and increments its revision number. You can then import a later revision to reinstate the rule.

Import custom Snort 2 rules

Import custom standard text rules so you can enable them in Snort 2 intrusion policies.

This procedure applies only to custom standard text rules for Snort 2 intrusion policies. In a multidomain deployment, you can import rules in any domain, and can view rules imported in the current domain and ancestor domains. For Snort 3, refer to the rule-tuning topics in *Custom Snort 3 Intrusion Policies for Access Control*.

Before you begin

- Make sure your custom rule file follows the guidelines in Guidelines for importing custom Snort 2 rules, on page 263.
- Make sure your process for importing custom rules complies with your security policies.

Procedure

Step 1 Choose **System** (⚙) > **Content Updates** > **Rule Updates**.

In Versions 7.2.0–7.2.5, choose **System** (⚙) > **Updates** > **Rule Updates**.

You can also click **Import Rules** in the intrusion rules editor (**Objects** > **Intrusion Rules**).

Step 2 (Optional) Delete all existing custom rules.

To replace all existing custom rules with the rules in the new file, click **Delete All Local Rules**. The deleted rules move to the deleted rule category, and their revision numbers increase.

Step 3 Import the rule file.

a) Under **One-Time Rule Update/Rules Import**, choose **Rule update or text rule file to upload and install**.

b) Click **Choose File**, select your rule file, and click **Import**.

You can monitor import progress in the Message Center. Even if the Message Center shows no progress for several minutes or indicates that the import has failed, do not restart the import. Instead, contact Cisco TAC.

- Imported rules are added to the local rule category in a disabled state.
- Imported rules receive Generator ID (GID) 1 unless imported into a subdomain, where they receive a domain-specific GID from 1000 through 2000.
- New imported rules receive an available Snort ID (SID) of 1000000 or greater and a revision of 1. In a multidomain deployment, SIDs assigned within a domain might not be sequential when multiple users import rules concurrently.

What to do next

- Edit Snort 2 intrusion policies and enable the rules you imported.

  Do not enable an imported rule that uses the deprecated `threshold` keyword in an intrusion policy that also uses intrusion event thresholding. This combination causes policy validation to fail.
- Deploy configuration changes.

Intrusion rule update logs

The Rule Update Log on the Rule Updates page records each intrusion rule update and custom Snort 2 rule import, including its time, user, and status. Each log also identifies the affected rules and components. Deleting a log does not delete the imported objects.

In a multidomain deployment, you can view logs for the current domain and its descendant domains. You cannot view logs from ancestor or sibling domains.

Intrusion rule update log fields

This table explains the fields in intrusion rule update logs.

Table 10: Intrusion rule update log fields

Field	Description
Action	Identifies how the import affected the object.
Default action	For a rule, the default action defined by the update is Pass, Alert, or Drop. The field is blank for other object types.
Details	Identifies the affected component or rule. For a changed rule, the field displays the GID, SID, and previous revision number in `GID:SID:Rev` format. The field is blank for a rule that has not changed.
Domain	The domain whose intrusion policies can use the updated rule. Intrusion policies in descendant domains can also use the rule. This field is only present in a multidomain deployment.

Field	Description
GID	The generator ID for a rule. For example: • A GID of 1 identifies a standard text rule. In a multidomain deployment, GID 1 is reserved for the Global domain. • A GID of 3 identifies a shared object rule. • A GID from 1000 through 2000 identifies a rule imported into a subdomain.
Name	The name of the imported object. For a rule, the name corresponds to the rule's Message field. For a rule update component, the field displays the component name.
Policy	For an imported rule, the field displays `All` when the import succeeds and the rule can be enabled in all applicable system-provided default intrusion policies. The field is blank for other object types.
Rev	The revision number for a rule.
Rule update	The rule update file name.
SID	The Snort ID for a rule.
Time	The time and date the import began.
Type	Identifies the imported object as a rule update component, rule, or policy apply. A policy apply record indicates that the option to reapply all policies after the import was enabled.

Action values

The Action field in intrusion rule update logs uses these values.

Table 11: Action values in intrusion rule update logs

Action	Meaning
new	The rule was stored for the first time.
changed	A rule update component was modified, or a rule was imported with a higher revision number and the same GID and SID.
collision	The import was skipped because the revision conflicts with an existing component or rule.
deleted	The rule was deleted from the rule update.
enabled	A rule, preprocessor, or other feature was enabled in a system-provided default policy.
disabled	A rule was disabled in a system-provided default policy.
drop	A rule was set to Drop and Generate Events in a system-provided default policy.
error	The update or import failed.

Action	Meaning
apply	The option to reapply all policies after the import was enabled.

Searching intrusion rule update logs

A search from the detailed log view searches the entire database, not only the selected import. Use the time range to limit the results to relevant records.

Maintain your air-gapped deployment

Manually obtain and install essential content updates when the Firewall Management Center cannot access the internet.

Refer to:

- Manually update the VDB, on page 257
- Manually update intrusion rules, on page 262
- Manually update the GeoDB, on page 259

History for content updates

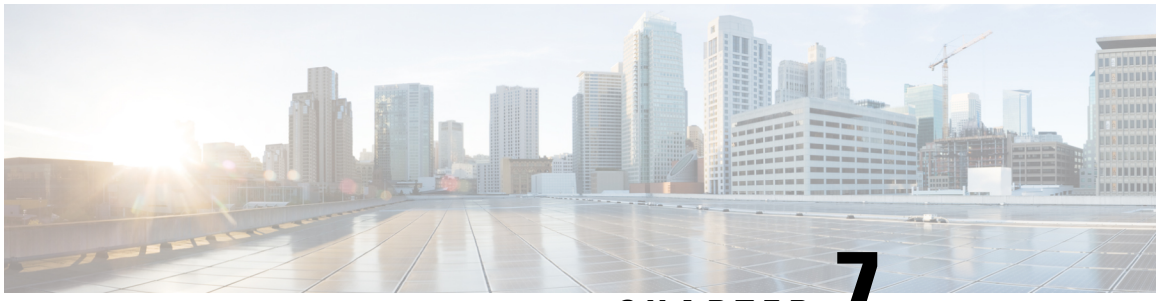
This table provides the feature history for content updates.

Table 12: History for content updates

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Content updates and product upgrades no longer share a page.	7.2.6 7.4.1	Any	Content updates and product upgrades no longer share a page. • **System**(⚙) > **Content Updates** is where you update intrusion rules, the VDB, and the GeoDB. • **System**(⚙) > **Product Upgrades** is where you upgrade the Firewall Management Center and all managed devices, as well as manage upgrade packages. • **System**(⚙) > **Updates** is deprecated. All Firewall Threat Defense upgrades now use the wizard.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Custom Snort 2 intrusion rule import warns when rules collide.	6.7.0	Any	The system now warns you of rule collisions when you import custom Snort 2 intrusion rules. Previously, the system would silently skip the rules that cause collisions—with the exception of Version 6.6.0.1, where a rule import with collisions would fail entirely. On the Rule Updates page, if a rule import had collisions, a warning icon is displayed in the Status column. For more information, hover your pointer over the warning icon and read the tooltip. Note that a collision occurs when you try to import an intrusion rule that has the same SID/revision number as an existing rule. You should always make sure that updated versions of custom rules have new revision numbers. New/modified screens: We added a warning icon to **System**(⚙) > **Updates** > **Rule Updates**.
Automatic VDB update during initial setup.	6.6.0	Any	When you set up a new or reimaged Firewall Management Center, the system attempts to update the vulnerability database (VDB). This is a one-time operation. If the Firewall Management Center has internet access, we recommend you schedule tasks to perform automatic recurring VDB update downloads and installations.
Automatic intrusion rule updates.	6.6.0	Any	Initial setup enables daily intrusion rule updates. We recommend you review this task and adjust if necessary. For the updated rules to take effect you must deploy configurations.
Automatic software downloads and GeoDB updates.	6.5.0	Any	When you set up a new or reimaged Firewall Management Center, the system schedules weekly patch downloads and GeoDB updates.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Signed SRU, VDB, and GeoDB updates.	6.4.0	Any	So the system can verify that you are using the correct update files, Version 6.4+ uses signed updates for intrusion rules (SRU), the vulnerability database (VDB), and the geolocation database (GeoDB). Earlier versions continue to use unsigned updates. Unless you manually download updates, for example, in an air-gapped deployment—you should not notice any difference in functionality. If, however, you do manually download and install SRU, VDB, and GeoDB updates, make sure you download the correct package for your current version. Signed update files begin with 'Cisco' instead of 'Sourcefire,' and terminate in .sh.REL.tar instead of .sh, as follows: • SRU: Cisco_Firepower_SRU-*date-build*-vrt.sh.REL.tar • VDB: Cisco_VDB_Fingerprint_Database-4.5.0-*version*.sh.REL.tar • GeoDB: Cisco_GEODB_Update-*date-build*.sh.REL.tar We will provide both signed and unsigned updates until the end-of-support for versions that require unsigned updates. Do not untar signed (.tar) packages. If you accidentally upload a signed update to an older Firewall Management Center or ASA FirePOWER device, you must manually delete it. Leaving the package takes up disk space, and also may cause issues with future upgrades.
Snort restart warnings before VDB updates.	6.2.3	Any	The system now warns you VDB updates restart the Snort process. This interrupts traffic inspection and, depending on how the managed device handles traffic, possibly interrupts traffic flow. You can cancel the install until a more convenient time, such as during a maintenance window. These warnings can appear: • After you download and manually install a VDB. • When you create a scheduled task to install the VDB. • When the VDB installs in the background, such as during a previously scheduled task or as part of a software upgrade.
Deprecated: Geolocation details	6.2.3	Any	We no longer provide the geolocation IP package, which contained contextual data associated with routable IP addresses. This saves disk space and does not affect geolocation rules or traffic handling in any way. Any contextual data is now stale, and upgrading to most later versions deletes the IP package. Options to download the IP package or view contextual data have no effect, and are removed in later versions.



CHAPTER 7

Licenses

This chapter provides in-depth information about the different license types, service subscriptions, licensing requirements and more.

- Licenses, on page 271
- Prerequisites for licensing, on page 283
- Create a Cisco account, on page 286
- Create a smart account and add licenses, on page 287
- Configure Smart Licensing, on page 288
- Specific License Reservation, on page 299
- Configure Legacy Firewall Management Center PAK-Based Licenses, on page 310
- Additional information about licensing, on page 311
- History for licenses, on page 311

Licenses

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 software 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.

Smart Software Manager and accounts

When you purchase one or more licenses, you manage them in the Smart Software Manager: https://software.cisco.com/#module/SmartLicensing. The Smart Software Manager lets you create a primary

account for your organization. If you do not yet have an account, click the link to set up a new account. For instructions, refer to Create a Cisco Account.

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 devices.

You manage licenses by virtual account. Only that virtual account's devices 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 devices between virtual accounts.

Licensing options for air-gapped deployments

This table compares the available licensing options for environments without internet access. Your sales representative may have additional advice for your specific situation.

Table 13: Comparison of licensing options for air-gapped networks

Smart Software Manager On-Prem	Specific License Reservation
Scalable for a large number of products	Best for a small number of devices
Automated licensing management, usage and asset management visibility	Limited usage and asset management visibility
No incremental operational costs to add devices	Linear operational costs over time to add devices
Flexible, easier to use, less overhead	Significant administrative and manual overhead for moves, adds, and changes
Out-of-compliance status is allowed initially and at various expiration states	Out-of-compliance status impacts system functioning
For more information, refer to Register the Firewall Management Center with the Smart Software Manager On-Prem, on page 291	For more information, refer to Specific License Reservation, on page 299

Licensing for Management Center and devices

The Firewall Management Center registers with the Smart Software Manager, and then assigns licenses for each managed device. Devices do not register directly with the Smart Software Manager.

A physical Firewall Management Center does not require a license for its own use. The Firewall Management Center Virtual does require a platform license.

Periodic communications with the Smart Software Manager

You use a Product Instance Registration Token to register the Firewall Management Center with the Smart Software Manager. The Smart Software Manager issues an ID certificate for communication between the Firewall Management Center and the Smart Software Manager. This certificate is valid for one year, although it will be renewed every six months. If an ID certificate expires (after a year with no communication), the Firewall Management Center may be removed from your account.

The Firewall Management Center communicates with the Smart Software Manager on a periodic basis. If you make changes in the Smart Software Manager, you can refresh the authorization on the Firewall Management Center so the changes immediately take effect. You also can wait for the Firewall Management Center to communicate as scheduled. Your Firewall Management Center must either have direct internet access to the Smart Software Manager, or use one of the options described in Licensing options for air-gapped deployments, on page 272.

In non-airgapped deployments, normal license communication occurs every 30 days. With the grace period, your Firewall Management Center will operate for up to 90 days without contacting the Smart Software Manager. Ensure that the Firewall Management Center contacts the Smart Software Manager within 90 days, or the Firewall Management Center will revert to an unregistered state.

If your device is unable to communicate due to network issues, it continues working for up to 90 days. Once communication resumes and the ID certificate is validated, license entitlement is restored.

Evaluation mode

Evaluation mode is a temporary license that

- operates for 90 days before the Firewall Management Center registers with the Smart Software Manager,
- allows you to assign feature licenses to managed devices while they remain in compliance for the duration of evaluation mode, and
- ends when the Firewall Management Center becomes unregistered after the 90-day period expires.

If you register the Firewall Management Center with the Smart Software Manager, the evaluation mode ends. If you later deregister the Firewall Management Center, you cannot resume evaluation mode, even if you did not initially use all 90 days.

For more information about the unregistered state, refer to Unregistered state, on page 274.

Note You cannot receive an evaluation license for Strong Encryption (3DES/AES); you must register with the Smart Software Manager to receive the export-compliance token that enables the Strong Encryption (3DES/AES) license.

Note The Talos certificate for Evaluation Mode in Secure Firewall Version 7.6.0 is set to expire on March 31, 2025. After this date, access to Talos-hosted services in Evaluation Mode (specifically those related to web reputation/categorization lookups) will be discontinued.

Out-of-compliance state

The Firewall Management Center can become out of compliance in these situations:

- Over-utilization—When the managed devices or the Firewall Management Center Virtual uses unavailable licenses.

Note Even if one license is unavailable for a device, the Firewall Management Center will be in the **Out-of-Compliance** state. Deployment of all configurations will work. For example, you have two Malware licenses and three managed devices. All three devices will be in the **Out-of-Compliance** state, and Malware feature will not work. This behavior is applicable for all licenses.

- License expiration—When a managed device term-based license expires.

In an out-of-compliance state, these effects occur:

- Firewall Management Center Virtual platform license—Operation is not affected.
- All managed device licenses—Operation is not affected.

After you resolve the licensing problem, the Firewall Management Center will show that it is now in compliance after its regularly scheduled authorization with the Smart Software Manager. To force an authorization, click **Re-Authorize** on the **System** (⚙) > **Licenses** > **Smart Licenses** page.

Unregistered state

The Firewall Management Center can become unregistered in these situations:

- Evaluation mode expiration—Evaluation mode expires after 90 days.
- Manual deregistration of the Firewall Management Center
- Lack of communication with the Smart Software Manager—The Firewall Management Center does not communicate with the Smart Software Manager for 1 year. Note: After 90 days, the Firewall Management Center authorization expires, but it can successfully resume communication within one year to automatically re-authorize. After a year, the ID certificate expires, and the Firewall Management Center is removed from your account so you will have to manually re-register the Firewall Management Center.

In an unregistered state, the Firewall Management Center cannot deploy any configuration changes to devices *for features that require licenses*.

End-user license agreement

The Cisco end-user license agreement (EULA) and any applicable supplemental agreement (SEULA) that governs your use of this product are available from http://www.cisco.com/go/softwareterms.

License types and restrictions

This section describes the types of licenses available for Secure Firewall.

Table 14: Smart Licenses

License You Assign	Subscription You Purchase	Duration	Granted Capabilities
Base	Based on license type	Perpetual or Subscription **Note** Base subscription licenses are supported only on Firewall Threat Defense Virtual .	Except for Specific License Reservation and the Secure Firewall 3100 , Base perpetual licenses are automatically assigned with all Firewall Threat Defense s. User and application control Switching and routing NAT For details, see Base licenses, on page 277 .
Threat	• T • TC (Threat + URL) • TMC (Threat + Malware defense + URL)	Subscription	Intrusion detection and prevention File control Security Intelligence filtering For details, see Threat licenses, on page 278
Malware defense	• TM (Threat + Malware defense) • TMC (Threat + Malware defense + URL) • AMP	Subscription	Malware defense Secure Malware Analytics File storage For details, see Malware Defense Licenses, on page 277 and *License Requirements for File and Malware Policies* in the *Cisco Secure Firewall Management Center Device Configuration Guide* .
URL Filtering	• TC (Threat + URL) • TMC (Threat + Malware defense + URL) • URL	Subscription	Category and reputation-based URL filtering For details, see URL Filtering licenses, on page 279 .
Firewall Management Center Virtual	Based on license type	• Regular Smart Licensing—Perpetual • Specific License Reservation—Subscription	The platform license determines the number of devices the Firewall Management Center Virtual can manage. For details, see Firewall Management Center Virtual licenses, on page 276 .

License You Assign	Subscription You Purchase	Duration	Granted Capabilities
Export-Controlled Features	No subscription required	Perpetual	Features that are subject to national security, foreign policy, and anti-terrorism laws and regulations; see License for export-controlled functionality, on page 280 .
Remote Access VPN: • AnyConnect Apex • AnyConnect Plus • AnyConnect VPN Only	Based on license type	Subscription or perpetual	Remote access VPN configuration. Your account must allow export-controlled functionality to configure remote access VPN. You select whether you meet export requirements when you register the device. The Firewall Threat Defense can use any valid AnyConnect Client license. The available features do not differ based on license type. For more information, see AnyConnect Client licenses, on page 279 and *VPN Licensing* in the *Cisco Secure Firewall Management Center Device Configuration Guide* .

Note Subscription licenses are term-based licenses.

Firewall Management Center Virtual licenses

The Firewall Management Center Virtual requires a platform license that correlates with the number of devices it can manage.

The Firewall Management Center Virtual supports Smart Licensing.

In regular Smart Licensing, these licenses are perpetual.

In Specific License Reservation (SLR), these licenses are subscription-based.

Starting from Firewall Management Center version 7.3, in regular Smart Licensing and Specific License Reservation, the licenses can be perpetual or subscription licenses. The subscription term can be 3,5, or 10 years. Select the type of license based on your requirement. By default, subscription licenses are selected.

Note For the add-on license requirements of your new devices on Firewall Management Center Virtual, it is recommended to migrate to a higher Firewall Management Center Virtual model that supports additional devices.

Base licenses

The Base license allows you to configure and manage firewall devices with essential networking and security capabilities, including switching, routing, high availability, clustering, user control, application control, and database updates.

The Base license allows you to:

- Configure your devices to perform switching and routing (including DHCP relay and NAT)
- Configure devices as a high availability pair
- Configure clustering
- Implement user and application control by adding user and application conditions to access control rules
- Update the Vulnerability database (VDB) and geolocation database (GeoDB).
- Download intrusion rules such as SRU/LSP. However, you cannot deploy access control policy or rules that have intrusion policy to the device unless Threat license is enabled.

Secure Firewall 3100

You obtain a Base license when you purchase the Secure Firewall 3100.

Other Models

Except in deployments using Specific License Reservation, a Base license is automatically added to your account when you register a device to the Firewall Management Center. For Specific License Reservation, you need to add the Base license to your account.

Malware Defense Licenses

A Malware defense license lets you perform Malware Defense and Secure Malware Analytics. With this feature, you can use devices to detect and block malware in files transmitted over your network.

To support this feature license, you can purchase the Malware defense (AMP) service subscription as a stand-alone subscription or in combination with Threat (TM) or Threat and URL Filtering (TMC) subscriptions. Threat license is a prerequisite for a Malware defense license.



Note Managed devices with Malware defense licenses enabled periodically attempt to connect to the Secure Malware Analytics Cloud even if you have not configured dynamic analysis. Because of this, the device's Interface Traffic dashboard widget shows transmitted traffic; this is expected behavior.

You configure Malware Defense as part of a file policy, which you then associate with one or more access control rules. File policies can detect your users uploading or downloading files of specific types over specific application protocols. Malware Defense allows you to use local malware analysis and file preclassification to inspect a restricted set of those file types for malware. You can also download and submit specific file types to the Secure Malware Analytics Cloud for dynamic and Spero analysis to determine whether they contain malware. For these files, you can view the network file trajectory, which details the path the file has taken through your network. The Malware license also allows you to add specific files to a file list and enable the file list within a file policy, allowing those files to be automatically allowed or blocked on detection.

Note that a Malware defense license is required only if you deploy Malware Defense and Secure Malware Analytics. Without a Malware defense license, the Firewall Management Center can receive Secure Endpoint malware events and indications of compromise (IOC) from the Secure Malware Analytics Cloud.

Refer to important information at *License Requirements for File and Malware Policies* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

When you disable this license:

- 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 Malware Defense configurations.
- For a very brief time after a Malware defense 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.

If the license expires, your entitlement for the above capabilities ceases and the Firewall Management Center moves to the out-of-compliance state.

Threat licenses

A Threat license allows you to perform intrusion detection and prevention, file control, and Security Intelligence filtering.

License capabilities

- *Intrusion detection and prevention* allows you to analyze network traffic for intrusions and exploits and, optionally, drop offending packets.
- *File control* allows you to detect and, optionally, block users from uploading (sending) or downloading (receiving) files of specific types over specific application protocols. *Malware defense*, which requires a Malware defense license, allows you to inspect and block a restricted set of those file types based on their dispositions.
- *Security Intelligence filtering* allows you to block —deny traffic to and from—specific IP addresses, URLs, and DNS domain names, before the traffic is subjected to analysis by access control rules. Dynamic feeds allow you to immediately block connections based on the latest intelligence. Optionally, you can use a "monitor-only" setting for Security Intelligence filtering.

License purchase options

You can purchase a Threat license as a stand-alone subscription (T) or in combination with URL Filtering (TC), Malware defense (TM), or both (TMC).

Disabled license behavior

When you disable this license:

- The Firewall Management Center stops acknowledging intrusion and file events from the affected devices. As a consequence, correlation rules that use those events as a trigger criteria stop firing.
- The Firewall Management Center does not contact the internet for either Cisco-provided or third-party Security Intelligence information.

- You cannot re-deploy existing intrusion policies until you re-enable Threat.

Expired license behavior

If the license expires, your entitlement for the above capabilities ceases and the Firewall Management Center moves to the out-of-compliance state.

URL Filtering licenses

The URL Filtering license allows you to write access control rules that determine the traffic that can traverse your network based on URLs requested by monitored hosts, correlated with information about those URLs. To support this feature license, you can purchase the URL Filtering service subscription as a stand-alone subscription or in combination with Threat (TC) or Threat and Malware defense (TMC) subscriptions. Threat license is a prerequisite for this license.

Tip Without a URL Filtering license, you can specify individual URLs or groups of URLs to allow or block. This option gives you granular, custom control over web traffic, but does not allow you to use URL category and reputation data to filter network traffic.

Although you can add category and reputation-based URL conditions to access control rules without a URL Filtering license, the Firewall Management Center will not download URL information. You cannot deploy the access control policy until you first add a URL Filtering license to the Firewall Management Center, then enable it on the devices targeted by the policy.

When you disable this license:

- You may lose access to filtering network traffic based on the URL category and reputation ACP rules. The manual URL filtering options will continue to be supported.
- Access control rules with URL conditions immediately stop filtering URLs.
- Your Firewall Management Center can no longer download updates to URL data.
- You cannot re-deploy existing access control policies if they include rules with category and reputation-based URL conditions.

If the license expires, your entitlement for the above capabilities ceases and the Firewall Management Center moves to the out-of-compliance state.

AnyConnect Client licenses

You can configure remote access VPN using the AnyConnect Client and standards-based IPSec/IKEv2.

To enable remote access VPN, you must purchase and enable one of these licenses: AnyConnect Plus, AnyConnect Apex, or AnyConnect VPN Only. You can select AnyConnect Plus and AnyConnect Apex if you have both licenses and you want to use them both. The AnyConnect VPN Onlylicense cannot be used with **Apex** or **Plus**. The AnyConnect Client license must be shared with the Smart Account. For more instructions, refer to http://www.cisco.com/c/dam/en/us/products/collateral/security/anyconnect-og.pdf.

You cannot deploy the remote access VPN configuration to the device if the specified device does not have the entitlement for a minimum of one of the specified AnyConnect Client license types. If the registered license moves out of compliance or entitlements expire, the system displays licensing alerts and health events.

While using remote access VPN, your Smart Account must have the export controlled features (strong encryption) enabled. The Firewall Threat Defense requires strong encryption (which is higher than DES) for successfully establishing remote access VPN connections with AnyConnect Clients.

You cannot deploy remote access VPN if these conditions are true:

- Smart Licensing on the Firewall Management Center is running in evaluation mode.
- Your Smart Account is not configured to use export-controlled features (strong encryption).

License for export-controlled functionality

This topic provides information about software features that are subject to export control laws and regulations, and describes how to determine and enable export-controlled functionality for your system.

Features that require export-controlled functionality

Certain software features are subject to national security, foreign policy, and anti-terrorism laws and regulations. These export-controlled features include:

- Security certifications compliance
- Remote access VPN
- Site-to-site VPN with strong encryption
- SSH platform policy with strong encryption
- SSL policy with strong encryption
- Functionality such as SNMPv3 with strong encryption

Verify export-controlled functionality in your system

To determine whether export-controlled functionality is currently enabled for your system: Choose **System** (⚙) > **Licenses** > **Smart Licenses** and see if **Export-Controlled Features** displays **Enabled**.

Enable export-controlled functionality

If **Export-Controlled Features** shows **Disabled** and you want to use features that require strong encryption, there are two ways to enable strong cryptographic features. Your organization may be eligible for one or the other (or neither), but not both.

- If there is *no* option to enable export-controlled functionality when you generate a new Product Instance Registration Token in the Smart Software Manager, contact your account representative.

  When approved by Cisco, you can manually add a strong encryption license to your account so you can use export-controlled features. For more information, see Enable the export control feature for accounts without global permission, on page 293

- If the option "Allow export-controlled functionality on the products registered with this token" appears when you generate a new Product Instance Registration Token in the Smart Software Manager, make sure you check it before generating the token.

  If you did not enable export-controlled functionality for the Product Instance Registration Token that you used to register the Firewall Management Center, then you must deregister and then re-register the

Firewall Management Center using a new Product Instance Registration Token with export-controlled functionality enabled.

If you registered devices to the Firewall Management Center in evaluation mode or before you enabled strong encryption on the Firewall Management Center, reboot each managed device to make strong encryption available. In a high availability deployment, the active and standby devices must be rebooted together to avoid an Active-Active condition.

The entitlement is perpetual and does not require a subscription.

More information

For general information about export controls, see https://www.cisco.com/c/en/us/about/legal/global-export-trade.html.

Firewall Threat Defense Virtual Licenses

Learn about the performance-tiered license entitlements for Firewall Threat Defense Virtual, including supported vCPU and memory configurations, compatible VM resource footprints, and RA VPN session limits by entitlement tier.

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 15: 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
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

Review and follow these guidelines and limitations 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 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 and later, 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 Defense , 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.
- You assign feature licenses to the cluster as a whole, not to individual nodes. However, each node of the cluster consumes a separate license for each feature. The clustering feature itself does not require any licenses.
- 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.

Edit performance tier of Firewall Threat Defense Virtual

The Firewall Threat Defense Virtual supports performance-tiered licensing that provides different throughput levels and VPN connection limits based on deployment requirements. You can edit the performance tier using this procedure.

Before you begin

Ensure that you review Firewall Threat Defense Virtual Performance Tier Licensing Guidelines and Limitations , on page 281.

Procedure

Step 1 Choose **System** (⚙) > **Licenses** > **Smart Licenses**.

Step 2 Click **Edit Performance Tier**.

Step 3 Choose a tier from the **Target Performance Tier** drop-down list.

Step 4 From the **Available Devices** list, select the required devices and click **Add**.

Step 5 Click **Apply**.

Prerequisites for licensing

This topic lists the prerequisites, supported domains, and user roles required for licensing.

General prerequisites

- Make sure NTP is configured on the Firewall Management Center and managed devices. Time must be synchronized for registration to succeed.
- For a Firepower 4100/9300, you must configure NTP on the chassis using the same NTP server for the chassis as for the Firewall Management Center.

  For Specific License Reservation requirements, see Prerequisites for Specific License Reservation, on page 299.

Supported domains

Global, except where indicated.

User roles

Admin

Prerequisites for licensing in high availability, clustering, and multi-instance

This section describes the licensing requirements for High Availability (for device High Availability and also Firewall Management Center Virtual High Availability), clustering, and multi-instance deployments.

Licensing for Firewall Management Center high availability

Ensure proper licensing is configured before implementing Firewall Management Center high availability.

License assignment examples for high availability configurations

Example: If you want to enable advanced malware protection for two devices managed by a Firewall Management Center pair, buy two Malware licenses and two TM subscriptions, register the active Firewall Management Center with the Smart Software Manager, then assign the licenses to the two devices on the active Firewall Management Center.

Hardware management center license requirements

No special license is required for hardware Firewall Management Centers in a high availability pair.

Virtual management center license requirements

You need two identically licensed Firewall Management Center Virtuals. Firewall Management Center Virtual requires an individual license for each device it manages.

Each registered device in a Firewall Management Center Virtual high availability pair must have its own license. Although only the active FMC registers with Smart Software Manager at any time, entitlement usage accounts for both Firewall Management Center Virtuals in the high availability pair.

Entitlement calculation for Firewall Management Center Virtual high availability pairs is as follows:

Firewall Management Center Virtual entitlement = Number of Firewall Management Center Virtual peers in the high availability pair x Number of managed devices

In Specific License Reservation deployments, only the primary Firewall Management Center requires a Specific License Reservation.

Example: For the Firewall Management Center Virtual high availability pair managing 2 devices, you can use:

- Two (2) Firewall Management Center Virtual entitlements
- 2 device licenses

Note Firepower MCv Device Licenses seen on Smart Software Manager is 4

Example: For two Firewall Management Center Virtual high availability pair managing 2 devices each, you can use:

- Four (4) Firewall Management Center Virtual entitlements
- 4 device licenses

Note Firepower MCv Device Licenses seen on Smart Software Manager is 8

If you break the high availability pair, the Firewall Management Center Virtual entitlements associated with the secondary Firewall Management Center Virtual are released. (In the example, you would then have two standalone Firewall Management Center Virtual devices.)

License requirements for Firewall Threat Defense device high availability

Both Firewall Threat Defense units in a high availability configuration must have the same licenses.

High availability configurations require two license entitlements: one for each device in the pair.

Before high availability is established, it does not matter which licenses are assigned to the secondary/standby device. During high availability configuration, the Firewall Management Center releases any unnecessary licenses assigned to the standby unit and replaces them with identical licenses assigned to the primary/active unit.

For example, if the active unit has a Base license and a Threat license, and the standby unit has only a Base license, the Firewall Management Center communicates with the Smart Software Manager to obtain an

available Threat license from your account for the standby unit. If your license account does not include enough purchased entitlements, your account becomes Out-of-Compliance until you purchase the correct number of licenses.

License for device clusters

Device clusters require specific licensing considerations for proper operation and feature functionality.

Performance tier license requirements

Each Firewall Threat Defense Virtual cluster node requires the same performance tier license. We recommend using the same number of CPUs and memory for all members, or else performance will be limited on all nodes to match the least capable member. The throughput level will be replicated from the control node to each data node so they match.

Feature license assignment

You assign feature licenses to the cluster as a whole, not to individual nodes. However, each node of the cluster consumes a separate license for each feature. The clustering feature itself does not require any licenses.

When you add the control node to the Firewall Management Center, you can specify the feature licenses you want to use for the cluster. Before you create the cluster, it doesn't matter which licenses are assigned to the data nodes; the license settings for the control node are replicated to each of the data nodes. You can modify licenses for the cluster by clicking **Devices** > **Device Management**, clicking **Edit** ( ) for the cluster, and then in the **License** area, clicking **Edit** ( ).

Note If you add the cluster before the Firewall Management Center is licensed (and running in Evaluation mode), then when you license the Firewall Management Center, you can experience traffic disruption when you deploy policy changes to the cluster. Changing to licensed mode causes all data units to leave the cluster and then rejoin.

Note APEX, PLUS, and VPN-Only licenses are not supported for clustering. Before you create a cluster, remove unsupported license assignments from the devices that will become cluster nodes.

Licensing for multi-instance deployments

This topic provides licensing information specific to multi-instance deployments where multiple container instances require appropriate license allocation and management.

All licenses are consumed per security engine/chassis (for the Firepower 4100) or per security module (for the Firepower 9300), and not per container instance. See the details:

- Base licenses are automatically assigned: one per security module/engine.
- Feature licenses are manually assigned to each instance; but you only consume one license per feature per security module/engine. For example, for the Firepower 9300 with 3 security modules, you only need one URL Filtering license per module for a total of 3 licenses, regardless of the number of instances in use.

For example:

Table 16: Sample license usage for container instances on a Firepower 9300

Firepower 9300	Instance	Licenses
Security Module 1	Instance 1	Base, URL Filtering, Malware
	Instance 2	Base, URL Filtering
	Instance 3	Base, URL Filtering
Security Module 2	Instance 4	Base, Threat
	Instance 5	Base, URL Filtering, Malware, Threat
Security Module 3	Instance 6	Base, Malware, Threat
	Instance 7	Base, Threat

Table 17: Total number of licenses

Base	URL Filtering	Malware	Threat
3	2	3	2

Create a Cisco account

You need a Cisco account to access Smart Account functionality and license any Cisco products.

You must have a Cisco account to request a Smart Account and license any Cisco products.

Procedure

Step 1 Open the URL https://id.cisco.com/signin/register to create a new account.

Step 2 Enter all the required fields to create an account.

Step 3 Click **Register**.

An email with an activation code is sent to verify your email address.

Note
If you haven't received an email yet, send an email to the registration support team at web-help@cisco.com.

Step 4 In the **Verify with your email** page, enter the activation code to complete the registration process and click **Verify**.

After successful registration, you will be redirected to the login page.

What to do next

Enter the newly created account details on the login page to request a Smart Account. Refer to Create a smart account and add licenses, on page 287.

Create a smart account and add licenses

Create a Smart Account to centrally manage and track your Cisco software licenses across your organization.

You should set up this account before you purchase licenses. Your account representative or reseller may have set up a Smart Account on your behalf. If so, obtain the necessary information to access the account from that person instead of using this procedure, then verify that you can access the account.

Before you begin

You must create a new Cisco account if you don't already have one. For instructions, refer to Create a Cisco Account.

For general information about Smart Accounts, refer to http://www.cisco.com/go/smartaccounts.

Procedure

Step 1 Go to the Create a Smart Account page.

You are prompted to log in with your Cisco account. In the **Create a Smart Account** page, your basic account information is displayed.

Step 2 Click the **My Account** icon appearing in the top right corner and click **Manage Profile**.



Step 3 Click **Personal** and configure these parameters:

a) In the **Your Company Details** section, click **Edit**.

b) In the **Company or organization** field, type your organization name.

c) If your company information is already present in our database, it will appear in the list. You can select your company.

In the **Address** drop-down list, select the address of your company.

d) If your company is not listed in our database, you can continue to enter your company information in the **Company or organization** field.

1. In the **Address** drop-down list, click the drop-down arrow and click **Add New Address**.
2. You can select one of the following **Address Type** options:

- **Company/Organization**: To provide your organization's address. Cisco verifies this address. If the address and company name cannot be validated with the country, you may not be able to proceed. Therefore, you must ensure that the correct address is provided.
- **Personal**: To provide your personal address.

e) Enter all the mandatory fields associated with your company and click **Update**.

The **Your Company Details** section displays the company details you entered.

A success message appears if your company details are validated.

Step 4 Click **Update**.

A success message appears if your company details are validated.

Step 5 Open the **Create a Smart Account** page, which was opened in the previous tab. Refresh the page if changes are not reflecting.

Alternatively, you can open this page using the https://software.cisco.com/software/company/smartaccounts/home?route=module/accountcreation URL and login using your credentials.

Step 6 Click **Create Account**.

The **Account Summary** page displays your account details.

Step 7 Click **Done**.

Step 8 Wait for an email telling you that your Smart Account is ready to set up. When it arrives, click the link it contains, as directed.

Step 9 Make sure your Smart Licensing account contains the available licenses you need.

What to do next

To configure Smart License using the Smart Software Manager, refer to Configure Smart Licensing, on page 288.

Configure Smart Licensing

This section describes how to use Smart Licensing using the Smart Software Manager or the Smart Software Manager On-Prem. To use Specific License Reservation, see Specific License Reservation, on page 299.

Register the Firewall Management Center for Smart Licensing

You can register the Firewall Management Center directly to the Smart Software Manager over the internet, or when using an air-gapped network, with the Smart Software Manager On-Prem.

Register the Firewall Management Center with the Smart Software Manager

Register the Firewall Management Center with the Smart Software Manager to enable Smart Licensing functionality and proper license management.

Before you begin

- Make sure your Smart Licensing account contains the available licenses you need.

  When you bought your device from Cisco or a reseller, your licenses should have been linked to your Smart Account. However, if you need to add licenses yourself, see Cisco Commerce Workspace.

- Ensure that the Firewall Management Center can reach the Smart Software Manager at smartreceiver.cisco.com.

- Configure NTP. During registration, a key exchange occurs between the Smart Agent and the Smart Software Manager, so time must be in sync for proper registration.

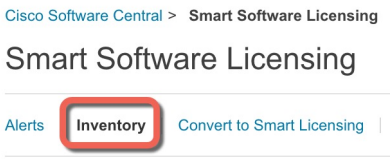
  For the Firepower 4100/9300, configure NTP on the chassis using the same NTP server as the Firewall Management Center.

- If your organization has multiple Firewall Management Centers, make sure each one has a unique name that clearly identifies and distinguishes it from others that may be registered to the same virtual account. This name is critical for managing your Smart License entitlements and ambiguous names will lead to problems later.

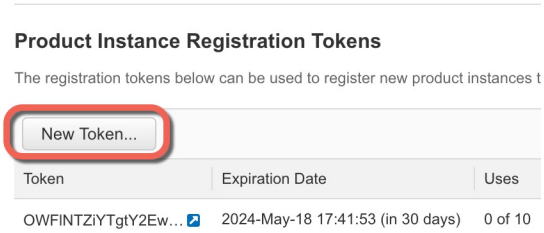
Procedure

Step 1 In the Smart Software Manager, request and copy a registration token for the virtual account to which you want to add this device.

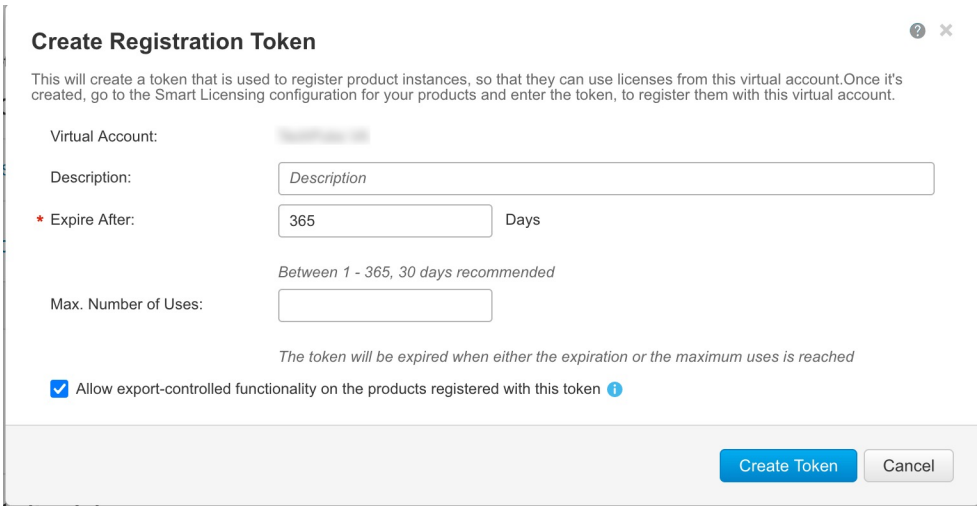
a) Click **Inventory**.



b) On the **General** tab, click **New Token**.



c) In the **Create Registration Token** dialog box enter these settings, and then click **Create Token**:



- **Description**
- **Expire After**—Cisco recommends 30 days.
- **Max. Number of Uses**
- **Allow export-controlled functionality on the products registered with this token**—Enables the export-compliance flag if you are in a country that allows for strong encryption. You must select this option now if you plan to use this functionality. If you enable this functionality later, you will need to re-register your device with a new product key and reload the device. If you do not see this option, your account does not support export-controlled functionality.

The token is added to your inventory.

d) Click the arrow icon to the right of the token to open the **Token** dialog box so you can copy the token ID to your clipboard. Keep this token ready for later in the procedure when you need to register the Firewall Threat Defense.

Figure 26: View token

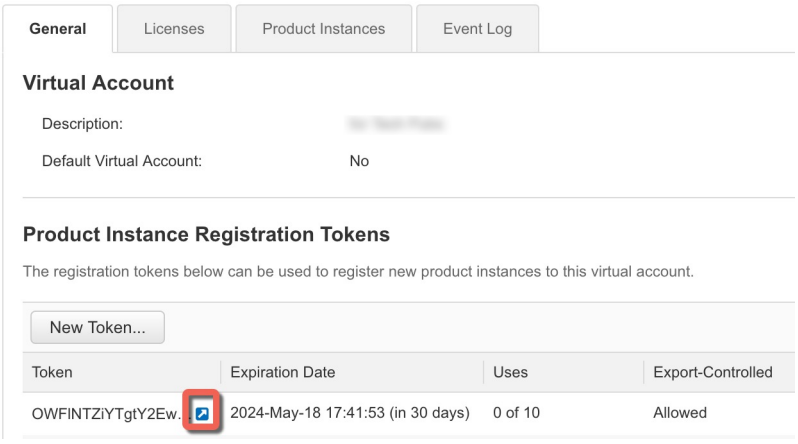
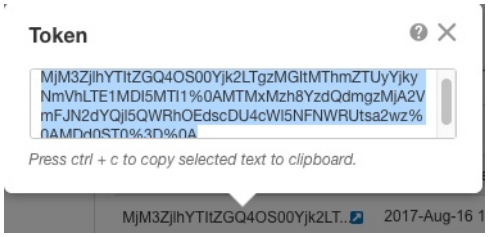


Figure 27: Copy token



Step 2 In the Firewall Management Center, choose **System** (⚙) > **Licenses** > **Smart Licenses**.

Step 3 Click **Register**.

Step 4 Paste the token you generated from Smart Software Manager into the **Product Instance Registration Token** field.

Make sure there are no empty spaces or blank lines at the beginning or end of the text.

Step 5 Decide whether to send usage data to Cisco.

- **Enable Cisco Success Network** is enabled by default. You can click **sample data** to see the kind of data Cisco collects. For more information, see Configure Firewall Management Center to share usage metrics and statistics with Cisco, on page 48.
- **Enable Cisco Support Diagnostics** is disabled by default. You can review the kind of data Cisco collects in the link provided above the check box. For more information, see Configure Firewall Management Center to share device health data with Cisco, on page 49.

Note

- When enabled, Cisco Support Diagnostics is enabled in the devices in the next sync cycle. The Firewall Management Center sync with the device runs once every 30 minutes.
- When enabled, Cisco Support Diagnostics is enabled automatically on any new device registered in this Firewall Management Center.

Step 6 Click **Apply Changes**.

What to do next

- Add a device to the Firewall Management Center; see *Add a Device to the Firewall Management Center* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.
- Assign licenses to your devices; see Assign licenses to multiple managed devices, on page 295.

Register the Firewall Management Center with the Smart Software Manager On-Prem

As described in Periodic communications with the Smart Software Manager, on page 272, the Firewall Management Center must communicate regularly with Cisco to maintain your license entitlement. If you have one of the following situations, you might want to use a Smart Software Manager On-Prem (formerly known as "Smart Software Satellite Server") as a proxy for connections to the Smart Software Manager:

- Your Firewall Management Center is offline or otherwise has limited or no connectivity (in other words, is deployed in an air-gapped network.)

  (For an alternate solution for air-gapped networks, see Licensing options for air-gapped deployments, on page 272.)

- Your Firewall Management Center has permanent connectivity, but you want to manage your Smart Licenses via a single connection from your network.

The Smart Software Manager On-Prem allows you to schedule synchronization or manually synchronize Smart License authorization with the Smart Software Manager.

For more information about the Smart Software Manager On-Prem, refer to https://www.cisco.com/c/en/us/buy/smart-accounts/software-manager.html#~on-prem

Procedure

Step 1 Deploy and set up Smart Software Manager On-Prem.

- See the documentation for the Smart Software Manager On-Prem, available from https://www.cisco.com/c/en/us/buy/smart-accounts/software-manager.html#~on-prem.
- Make a note of the CN of the TLS/SSL CERTIFICATE on your Smart Software Manager On-Prem.
- Go to http://www.cisco.com/security/pki/certs/clrca.cer and copy the entire body of the TLS/SSL CERTIFICATE (from "-----BEGIN CERTIFICATE-----" to "-----END CERTIFICATE-----") into a place you can access during configuration.

Step 2 Register the Firewall Management Center with the Smart Software Manager On-Prem.

a) Choose **Integration** > **Other Integrations** and click **Smart Software Manager On-Prem**.
b) Select **Connect to Cisco Smart Software Manager On-Prem Server**.
c) Enter the **URL** of your Smart Software Manager On-Prem, using the CN value you collected in the prerequisites of this procedure, in the following format:

   `https://FQDN_or_hostname_of_your_SSM_On-Prem/Transportgateway/services/DeviceRequestHandler`

   The FQDN or hostname must match the CN value of the CERTIFICATE presented by your Smart Software Manager On-Prem.

d) Add a new **SSL CERTIFICATE** and paste the CERTIFICATE text that you copied earlier, and click **Apply**.
e) Choose **System** (⚙) > **Licenses** > **Smart Licenses** and click **Register**.
f) Create a new token on Smart Software Manager On-Prem and copy the token.
g) Paste the token into the form in the Firewall Management Center page.
h) Click **Apply Changes**.

   The Firewall Management Center is now registered to Smart Software Manager On-Prem.

Step 3 After you assign licenses to devices, synchronize Smart Software Manager On-Prem to the Smart Software Manager.

Refer to the Smart Software Manager On-Prem documentation, above.

Step 4 Schedule ongoing synchronization times.

Enable the export control feature for accounts without global permission

If your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account.

Before you begin

- Make sure that your deployment does **not** already support the export-controlled functionality.

  If your deployment supports export-controlled features, you will see an option that allows you to enable export-controlled functionality in the **Create Registration Token** page in the Smart Software Manager. For more information, refer to https://www.cisco.com/c/en/us/buy/smart-accounts/software-manager.html.

- Make sure your deployment is not using an evaluation license.
- In the Smart Software Manager, in the **Inventory** > **Licenses** page, verify that you have the license that corresponds to your Firewall Management Center:

Export Control License	Firewall Management Center Model
Cisco Virtual FMC Series Strong Encryption (3DES/AES)	All Firewall Management Center Virtuals
Cisco FMC 1K Series Strong Encryption (3DES/AES)	1000, 1600
Cisco FMC 2K Series Strong Encryption (3DES/AES)	2500, 2600
Cisco FMC 4K Series Strong Encryption (3DES/AES)	4500, 4600

Procedure

Step 1 Choose **System** (⚙) > **Licenses** > **Smart Licenses**.

Note
If you see the **Request Export Key**, your account is approved for the export-controlled functionality and you can proceed to use the required feature.

Step 2 Click **Request Export Key** to generate an export key.

Tip
If the export control key request fails, make sure that your virtual account has a valid Export Control license.

Disable the export control license by clicking **Return Export Key**

You can now deploy configurations or policies that use the export-controlled features.

What to do next

Remember The new export-controlled licenses and all features enabled by it do not take effect on the Firewall Threat Defense devices until the devices are rebooted. Until then, only the features supported by the older license will be active.

In High Availability deployments both the Firewall Threat Defense devices need to be rebooted simultaneously, to avoid an Active-Active condition.

Assign licenses to devices

You can assign most licenses when you register a device to the Firewall Management Center. You can also assign licenses per device, or for multiple devices.

Assign licenses to a single device

This task allows you to enable or disable specific licensed features on a managed device by assigning or removing licenses as needed.

You can enable or disable specific licensed features on a managed device by assigning or removing licenses as needed. Although there are some exceptions, you cannot use the features associated with a license if you disable it on a managed device.

Note For container instances on the same security module/engine, you apply the license to each instance; note that the security module/engine consumes only one license per feature for all instances on the security module/engine.

Note For the Firewall Threat Defense cluster, you apply the licenses to the cluster as a whole; note that each unit in the cluster consumes a separate license per feature.

Before you begin

You must have Admin or Network Admin privileges to perform this task. When operating with multiple domains, you must do this task in leaf domains.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to assign or disable a license, click **Edit** ( ).

Step 3 Click **Device**.

Step 4 Next to the **License** section, click **Edit** ( ).

Step 5 Check or clear the appropriate check boxes to assign or disable licenses for the device.

Step 6 Click **Save**.

Step 7 Deploy configuration changes; see the *Cisco Secure Firewall Management Center Device Configuration Guide*.

What to do next

Verify license status:

1. Choose **System** ( ) > **Licenses** > **Smart Licenses**.
2. Enter the hostname or IP address of the device into the filter at the top of the Smart Licenses table
3. Verify that only a green circle with a **Check Mark** ( ) appears for each device, for each license type.

   If you see any other icon, hover over the icon for more information.

Assign licenses to multiple managed devices

Devices managed by the Firewall Management Center obtain their licenses via the Firewall Management Center, not directly from the Smart Software Manager.

Note For container instances on the same security module/engine, you apply the license to each instance; note that the security module/engine consumes only one license per feature for all instances on the security module/engine.

Note For the Firewall Threat Defense cluster, you apply the licenses to the cluster as a whole; note that each unit in the cluster consumes a separate license per feature.

Use this procedure to enable licensing on multiple devices at once.

Procedure

Step 1 Choose **System** ( ) > **Licenses** > **Smart Licenses** or **Specific Licenses**.

Step 2 Click **Edit Licenses**.

Step 3 For each type of license you want to add to a device:

a) Click the tab for that type of license.
b) Click a device in the list on the left.
c) Click **Add** to move that device to the list on the right.
d) Repeat for each device to receive that type of license.

For now, don't worry about whether you have licenses for all of the devices you want to add.

e) Repeat this subprocedure for each type of license you want to add.

f) To remove a license, click the **Delete** ( ) next to the device.

g) Click **Apply**.

What to do next

Verify that your licenses are correctly installed. Follow the procedure in Monitor Smart Licensing, on page 297.

Manage Smart Licensing

This section describes how to manage Smart Licensing.

Deregister the Firewall Management Center

Deregister your Firewall Management Center from the Smart Software Manager to release all of the license entitlements back to your Smart Account so they can be used for other devices.

Deregister if you need to decommission the Firewall Management Center or reimage it.

Refer to Unregistered state, on page 274 for more information about license enforcement in an unregistered state.

Procedure

Step 1 Choose **System** ( ) > **Licenses** > **Smart Licenses**.

Step 2 Click **Deregister** ( ).

Synchronize or reauthorize the Firewall Management Center

By default, the ID certificate is automatically renewed every 6 months, and the license entitlement is renewed every 30 days. You might want to manually renew the registration for either of these items if you have a limited window for internet access, or if you make any licensing changes in the Smart Software Manager, for example.

Procedure

Step 1 Choose **System** ( ) > **Licenses** > **Smart Licenses**.

Step 2 To renew the ID certificate, click **Synchronize** ( ).

Step 3 To renew the license entitlements, click **Re-Authorize**.

Monitor Smart License status

The **Smart License Status** section of the **System** (⚙) > **Licenses** > **Smart Licenses** page provides an overview of license usage on the Firewall Management Center.

Usage authorization

Possible status values are:

- **In-compliance** (🟢) — All licenses assigned to managed devices are in compliance and the Firewall Management Center is communicating successfully with the Smart Software Manager.
- **License is in compliance but communication with licensing authority has failed**— Device licenses are in compliance, but the Firewall Management Center is not able to communicate with the Cisco licensing authority.
- **Out-of-compliance icon or unable to communicate with License Authority**— One or more managed devices is using a license that is out of compliance, or the Firewall Management Center has not communicated with the Smart Software Manager in more than 90 days.

Product registration

Specifies the last date when the Firewall Management Center contacted the Smart Software Manager and registered.

Assigned virtual account

Specifies the Virtual Account under the Smart Account that you used to generate the Product Instance Registration Token and register the Firewall Management Center. If this deployment is not associated with a particular virtual account within your Smart Account, this information is not displayed.

Export-Controlled features

If this option is enabled, you can deploy restricted features. For details, see License for export-controlled functionality, on page 280.

Cisco success network

Specifies whether you have enabled Cisco Success Network for the Firewall Management Center. If this option is enabled, you provide usage information and statistics to Cisco which are essential to provide you with technical support. This information also allows Cisco to improve the product and make you aware of unused available features so that you can maximize the value of the product in your network. See Configure Firewall Management Center to share usage metrics and statistics with Cisco, on page 48 for more information.

Monitor Smart Licensing

To view the license status for the Firewall Management Center and its managed devices, use the Smart Licenses page.

For each type of license in your deployment, the page lists the total number of licenses consumed, whether the license is in compliance or out of compliance, the device type, and the domain and group where the device is deployed. You can also view the Firewall Management Center's Smart License Status. Container instances on the same security module/engine only consume one license per security module/engine. Therefore, even though the Firewall Management Center lists each container instance separately under each license type, the number of licenses consumed for feature license types will only be one.

Other than the **Smart Licenses** page, there are a few other ways you can view licenses:

- The **Product Licensing** dashboard widget provides an at-a-glance overview of your licenses.

  See Adding widgets to a dashboard, on page 372 and Dashboard widget availability by user role, on page 359 and The product licensing widget, on page 369.

- The **Device Management** page (**Devices** > **Device Management**) lists the licenses applied to each of your managed devices.
- The **Smart License Monitor** health module communicates license status when used in a health policy.

Procedure

Step 1 Choose **System** ( ) > **Licenses** > **Smart Licenses**.

Step 2 In the **Smart Licenses** table, click the arrow at the left side of each **License Type** folder to expand that folder.

Step 3 In each folder, verify that each device has a green circle with a **Check Mark** ( ) in the **License Status** column.

Note
If you see duplicate Firewall Management Center Virtual licenses, each represents one managed device.

If all devices show a green circle with a **Check Mark** ( ), your devices are properly licensed and ready to use.

If you see any License Status other than a green circle with a **Check Mark** ( ), hover over the status icon to view the message.

What to do next

- If you had any devices that did not have a green circle with a **Check Mark** ( ), you may need to purchase more licenses.

Troubleshoot Smart Licensing

This topic lists some common Smart Licensing issues and provides troubleshooting steps.

Expected licenses do not appear in my smart account

If the licenses you expect to see are not in your Smart Account, try the following:

- Make sure they are not in a different Virtual Account. Your organization's license administrator may need to assist you with this.
- Check with the person who sold you the licenses to be sure that transfer to your account is complete.

Unable to connect to Smart Licensing server

Check the obvious causes first. For example, make sure your Firewall Management Center has outside connectivity. See Communication Ports and Internet Access, on page 1063.

Unexpected out-of-compliance notification or other error

- If a device is already registered to a different Firewall Management Center, you need to deregister the original Firewall Management Center before you can license the device under a new Firewall Management Center. See Deregister the Firewall Management Center, on page 296.
- Check if the term of the subscription license has expired.

Troubleshoot other issues

For solutions to other common issues, see Use FMC and FTD Smart License Registration and Common Issues to Troubleshoot.

Specific License Reservation

If you cannot have or do not want direct communications between the Firewall Management Center and the Smart Software Manager or the Smart Software Manager On-Prem, use Specific License Reservation. Specific License Reservation uses the same licenses as regular Smart Licensing.

Configuration process

To configure Specific License Reservation, you will generate a request code on the Firewall Management Center, access the Smart Software Manager with your browser, enter the request code to obtain an authorization code, then return the authorization code to the Firewall Management Center.

Note Various names are used at Cisco for Specific License Reservation. These include SLR, SPLR, PLR, and Permanent License Reservation. These terms may also be used at Cisco to refer to similar but not necessarily identical licensing models.

Prerequisites for Specific License Reservation

This topic lists the prerequisites for configuring Specific License Reservation.

Switching to Specific License Reservation

If you are currently using regular Smart Licensing, de-register the Firewall Management Center from Smart Software Manager before you implement Specific License Reservation; refer to Deregister the Firewall Management Center, on page 296.

All Smart Licenses that are currently deployed to the Firewall Management Center will be returned to the pool of available licenses in your account, and you can re-use them when you implement Specific License Reservation.

Specific License Reservation and high availability

Configure high availability before you assign licenses. If you already assigned licenses to devices on the secondary Firewall Management Center, be sure to unassign them.

If an SLR license is assigned to a primary Firewall Management Center, when the secondary Firewall Management Center becomes active after a failover, you cannot add the SLR license to the secondary Firewall Management Center.

- Perform a failover to make the primary Firewall Management Center active.
- Unassign and re-assign the license to the secondary Firewall Management Center.

Verify that your smart account is ready to deploy Specific License Reservation

To prevent problems when deploying your Specific License Reservation, complete this procedure before you make any changes in your Firewall Management Center.

Before you begin

- Ensure that you have met the requirements described in Prerequisites for Specific License Reservation, on page 299.
- Make sure you have your Smart Software Manager credentials.

Procedure

Step 1 Sign in to the Smart Software Manager:

https://software.cisco.com/#SmartLicensing-Inventory

Step 2 If applicable, select the correct account from the top right corner of the page.

Step 3 If necessary, click **Inventory**.

Step 4 Click **Licenses**.

Step 5 Verify the following:

- There is a **License Reservation** button.
- There are enough platform and feature licenses for the devices and features you will deploy, including Firewall Management Center Virtual entitlements for your devices, if applicable.

Step 6 If any of these items is missing or incorrect, contact your account representative to resolve the problem.

Note
Do not continue with this process until any problems are corrected.

Enable the Specific Licensing menu option

This procedure changes the Smart Licenses menu option to Specific Licenses in the Firewall Management Center.

Procedure

Step 1 Access the Firewall Management Center console using a USB keyboard and VGA monitor, or use SSH to access the management interface.

Step 2 Log into the Firewall Management Center CLI **admin** account.

Step 3 Enter the **expert** command to access the Linux shell.

Step 4 Execute the following command to access the Specific License Reservation options:

sudo manage_slr.pl

Example:

```
admin@fmc63betaslr: ~$ sudo manage_slr.pl
Password:


*************** Configuration Utility **************

1   Show SLR Status
2   Enable SLR
3   Disable SLR
0   Exit

************************************************************
Enter choice:
```

Step 5 Enable Specific License Reservation by selecting option **2**.

Step 6 Select option **0** to exit the manage_slr utility.

Step 7 Type **exit** to exit the Linux shell.

Step 8 Enter **exit** to exit the command line interface.

Step 9 Verify that you can access the **Specific License Reservation** page in the Firewall Management Center web interface:

- If the **System** (⚙) > **Licenses** > **Smart Licenses** page is currently displayed, refresh the page.
- Otherwise, choose **System** (⚙) > **Licenses** > **Specific Licenses**.

Enter the Specific License Reservation authorization code into the Firewall Management Center

Use specific license reservation when you need to reserve and use licenses offline or in environments where direct communication with Smart Software Manager is not available.

Procedure

Step 1 Generate the reservation request code.

a. In the Firewall Management Center, choose **System** (⚙) > **Licenses** > **Specific Licenses**.

b. Click **Generate**.

c. Make a note of the **Reservation Request Code**.

Step 2 Generate the reservation authorization code.

a. Go to the Cisco Smart Software Manager: https://software.cisco.com/#SmartLicensing-Inventory

b. If necessary, select the correct account from the top right of the page.

c. If necessary, click **Inventory**.

d. Click **Licenses** and click **License Reservation**.

e. Enter the code that you generated from Firewall Management Center into the **Reservation Request Code** box and click **Next**.

f. Select **Reserve a specific license** and scroll down to display the entire License grid.

g. Under **Quantity To Reserve**, enter the number of each platform and feature license needed for your deployment.

Note

- You must explicitly include a Base license for each managed device, or, for multi-instance deployments, for each container.
- If you are using the Firewall Management Center Virtual, you must include a platform entitlement for each container (in multi-instance deployments) or each managed device (all other deployments).
- If you use strong encryption functionality:
  - If your entire Smart Account is enabled for export-controlled functionality, you do not need to do anything here.
  - If your organization's entitlement is per-Firewall Management Center, you must select the appropriate license.

    For the correct license name to choose for your Firewall Management Center, see the prerequisites in Enable the export control feature for accounts without global permission, on page 293.

Note

-
- You must explicitly include a Base license for each managed device, or, for multi-instance deployments, for each container.
- If you are using the Firewall Management Center Virtual, you must include a platform entitlement for each container (in multi-instance deployments) or each managed device (all other deployments).

- If you use strong encryption functionality:
  - If your entire Smart Account is enabled for export-controlled functionality, you do not need to do anything here.
  - If your organization's entitlement is per-Firewall Management Center, you must select the appropriate license.

    For the correct license name to choose for your Firewall Management Center, see the prerequisites in Enable the export control feature for accounts without global permission, on page 293.

h. Click **Next** and click **Generate Authorization Code**.

At this point, the license is now in use according to the Smart Software Manager.

i. Download the Authorization Code in preparation for entering it into the Firewall Management Center.

Step 3 Enter the authorization code in the Firewall Management Center.

a. In the Firewall Management Center, click **Browse** to upload the text file with the authorization code that you generated from the Smart Software Manager.

b. Click **Install**

.

c. Verify that the **Specific License Reservation** page shows the **Usage Authorization** status as authorized.

Step 4 Click the **Reserved License** tab to verify the licenses selected while generating the **Authorization Code**.

If you do not see the licenses you require, then add the necessary licenses. For more info, refer to Update the Specific Licenses for Firepower Management Center.

Assign specific licenses to managed devices

Use this procedure to quickly assign licenses to multiple managed devices at one time.

You can also use this procedure to disable or move licenses from one device to another. If you disable a license for a device, you cannot use the features associated with that license on that device.

Procedure

Step 1 Choose **System** (⚙) > **Licenses** > **Specific Licenses**.

Step 2 Click **Edit Licenses**.

Step 3 Click each tab and assign licenses to devices as needed.

Step 4 Click **Apply**.

Step 5 Click the **Assigned Licenses** tab and verify that your licenses are correctly installed on each device.

Step 6 Deploy configuration changes; see the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Manage specific license reservation

This section describes how to manage Specific License Reservation.

Maintain your Specific License Reservation deployment

To update the threat data and software that keep your deployment effective, see Maintain your air-gapped deployment, on page 267.

To ensure that all functionality continues to work without interruption, monitor your license expiration dates (in the **Reserved Licenses** tab). If any of the licenses expire, the Firewall Management Center will be in the Out of Compliance state if the usage count is greater than the available count.

Update a Specific License Reservation

After you have successfully deployed Specific Licenses on your Firewall Management Center, you can add or remove entitlements at any time using this procedure.

Use this procedure if you need to renew your licenses after they expire. If you do not have the required licenses, these actions are restricted:

- Device registration
- Policy deployment

Procedure

Step 1 In the Firewall Management Center, obtain the unique product instance identifier of this Firewall Management Center:

a) Select **System** (⚙) > **Licenses** > **Specific Licenses**.
b) Make a note of the **Product Instance** value.

You will need this value several times during this process.

Step 2 In the Smart Software Manager, identify the Firewall Management Center to update:

a) Go to the Smart Software Manager:

https://software.cisco.com/#SmartLicensing-Inventory

b) If necessary, click **Inventory**.
c) Click **Product Instances**.
d) Look for a product instance that has **FP** in the **Type** column and a generic SKU (not a hostname) in the **Name** column. You may also be able to use the values in other table columns to help determine which Firewall Management Center is the correct Firewall Management Center. Click the name.
e) Look at the **UUID** and see if it is the UUID of the Firewall Management Center that you are trying to modify.

If not, you must repeat these steps until you find the correct Firewall Management Center.

Step 3 When you have located the correct Firewall Management Center in the Smart Software Manager, update the reserved licenses and generate a new authorization code:

a) On the page that shows the correct UUID, from the **Actions** drop-down list, choose **Update Reserved Licenses**.

b) Update the reserved licenses as needed.

Note

- You must explicitly include a Base license for each managed device, or, for multi-instance deployments, for each container.
- If you are using the Firewall Management Center Virtual, you must include a platform entitlement for each container (in multi-instance deployments) or each managed device (all other deployments).
- If you use strong encryption functionality:
  - If your entire Smart Account is enabled for export-controlled functionality, you do not need to do anything here.
  - If your organization's entitlement is per-Firewall Management Center, you must select the appropriate license.

    For the correct license name to choose for your Firewall Management Center, see the prerequisites in Enable the export control feature for accounts without global permission, on page 293.

c) Click **Next** and verify the details.

d) Click **Generate Authorization Code**.

e) Download the Authorization Code in preparation for entering it into the Firewall Management Center.

f) Leave the **Update Reservation** page open. You will return to it later in this procedure.

Step 4 Update the Specific Licenses in the Firewall Management Center.

a) Choose **System** (⚙) > **Licenses** > **Specific Licenses**.

b) Click **Edit SLR**.

c) Click **Browse** to upload the newly generated authorization code.

d) Click **Install** to update the licenses.

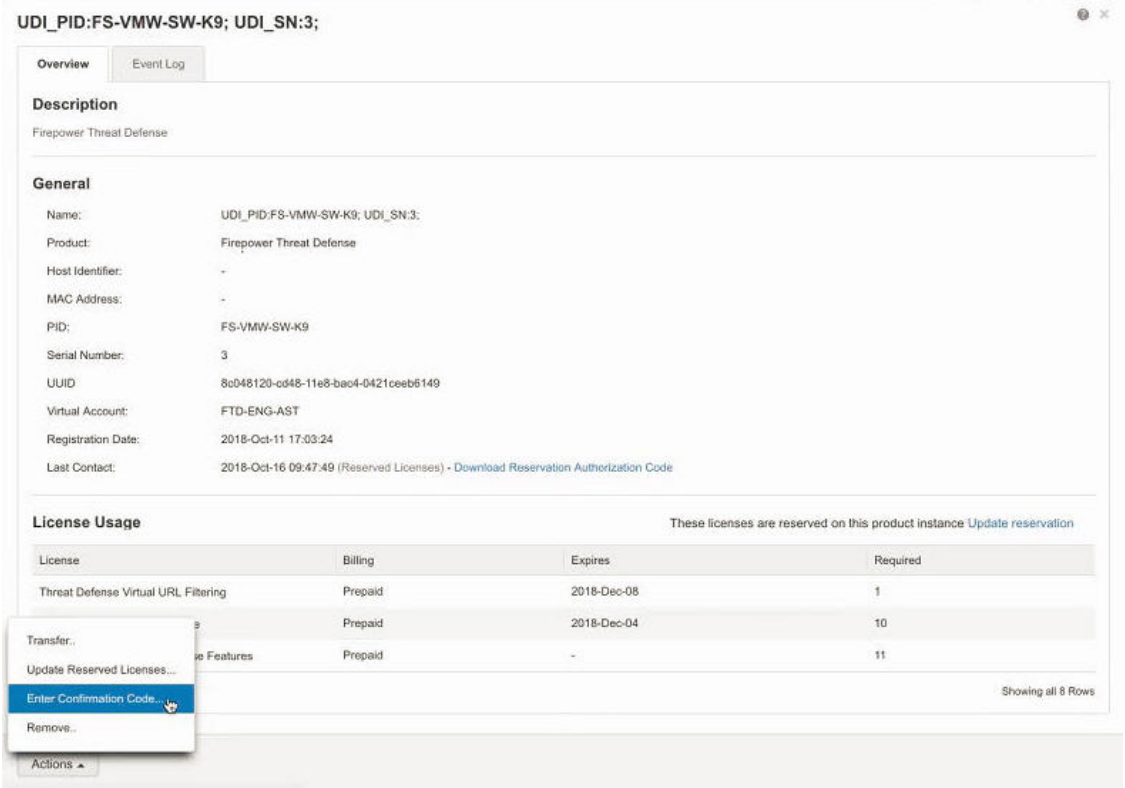
After successful installation of the authorization code, ensure that the licenses shown in the **Reserved** column of Firewall Management Center, matches with the licenses that you have reserved in the Smart Software Manager.

e) Make a note of the **Confirmation Code**.

Step 5 Enter the confirmation code in the Smart Software Manager:

a) Return to the Smart Software Manager page that you left open earlier in this procedure.

b) From the **Actions** drop-down list, choose **Enter Confirmation Code**:



c) Enter the confirmation code that you generated from the Firewall Management Center.

Step 6 In the Firewall Management Center, verify that your licenses are reserved as you expect them, and that each feature for each managed device shows a green circle with a **Check Mark** ( ).

If necessary, refer to Monitor Specific License Reservation status, on page 308 for more information.

Step 7 Deploy configuration changes; see the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Deactivate and return the Specific License Reservation

If you no longer need a specific license, you must return it to your Smart Account. If you want to register your Smart Licensing account, you must disable the Specific License Reservation (Step 6 of the procedure below).

This procedure releases all license entitlements associated with the Firewall Management Center back to your virtual account. After you de-register, no updates or changes on licensed features are allowed.

Important If you do not follow all of the steps in this procedure, the license remains in an in-use state and cannot be re-used.

Procedure

Step 1 In the Firewall Management Center Web interface, select **System** (⚙) > **Licenses** > **Specific Licenses**.

Step 2 Make a note of the **Product Instance** identifier for this Firewall Management Center.

Step 3 Generate a return code from the Firewall Management Center.

a) Click **Return SLR**.

Devices become unlicensed and the Firewall Management Center moves to the de-registered state. It generates a return code and allows you to re-register the Firewall Management Center with SLR.

b) Make a note of the **Return Code**.

Step 4 In the Smart Software Manager, identify the Firewall Management Center to deregister:

a) Go to the Smart Software Manager:

https://software.cisco.com/#SmartLicensing-Inventory

b) If necessary, click **Inventory**.

c) Click **Product Instances**.

d) Look for a product instance that has **FP** in the **Type** column and a generic SKU (not a hostname) in the **Name** column. You may also be able to use the values in other table columns to help determine which Firewall Management Center is the correct Firewall Management Center. Click the name.

e) Look at the **UUID** and see if it is the UUID of the Firewall Management Center that you are trying to modify.

If not, you must repeat these steps until you find the correct Firewall Management Center.

Step 5 When you have identified the correct Firewall Management Center, return the licenses to your Smart Account:

a) On the page that shows the correct UUID, from the **Actions** drop-down list, choose **Remove**.

b) Enter the reservation return code that you generated from the Firewall Management Center into the **Remove Product Instance** dialog box.

c) Click **Remove Product Instance**.

The specific reserved licenses are returned to the available pool in your Smart Account and this Firewall Management Center is removed from the Smart Software Manager Product Instances list.

Step 6 Disable the Specific License in the Firewall Management Center Linux shell:

a) Access the Firewall Management Center console using a USB keyboard and VGA monitor, or use SSH to access the management interface.

b) Log in to the Firewall Management Center CLI **admin** account. This gives you access to the command line interface.

c) Enter the **expert** command to access the Linux shell.

d) Execute this command:

sudo manage_slr.pl

Example:

```
admin@fmc63betaslr: ~$ sudo manage_slr.pl
Password:
```

```
*************** Configuration Utility **************

1   Show SLR Status
2   Enable SLR
3   Disable SLR
0   Exit

************************************************************
Enter choice:
```

e) Select menu option **3** to disable the Specific License Reservation.
f) Select option **0** to exit the manage_slr utility.
g) Enter **exit** to exit the Linux shell.
h) Enter **exit** to exit the command line interface.

Monitor Specific License Reservation status

The **System** (⚙) > **Licenses** > **Specific Licenses** page provides an overview of license usage on the Firewall Management Center, as described below.

Usage authorization

Possible status values are:

- **Authorized** — The Firewall Management Center is in compliance and registered successfully with the License Authority, which has authorized the license entitlements for the appliance.
- **Out-of-compliance** — If licenses are expired or if the Firewall Management Center has overused licenses even though they are not reserved, status shows as Out-of-Compliance. License entitlements are enforced in Specific License Reservation, so you must take action.

Product registration

Specifies registration status and the date that an authorization code was last installed or renewed on the Firewall Management Center.

Export-Controlled features

Specifies whether you have enabled export-controlled functionality for the Firewall Management Center.

For more information about Export-Controlled Features, see License for export-controlled functionality, on page 280.

Product instance

The Universally Unique Identifier (UUID) of this Firewall Management Center. This value identifies this device in the Smart Software Manager.

Confirmation code

The **Confirmation Code** is needed if you update or deactivate and return Specific Licenses.

Assigned licenses tab

Shows the licenses assigned to each device and the status of each.

Reserved licenses tab

Shows the number of licenses used and available to be assigned, and license expiration dates.

Troubleshoot Specific License Reservation

This topic provides troubleshooting guidance for resolving Specific License Reservation issues.

How do I identify a particular Firewall Management Center in the Product Instance list in Smart Software Manager?

In the Product Instances page in Smart Software Manager, if you cannot identify the product instance based on a value in one of the columns in the table, you must click the name of each generic product instance of type **FP** to view the product instance details page. The **UUID** value on this page uniquely identifies one management center.

In the Firewall Management Center web interface, the UUID for the management center is the **Product Instance** value displayed in the **System** (⚙) > **Licenses** > **Specific Licenses** page.

I do not see a license reservation button in the Smart Software Manager

If you do not see the **License Reservation** button, then your account is not authorized for Specific License Reservation. If you have already enabled Specific License Reservation in the Linux shell and generated a request code, perform these steps:

1. If you have already generated a **Request Code** in the management center web interface, cancel the request code.
2. Disable Specific License Reservation in the management center Linux shell as described within the section Deactivate and return the Specific License Reservation, on page 306.
3. Register the management center with the Smart Software Manager in regular mode using smart token.
4. Contact Cisco TAC to enable Specific License for your smart account.

I was interrupted in the middle of the licensing process. How can i pick up where I left off?

If you have generated but not yet downloaded an Authorization code from the Smart Software Manager, you can go to the **Product Instance** page in the Smart Software Manager, click the product instance, then click **Download Reservation Authorization Code**.

I am unable to register devices to the Firewall Management Center Virtual

Make sure you have enough Firewall Management Center Virtual entitlements in your Smart Account to cover the devices you want to register, then update your deployment to add the necessary entitlements.

Refer to Update a Specific License Reservation, on page 304.

I have enabled Specific Licensing, but now I do not see a Smart License page.

This is the expected behavior. When you enable Specific Licensing, Smart Licensing is disabled. You can use the Specific License page to perform licensing operations.

If you want to use Smart Licensing, you must return the Specific License. For more information refer to, Deactivate and return the Specific License Reservation, on page 306.

What if I do not see a Specific License page in the Firewall Management Center Virtual?

You need to enable Specific License to view the Specific License page. For more information see, Enable the Specific Licensing menu option, on page 301.

I have disabled specific licensing, but forgot to copy the return code. what should I do?

The **Return Code** is saved in the Firewall Management Center Virtual. You must re-enable the Specific License from the Linux shell (see Enable the Specific Licensing menu option, on page 301), then refresh the Firewall Management Center Virtual web interface. Your **Return Code** will be displayed.

Configure Legacy Firewall Management Center PAK-Based Licenses

The Firewall Management Center supports either a Smart license or a legacy PAK (Product Activation Key) license for its platform license. This procedure describes how to apply a PAK-based license.

After re-registration of your Smart Account, you must manually add the classic licenses for all classic devices.

Before you begin

- Make sure you have the product activation key (PAK) from the Software Claim Certificate that Cisco provided when you purchased the license. If you have a legacy, pre-Cisco license, contact Support.

Procedure

Step 1 Note the license key for the Firewall Management Center.

The license key uniquely identifies the Firewall Management Center in the Smart Software Manager. It is composed of a product code (for example, 66) and the MAC address of the management port (eth0) of the Firewall Management Center; for example, 66:00:00:77:FF:CC:88.

a) Choose **System** (⚙) > **Licenses** > **Classic Licenses**.
b) Click **Add New license**.
c) Note the value in the **license Key** field at the top of the **Add Feature license** dialog.

Step 2 Choose **System** (⚙) > **Licenses** > **Classic Licenses**.

Step 3 Click **Add New license**.

Step 4 Continue as appropriate:

- If you have already obtained the license text, skip to Step 8.
- If you still need to obtain the license text, go to the next step.

Step 5 Click **Get license** to open the License Registration Portal.

Note
If you cannot access the Internet using your current computer, switch to a computer that can, and browse to http://cisco.com/go/license.

Step 6 Generate a license from the PAK in the License Registration Portal: https://cisco.com/go/license.

This step requires the PAK you received during the purchase process, as well as the LICENSE key for the Firewall Management Center.

For more information on using this portal, see:

https://slexui.cloudapps.cisco.com/SWIFT/LicensingUI/Quickstart

You will need your account credentials in order to access these links.

Step 7 Copy the license text from either the License Registration Portal display, or the email the License Registration Portal sends you.

Important
The licensing text block in the portal or email message may include more than one license. Each license is bounded by a BEGIN LICENSE line and an END LICENSE line. Make sure that you copy and paste only one license at a time.

Step 8 Return to the **Add Feature License** page in the management center virtual's web interface.

Step 9 Paste the license text into the **License** field, click **Verify License**, and click **Submit License**.

If the license is invalid, make sure that you correctly copied the license text.

Additional information about licensing

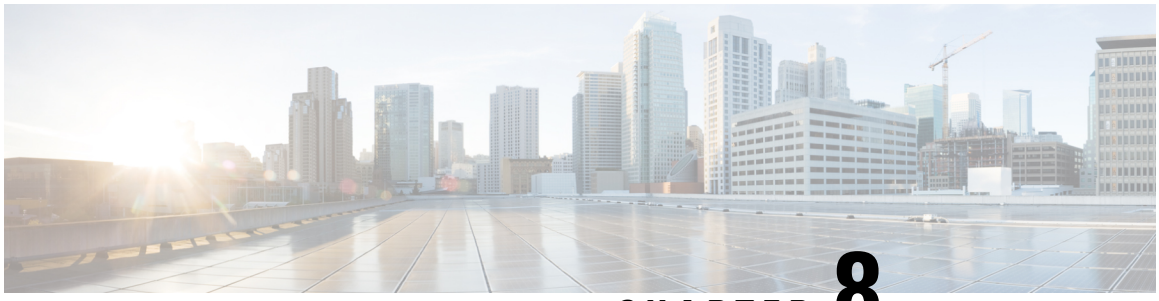
For additional information to help resolve common licensing questions, see these documents:

- Licensing FAQ
- License Roadmap

History for licenses

This topic provides a chronological history of license-related features, enhancements, and changes implemented across different software versions.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Performance tier licensing for the Firewall Threat Defense Virtual	7.0	Any	Performance-tiered licensing provides different throughput levels and VPN connection limits based on deployment requirements. License tiers map to new Firewall Threat Defense Virtual models.
Licensing for multi-instance capability for the Firewall Threat Defense on the Firepower 4100/9300	6.3	Any	You can now deploy multiple Firewall Threat Defense container instances on a Firepower 4100/9300. You only need a single license per feature per security module/engine. The base license is automatically assigned to each instance. New/Modified screens: **System** > **Licenses** > **Smart Licenses**. Supported platforms: Firewall Threat Defense on the Firepower 4100/9300
Specific License Reservation for air-gapped deployments	6.3	Any	Customers whose deployments cannot connect to the internet to communicate with the Cisco License Authority can use a Specific License Reservation. New/Modified screens: **System** > **Licenses** > **Specific Licenses** (This option is not available by default.) Supported platforms: Firewall Management Center, Firewall Threat Defense
Export-controlled functionality for restricted customers	6.3	Any	Certain customers whose Smart Accounts are not otherwise eligible to use restricted functionality can purchase term-based licenses, with approval. Supported platforms: Firewall Management Center, Firewall Threat Defense



CHAPTER 8

High availability

The following topics describe how to configure Active/Standby high availability of Cisco Secure Firewall Management Centers:

- High availability in Firewall Management Center, on page 313
- Requirements for Firewall Management Center high availability, on page 321
- Prerequisites for Firewall Management Center high availability, on page 325
- Establish high availability for Firewall Management Center, on page 325
- View Firewall Management Center high availability status, on page 331
- Configure synchronization of configurations on Firewall Management Center high availability pairs, on page 332
- Recommendation: Configure external access to the Firewall Management Center database only on the active peer, on page 332
- Resolve device registration using CLI in Firewall Management Center high availability, on page 333
- Switch peer roles in the Firewall Management Center high availability pair, on page 334
- Pause communication between paired Firewall Management Centers, on page 334
- Restart communication between paired Firewall Management Centers, on page 335
- Change the IP address of the Firewall Management Center in a high availability pair, on page 335
- Disable Firewall Management Center high availability, on page 336
- Replace Firewall Management Centers in a high availability pair, on page 337
- Restore management center in a high-availability pair (no hardware failure), on page 341
- History for Firewall Management Center high availability, on page 344

High availability in Firewall Management Center

A high availability pair in Firewall Management Center is a redundancy feature that

- designates redundant Firewall Management Centers to manage devices
- supports Active/Standby roles. One appliance is active and manages devices, while the standby unit does not actively manage devices, and
- enables event data streaming from managed devices to both units, enabling uninterrupted monitoring if one fails.

High availability reference information

High availability in Firewall Management Center allows configuration of a secondary unit to take over the functionality of a primary unit if the primary fails. When the primary fails, you must promote the secondary to become the active unit.

Event data streams from managed devices to both units in the high availability pair. If one unit fails, you can monitor your network without interruption using the other unit.

Firewall Management Centers configured as a high availability pair do not need to be on the same trusted management network, nor do they have to be in the same geographic location.



Caution Because the system restricts some functionality to the active Firewall Management Center, if that appliance fails, you must promote the standby Firewall Management Center to active.



Note Triggering a switchover on Firewall Management Center immediately after a successful change deployment can lead to preview configuration not working on the new active Firewall Management Center. This does not impact policy deploy functionality. It is recommended to trigger a switchover on the Firewall Management Center after the necessary sync is completed.

Similarly, when Firewall Management Center HA synchronization is in degraded state, triggering a switchover or changing roles could cause Firewall Management Center HA to damage the database, which may result in catastrophic failure. We recommend that you immediately contact Cisco Technical Assistance Center (TAC) for further assistance to resolve this issue.

This HA synchronization can end up in degraded state due to various reasons. The Replace Firewall Management Centers in a high availability pair, on page 337 section in this chapter covers some of the failure scenarios and the subsequent procedure to fix the issue. If the reason or scenario of degraded state matches to the scenarios explained, follow the steps to fix the issue. For other reasons, we recommend that you contact TAC.

Remote access VPN high availability

If the primary device has Remote Access VPN configuration with an identity certificate enrolled using a CertEnrollment object, the secondary device must have an identity certificate enrolled using the same CertEnrollment object. The CertEnrollment object can have different values for the primary and secondary devices due to device-specific overrides. The limitation is only that the same CertEnrollment object is enrolled on the two devices before high availability formation.

SNMP behavior in Firewall Management Center high availability

In an SNMP-configured HA pair, when you deploy an alert policy, the active Firewall Management Center sends the SNMP traps. When the primary Firewall Management Center fails, the secondary Firewall Management Center which becomes the active unit starts sending the SNMP traps without the need for any additional configuration.

High availability example

For example, if the primary Firewall Management Center fails, you can promote the secondary unit to active and continue managing devices and monitoring events without interruption.

Roles and statuses in Firewall Management Center high availability

Primary and secondary roles

When setting up Firewall Management Centers in a high availability pair, you configure one Firewall Management Center to be primary and the other as secondary. During configuration, the primary unit's policies are synchronized to the secondary unit. After this synchronization, the primary Firewall Management Center becomes the active peer, while the secondary Firewall Management Center becomes the standby peer. The two units act as a single appliance for managed device and policy configuration.

Active and standby statuses

The main difference between the two Firewall Management Centers in a high availability pair is whether the peer is active or standby. The active Firewall Management Center remains fully functional, where you can manage devices and policies. Configuration options are hidden on the standby Firewall Management Center, and you cannot make any changes.

Event processing on Firewall Management Center high availability pairs

An event processing configuration is a high-availability feature that

- allows both Firewall Management Centers in a high-availability pair to receive events from managed devices
- ensures that the appliances do not share management IP addresses, and
- provides continuous processing of events without intervention if one Firewall Management Center fails.

AMP cloud connections and malware information

AMP cloud connections and malware information are security features that

- operate independently for Firewall Management Centers in a high availability pair
- do not share malware dispositions between paired Firewall Management Centers, and
- require both primary and secondary Firewall Management Centers to have access to the AMP cloud to ensure consistent malware dispositions.

AMP cloud connection requirements

Both primary and secondary Firewall Management Centers must have access to the AMP cloud to ensure continuity of operations and consistent malware dispositions.

URL filtering and security intelligence

Synchronization of URL filtering information

URL filtering and Security Intelligence configurations and information are synchronized between Secure Firewall Management Centers in a high availability deployment.

The primary Secure Firewall Management Center downloads URL category and reputation data required for Security Intelligence feed updates.

Updating Threat Intelligence after primary failure

After a primary failure, update threat intelligence data by confirming that the secondary Secure Firewall Management Center can access the internet and is promoted to active using the web interface.

Processing user data during Firewall Management Center failovers

Summary

The key components involved in processing user data during Firewall Management Center failovers are:

- Primary Firewall Management Center: The main device responsible for managing user and SGT mappings.
- Secondary Firewall Management Center: Takes over propagation of mappings when the primary fails.
- TS Agent identity source: Provides user-to-IP mappings.
- ISE/ISE-PIC identity source: Provides SGT mappings.
- Unknown users: Users not yet seen by identity sources.

This process ensures continuity of user and SGT mapping propagation during failover events and proper handling of unknown users after downtime.

Workflow

The process of handling user data during failover involves these stages:

1. Primary Firewall Management Center fails.
2. The secondary Firewall Management Center propagates user-to-IP mappings from the TS Agent identity source. It also sends SGT mappings from the ISE or ISE-PIC identity source to managed devices.

   Users not yet seen by identity sources are identified as Unknown.
3. After downtime, Unknown users are re-identified and processed according to the rules in the identity policy.

Configuration management on Firewall Management Center high availability pairs

Configuration management on Firewall Management Center high availability pairs is a deployment feature that

- allows only the active Firewall Management Center to manage devices and apply policies
- keeps both Firewall Management Centers continuously synchronized, and
- requires you to manually promote the standby appliance to active when the current active appliance fails.

Firewall Management Center high availability disaster recovery

A Firewall Management Center high availability disaster recovery is a failover mechanism that requires manual switchover between primary and secondary Firewall Management Center units.

Manual switchover in disaster recovery

Manual switchover is required when a disaster recovery situation occurs in a high availability pair.

- When the primary Firewall Management Center, FMC1, fails, use the web interface of the secondary Firewall Management Center, FMC2, to switch peers.
- This process is applicable if either the primary or secondary unit fails.

For more information, see Switch peer roles in the Firewall Management Center high availability pair, on page 334.

For restoring a failed unit, see Replace Firewall Management Centers in a high availability pair, on page 337.

Configure single sign-on for high availability pairs

Firewall Management Centers in a high availability configuration can support Single Sign-On, but you must keep these considerations in mind:

- SSO configuration is not synchronized between the members of the high availability pair. Configure SSO separately on each member.
- Both Firewall Management Centers in a high availability pair must use the same identity provider (IdP) for SSO. Configure a service provider application at the IdP for each Firewall Management Center for SSO.
- In a high availability pair of Firewall Management Centers where both support SSO, you must log in to the primary Firewall Management Center for the first time.
- When configuring SSO for Firewall Management Centers in a high availability pair:
  - If you configure SSO on the primary Firewall Management Center, you are not required to configure SSO on the secondary Firewall Management Center.
  - If you configure SSO on the secondary Firewall Management Center, you are required to configure SSO on the primary Firewall Management Center as well. (This is because SSO users must log in to the primary Firewall Management Center at least once before logging into the secondary Firewall Management Center.)

Related Topics

SAML Single Sign-On, on page 169

Firewall Management Center high availability behaviors during a backup

A Firewall Management Center high availability backup behavior is a system operation that

- pauses synchronization between peers during the backup process
- allows continued use of the active Firewall Management Center, but not the standby peer, and
- briefly disables processes on the active peer after backup completion, displaying a holding page until all processes resume.

Firewall Management Center high availability split-brain

A high availability split-brain is a failure state in high availability pairs that

- occurs when both Firewall Management Center appliances assume the active role,
- results from the original active appliance coming back online after a failure, and
- requires manual intervention to select an active appliance and demote the other to standby.

Device registration and policy configuration impact

The split-brain state impacts device registrations and policy configurations for the intended standby appliance.

Note The intended standby loses all of its device registrations and policy configurations when you resolve split-brain. For example, modifications to policies that exist on the intended standby but not on the intended active are lost. If the Firewall Management Center is in a high availability split-brain scenario where both appliances are active, and you register managed devices and deploy policies before you resolve split-brain, you must export any policies and unregister any managed devices from the intended standby Firewall Management Center before re-establishing high availability. After re-establishing high availability, register the managed devices and import policies to the intended active Firewall Management Center.

Split-brain scenario example

If the active Firewall Management Center in a high-availability pair fails due to power or network issues, promote the standby Firewall Management Center to active. When the original active peer becomes operational, both peers assume active status, which results in split-brain.

Degraded state scenario

If the active Firewall Management Center goes down or disconnects due to a network failure, you may break high availability or switch roles. The standby Firewall Management Center enters a degraded state, which is not split-brain.

Troubleshoot Firewall Management Center high availability

This reference provides troubleshooting information for common Firewall Management Center high availability operation errors, including error descriptions and recommended solutions.

Error	Description	Solution
You must reset your password on the active Firewall Management Center before you can log in to the standby.	You attempted to log in to the standby Firewall Management Center when a force password reset is enabled for your account.	Because the database is read-only for a standby Firewall Management Center, reset the password on the login page of the active Firewall Management Center.
500 Internal	This error may occur when you try to access the web interface during critical high availability operations, such as switching peer roles or pausing and resuming synchronization.	Wait until the operation completes before using the web interface.
System processes are starting, please wait Also, the web interface does not respond.	This error may occur when the system reboots manually or while recovering from a power down during high availability or data synchronization operations.	1. Access the Firewall Management Center shell and use the `manage_hadc.pl` command to access the Firewall Management Center high availability configuration utility. **Note** Run the utility as a root user, using `sudo`. 2. Pause mirroring operations by using option 5. Reload the Firewall Management Center web interface. 3. Use the web interface to resume synchronization. Choose **Integration** > **Other Integrations**, then click the **High Availability** tab and choose **Resume Synchronization**.

Error	Description	Solution
Device Registration Status:Host <*string*> is not reachable (occurs during initial configuration)	During the initial configuration of a Firewall Threat Defense, if the Firewall Management Center IP address and NAT ID are specified, the **Host** field can be left blank. However, in an HA environment with both the Firewall Management Centers behind a NAT, this error occurs when you add the Firewall Threat Defense on the secondary Firewall Management Center.	**1.** Delete the Firewall Threat Defense from primary Firewall Management Center. See *Delete a Device from the* Firewall Management Center in Cisco Secure Firewall Management Center Device Configuration Guide. **2.** Remove managers from Firewall Threat Defense using the **configure manager delete** command. See *Cisco Secure Firewall Threat Defense Command Reference*. **3.** Add Firewall Threat Defense to the Firewall Management Center with the IP address OR name of the Firewall Threat Defense device in the **Host** field. See *Add a Device to the* Firewall Management Center in Cisco Secure Firewall Management Center Device Configuration Guide.

Error	Description	Solution
Device Registration Status:Host <*string*> is not reachable (occurs when adding device to the secondary Firewall Management Center)	The error occurs when adding Firewall Threat Defense device to the secondary Firewall Management Center center in a high-availability deployment where both the secondary Firewall Management Center and the Firewall Threat Defense device are behind NAT.	On the standby Firewall Management Center web interface, click **Integration** > **Other Integrations** > **High Availability**. Under the pending device registration table, click the IP address of the pending device, then update it to the public IP address of the Firewall Threat Defense. Alternatively, follow these steps: 1. Access the Firewall Threat Defense shell and use the `show managers` command to get the standby Firewall Management Center entry identifier value. 2. In the Firewall Threat Defense shell, edit the standby Firewall Management Center hostname to the public IP address. Execute the `configure manager edit <standby_uuid> hostname <standby_ip>` command using the entry identifier value and the host IP address. For more information, see Resolve device registration using CLI in Firewall Management Center high availability, on page 333.

Requirements for Firewall Management Center high availability

Model support

See Requirements for hardware high availability, on page 322.

Virtual Model Support

See Understand virtual platform requirements, on page 322.

Supported domains

Global

User roles

Admin

Requirements for hardware high availability

Restrictions for hardware high availability

- All Firewall Management Center hardware supports high availability. Ensure that both peers are the same model.
- The peers may be physically and geographically separated from each other in different data centers.
- The bandwidth requirement for high availability configuration depends on various factors such as the size of the network, the number of managed devices, the volume of events and logs, and the size and frequency of configuration updates.
- For a typical Firewall Management Center high availability deployment, in case of high latency networks of close to 100 ms, a minimum of 5 Mbps network bandwidth between the peers is recommended.
- You can enhance the high availability synchronization speed by reducing the number of configuration versions saved on your Firewall Management Center. For more information, see *Set the Number of Configuration Versions* in Cisco Secure Firewall Management Center Device Configuration Guide. Note that this option is not supported on Secure Firewall Management Center versions 7.3.0 and 7.4.0.
- Ensure that both Firewall Management Centers have unique UUIDs. To check the UUID, review this file: `/etc/sf/ims.conf`.
- Do not restore a backup of the primary peer to the secondary.

These requirements apply to all Firewall Management Center hardware deployments configured for high availability.

Ensuring hardware compatibility, adequate bandwidth, and proper configuration versioning helps maintain reliable high availability and prevents synchronization issues.

See also License requirements for Firewall Management Center high availability configurations, on page 324 for additional licensing requirements.

Understand virtual platform requirements

High availability is supported for specific public and private cloud platforms. This reference lists the supported platforms and outlines licensing requirements for virtual deployments.

High availability is supported for these public cloud platforms:

- Amazon Web Services (AWS)
- Oracle Cloud Infrastructure (OCI)

High availability is also supported for these on-premises or private cloud platforms:

- Cisco HyperFlex
- VMware vSphere/VMware ESXi

Note Both Firewall Management Center Virtual peers in a high availability pair must be deployed on the same virtual platform. You cannot configure a high availability pair with peers deployed on different virtual platforms.

The Firewall Management Centers must have the same device management capacity (not supported on FMCv2) and be identically licensed. You also need one Firewall Threat Defense entitlement for each managed device. For more information, see License requirements for Firewall Management Center high availability configurations, on page 324.

Note If you are managing Version 7.0.x Classic devices only (NGIPSv or ASA FirePOWER), you do not need FMCv entitlements.

Requirements for software version consistency in high availability

Requirements for high availability software and update versions

Confirm that both management centers in a high availability configuration have the same major, minor, and maintenance software version. Install identical versions of the intrusion rule update, vulnerability database update, and LSP (Lightweight Security Package) on each management center. Ensure port 8305 is accessible between the two management centers to enable communication.

- The two Firewall Management Centers in a high availability configuration must have the same major (first number), minor (second number), and maintenance (third number) software version.
- The two Firewall Management Centers in a high availability configuration must have the same version of the intrusion rule update installed.
- The two Firewall Management Centers in a high availability configuration must have the same version of the vulnerability database update installed.
- The two Firewall Management Centers in a high availability configuration must have the same version of the LSP (Lightweight Security Package) installed.
- The two management centers in a high availability configuration must have port 8305 accessible between them for communication.

Access the **Appliance Information** widget to verify the software version, the intrusion rule update version and the vulnerability database update. By default, the widget appears on the **Status** tab of the **Detailed Dashboard** and the **Summary Dashboard**. For more information, see The appliance information widget, on page 361.

Warning You cannot establish high availability if the software versions, intrusion rule update versions, and vulnerability database update versions are not identical on both Firewall Management Center.

License requirements for Firewall Management Center high availability configurations

A high availability configuration is a deployment model that

- requires each device to have the same licenses whether managed by a single or multiple Firewall Management Centers in a high availability pair (hardware or virtual),
- registers only the active Firewall Management Center with the Smart Software Manager, and
- communicates with Smart Software Manager to release and assign license entitlements during failover events.

License assignment examples for high availability configurations

Example: If you want to enable advanced malware protection for two devices managed by a Firewall Management Center pair, buy two Malware licenses and two TM subscriptions, register the active Firewall Management Center with the Smart Software Manager, then assign the licenses to the two devices on the active Firewall Management Center.

Hardware management center license requirements

No special license is required for hardware Firewall Management Centers in a high availability pair.

Virtual management center license requirements

You need two identically licensed Firewall Management Center Virtuals. Firewall Management Center Virtual requires an individual license for each device it manages.

Each registered device in a Firewall Management Center Virtual high availability pair must have its own license. Although only the active FMC registers with Smart Software Manager at any time, entitlement usage accounts for both Firewall Management Center Virtuals in the high availability pair.

Entitlement calculation for Firewall Management Center Virtual high availability pairs is as follows:

Firewall Management Center Virtual entitlement = Number of Firewall Management Center Virtual peers in the high availability pair x Number of managed devices

In Specific License Reservation deployments, only the primary Firewall Management Center requires a Specific License Reservation.

Example: For the Firewall Management Center Virtual high availability pair managing 2 devices, you can use:

- Two (2) Firewall Management Center Virtual entitlements
- 2 device licenses

Note Firepower MCv Device Licenses seen on Smart Software Manager is 4

Example: For two Firewall Management Center Virtual high availability pair managing 2 devices each, you can use:

- Four (4) Firewall Management Center Virtual entitlements
- 4 device licenses

Note Firepower MCv Device Licenses seen on Smart Software Manager is 8

If you break the high availability pair, the Firewall Management Center Virtual entitlements associated with the secondary Firewall Management Center Virtual are released. (In the example, you would then have two standalone Firewall Management Center Virtual devices.)

Prerequisites for Firewall Management Center high availability

Recommendation: Complete prerequisites before establishing high availability

Before you establish the Firewall Management Center high availability pair, ensure that you complete these steps:

- Export required policies from the intended secondary Firewall Management Center to the intended primary Firewall Management Center. For more information, see Export configurations, on page 517.
- Make sure that the intended secondary Firewall Management Center does not have any devices added to it. Delete devices from the intended secondary Firewall Management Center and register these devices to the intended primary Firewall Management Center. For more information see *Delete a Device from the Firewall Management Center* and *Add a Device to the Firewall Management Center* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.
- Import the policies into the intended primary Firewall Management Center. For more information, see Import configurations, on page 518.
- On the intended primary Firewall Management Center, verify the imported policies, edit them as needed and deploy them to the appropriate device. For more information, see *Deploy Configuration Changes* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.
- On the intended primary Firewall Management Center, associate the appropriate licenses to the newly added devices. For more information see Assign licenses to a single device, on page 294.

You can now proceed to establish high availability. For more information, see Establish high availability for Firewall Management Center, on page 325.

Establish high availability for Firewall Management Center

Establishing high availability for Firewall Management Center enables redundancy and ensures continuous management operations in case one management center fails.

High availability setup can take several hours, depending on bandwidth between peers, the number of policies, and the number of devices registered to the active Firewall Management Center that must be synchronized to

the standby Firewall Management Center. You can view the High Availability page to check the status of the high availability peers.

Before you begin

- Confirm that both the Firewall Management Centers adhere to the high availability system requirements. For more information, refer to Requirements for Firewall Management Center high availability, on page 321.
- Confirm that you completed the prerequisites for establishing high availability. For more information, refer to Prerequisites for Firewall Management Center high availability, on page 325.
- In a multidomain deployment, perform this task in the Global domain.

Procedure

Step 1 Log in to the Firewall Management Center that you want to designate as the secondary.

Step 2 Choose **Integration** > **Other Integrations**, and then choose **High Availability**.

Step 3 Under Role for this Firewall Management Center, choose **Secondary**, and complete these steps:

a. Enter the hostname or IP address of the primary Firewall Management Center in the **Primary Firewall Management Center Host** text box.

You can leave this empty if the primary Firewall Management Center does not have an IP address reachable from the peer Firewall Management Center (the address can be public or private). In this case, use both the **Registration Key** and the **Unique NAT ID** fields. You need to specify the IP address of at least one Firewall Management Center to enable HA connection.

b. Enter a one-time-use registration key in the **Registration Key** text box.

The registration key is a user-defined alphanumeric value up to 37 characters in length. Use this key to register both the secondary and the primary Firewall Management Centers.

c. If you did not specify the primary IP address, or if you do not plan to specify the secondary IP address on the primary Firewall Management Center, then in the **Unique NAT ID** field, enter a unique alphanumeric ID. Refer to NAT environments, on page 92 for more information.

Step 4 Click **Register**.

Step 5 Using an account with Admin access, log in to the Firewall Management Center that you want to designate as the primary, and then choose **Integration** > **Other Integrations**, and then choose **High Availability**.

Step 6 Under Role for this Firewall Management Center, choose **Primary**, and complete these steps:

a. Enter the hostname or IP address of the secondary Firewall Management Center in the **Secondary Firewall Management Center Host** text box.

You can leave this empty if the secondary Firewall Management Center does not have an IP address reachable from the peer Firewall Management Center (which can be a public or private IP address). In this case, use both the **Registration Key** and the **Unique NAT ID** fields. You must specify the IP address of at least one Firewall Management Center to enable HA connection.

b. Enter the same one-time-use registration key in the **Registration Key** text box you used in step 6.

c. If required, enter the same NAT ID that you used in step 7 in the **Unique NAT ID** text box.

Step 7 Click **Register**.

After you complete these steps, the two Firewall Management Centers operate with high availability, providing redundancy and seamless failover for device management.

What to do next

After establishing the Firewall Management Center high availability pair, devices registered to the active Firewall Management Center are automatically registered to the standby Firewall Management Center.



Note When a registered device has a NAT IP address, automatic device registration fails and the secondary Firewall Management Center High Availability page lists the device as local, pending. You can then assign a different NAT IP address to the device on the standby Firewall Management Center High Availability page. If automatic registration otherwise fails on the standby Firewall Management Center, but the device appears to be registered to the active Secure Firewall Management Center, see Resolve device registration using CLI in Firewall Management Center high availability, on page 333.

Configure high availability for Firewall Management Centers hosted on public cloud

Configure high availability between your Firewall Management Centers hosted on public clouds so that your devices can register with both management centers for redundancy and failover.

Set up high availability to maintain continuous management and device registration in public cloud deployments. Use this procedure to establish resilient management for your firewalls with public IP addresses or hostnames for your management centers.

Before you begin

Ensure you have access to both the primary and secondary Firewall Management Centers hosted on public cloud. Make sure you have the registration key and NAT ID information.

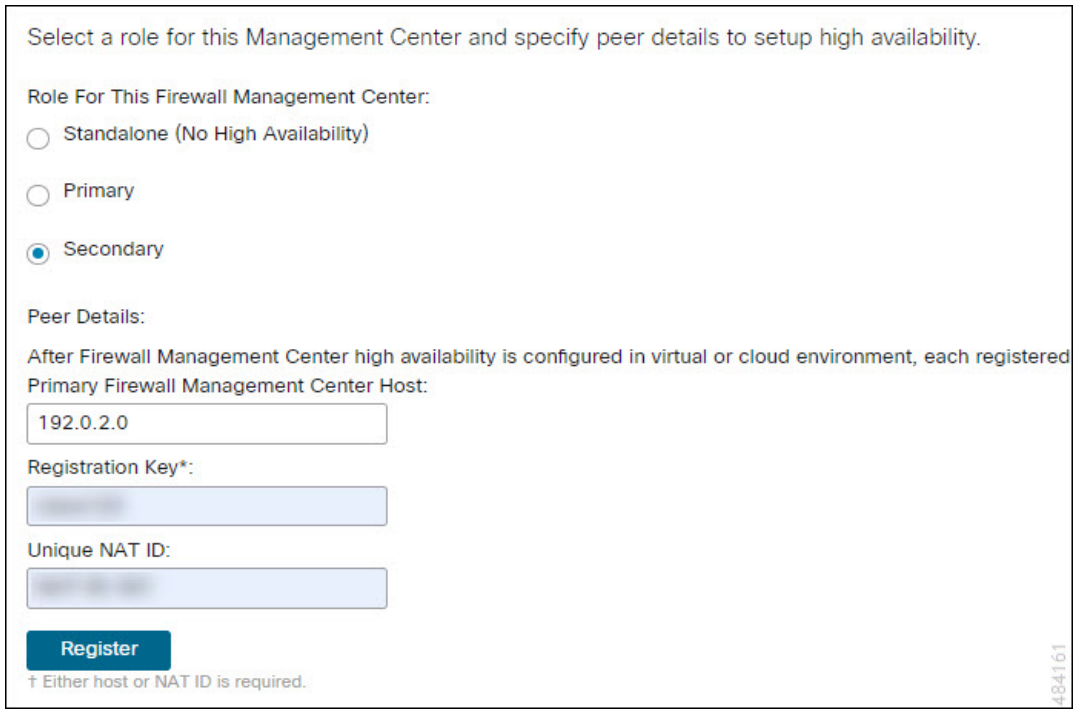
Follow these steps to configure high availability for Firewall Management Centers hosted on public cloud:

Procedure

Step 1 Configure high availability using public IP addresses or hostnames for both the primary and secondary management centers.

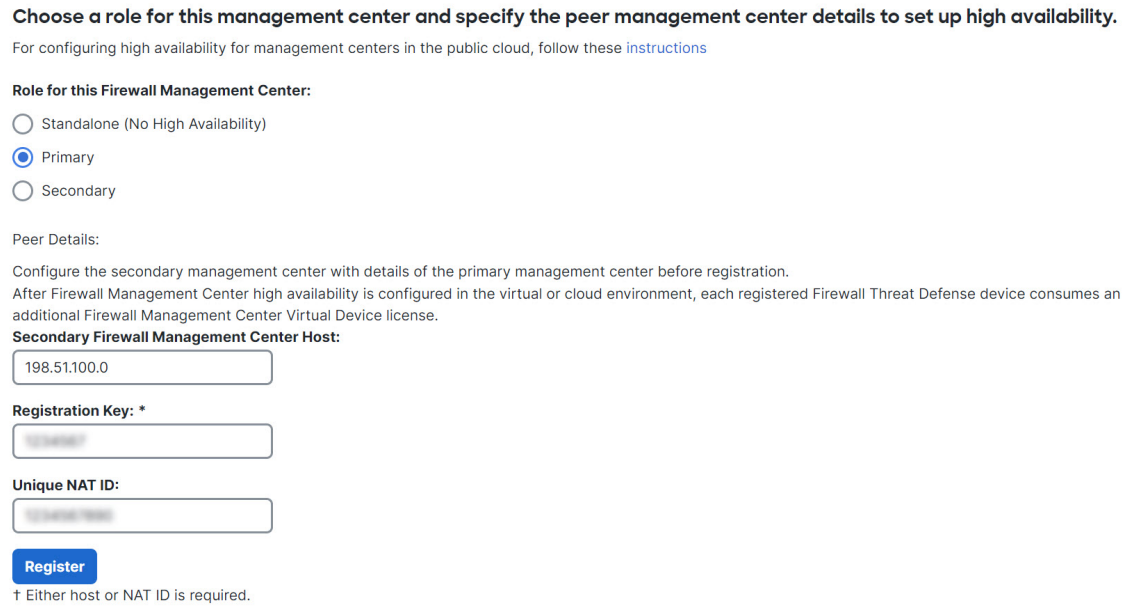
a. On the secondary Firewall Management Center, complete these steps:

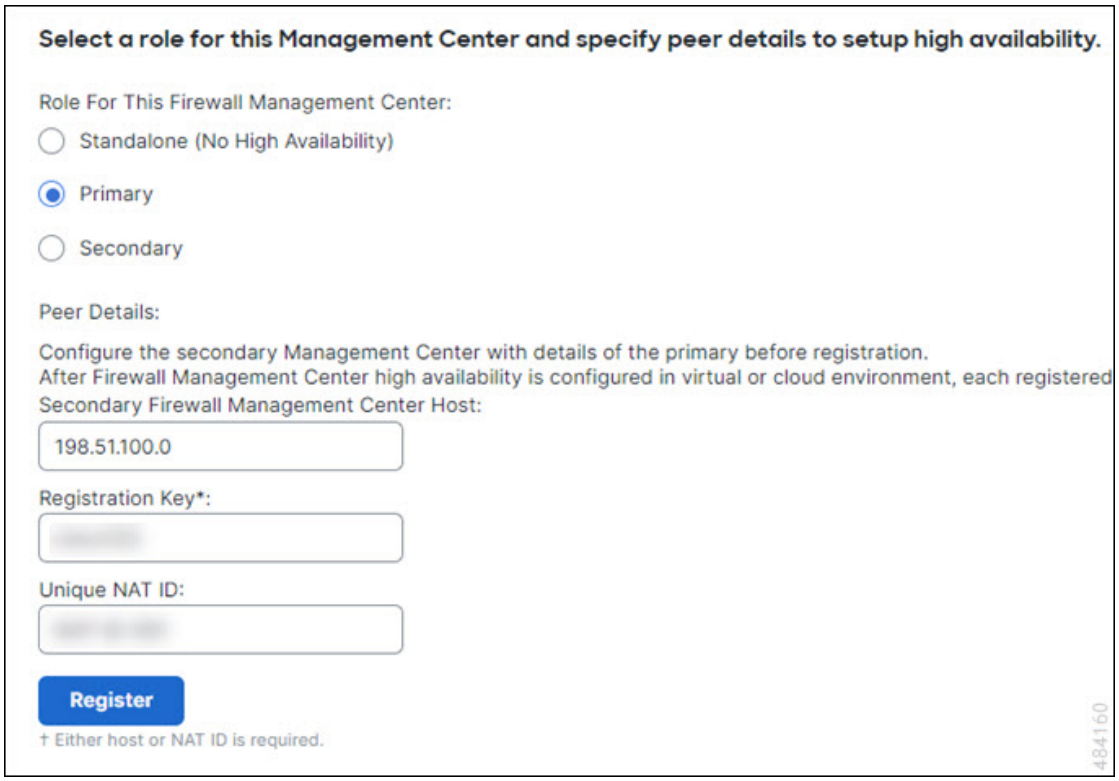
1. Choose **Secondary** as the **Role for this Firewall Management Center**.
2. Enter the public IP address or hostname for the secondary Firewall Management Center in the **Primary Firewall Management Center Host** field.
3. Enter your registration key.
4. Enter the same NAT ID that you used in the primary Firewall Management Center.



b. On the primary Firewall Management Center, do the following:

1. Choose **Primary** as the **Role for this Firewall Management Center**.
2. Enter the public IP address or hostname for the secondary Firewall Management Center in the **Secondary Firewall Management Center Host** field.
3. Enter the registration key.
4. Enter the unique NAT ID.



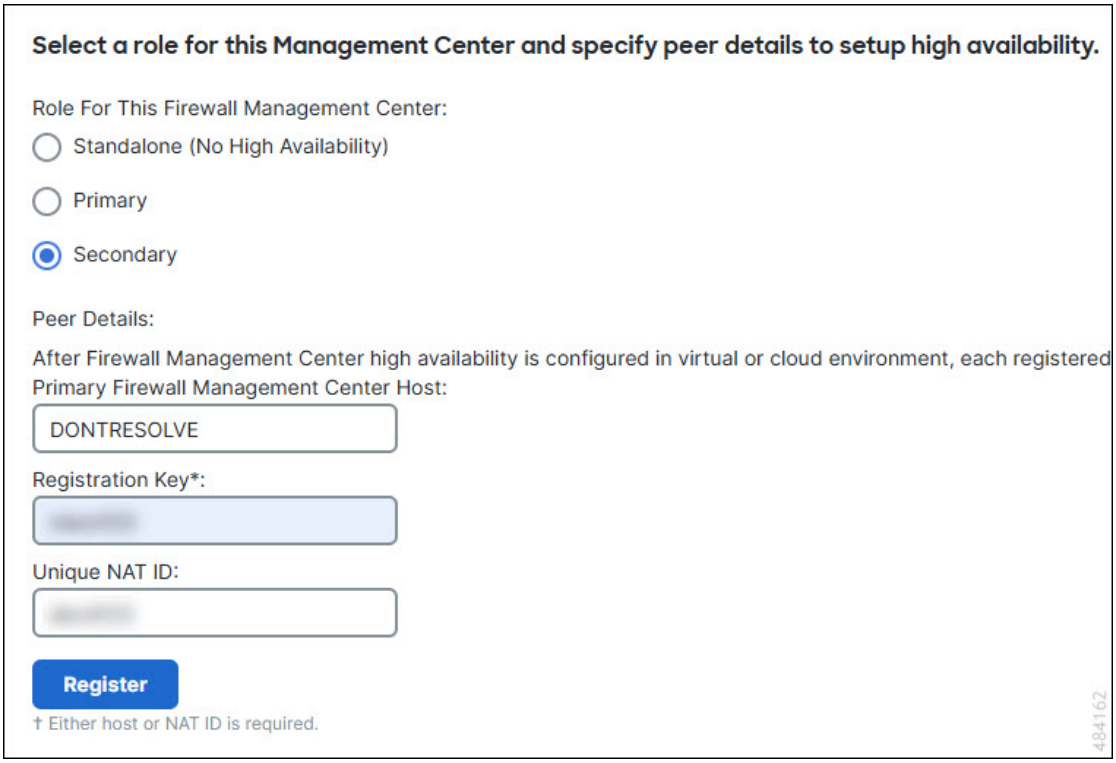


Step 2 Configure high availability using the public IP address or hostname for only the secondary management center.

a. On the secondary Firewall Management Center, do the following:

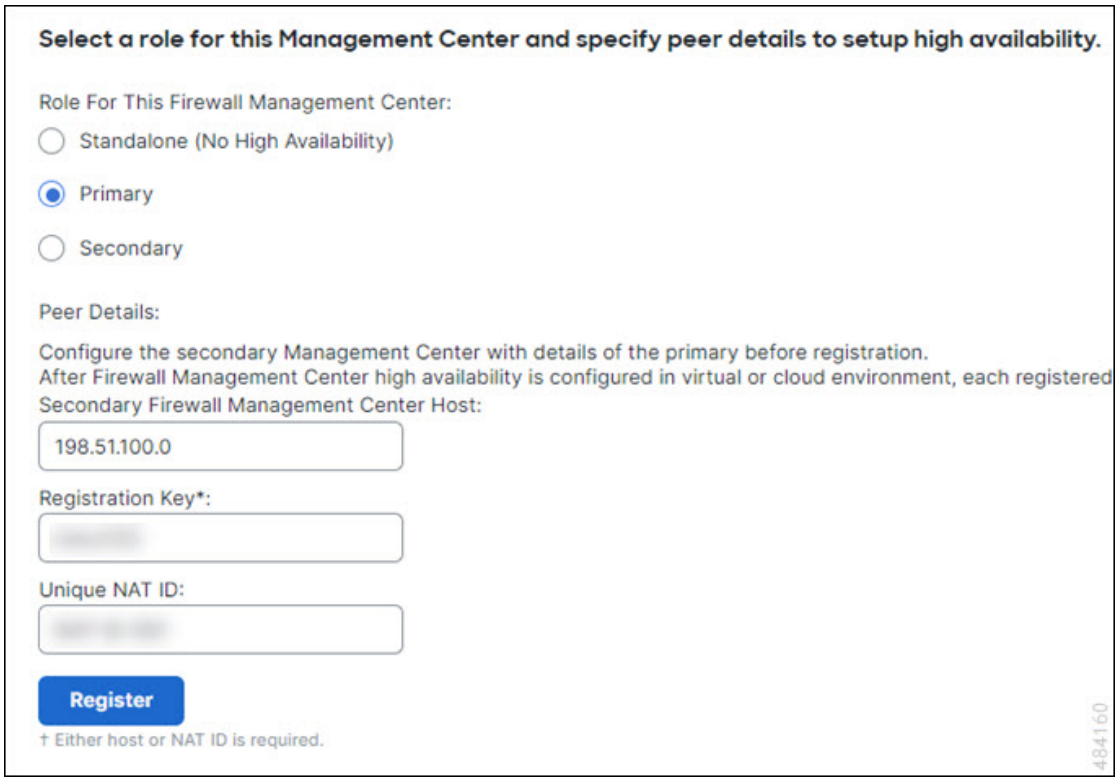
1. Choose **Secondary** as the **Role for this Firewall Management Center**.
2. Enter **DONTRESOLVE** in the **Primary Firewall Management Center Host** field.

3. Enter the registration key.
4. Enter the same NAT ID that you used in the primary Firewall Management Center.



b. On the primary Firewall Management Center, do the following:

1. Choose **Primary** as the **Role for this Firewall Management Center**.
2. Enter the public IP address or hostname for the secondary Firewall Management Center in the **Secondary Firewall Management Center Host** field.
3. Enter the registration key.
4. Enter the unique NAT ID.



View Firewall Management Center high availability status

View the high availability status of paired Firewall Management Center appliances to monitor their health, synchronization, and registration details.

After you identify your active and standby Firewall Management Centers, you can view information about the local Firewall Management Center and its peer.

Note In this context, Local Peer refers to the appliance where you are viewing the system status. Remote Peer refers to the other appliance, regardless of active or standby status.

Procedure

Step 1 Log into one of the Firewall Management Centers that you paired using high availability.

Step 2 Choose **Integration** > **Other Integrations**, and then choose **High Availability**.

Configure synchronization of configurations on Firewall Management Center high availability pairs

When you establish high availability between two Firewall Management Centers, configuration data is synced between them.

- License entitlements
- Access control policies
- Intrusion rules
- Malware and file policies
- DNS policies
- Identity policies
- SSL policies
- Prefilter policies
- Network discovery rules
- Application detectors
- Correlation policy rules
- Alerts
- Scanners
- Response groups
- Contextual cross-launch of external resources for investigating events
- Remediation settings, although you must install custom modules on both Firewall Management Centers. For more information on remediation settings, see Manage remediation modules, on page 1050.

Recommendation: Configure external access to the Firewall Management Center database only on the active peer

In a high availability setup, configure external access to the Firewall Management Center database only on the active peer. Configuring the standby peer for external database access leads to frequent disconnections.

This principle applies when configuring external access to the database in a high availability pair.

If you configure external access on the standby peer, frequent disconnections can disrupt database connectivity.

This recommendation ensures stable connectivity to the database and prevents unnecessary interruptions.

To restore connectivity after configuring the standby peer, pause and resume synchronization of the standby peer. For details on enabling external database access to Firewall Management Centers, see Enable external access to the database, on page 73.

Resolve device registration using CLI in Firewall Management Center high availability

This task enables you to resolve device registration issues in high availability deployments of Firewall Management Center using CLI commands.

If automatic device registration fails on the standby Firewall Management Center, but appears to be registered to the active Firewall Management Center, use this procedure to manually resolve registration. This is relevant when managing devices in high availability environments.



Warning If you do an RMA of secondary Firewall Management Center or add a secondary Firewall Management Center, the managed devices are unregistered, and their configuration can get deleted as a result.

Procedure

Step 1 Delete the device from the active Firewall Management Center. Refer to *Delete (Unregister) a Device from the Firewall Management Center* in Cisco Secure Firewall Management Center Device Configuration Guide.

Step 2 Trigger automatic registration of the device on the standby Firewall Management Center:

a. Log in to the CLI for the affected device.

b. Run the CLI command: **configure manager delete**.

This command disables and removes the current Firewall Management Center.

c. Run the CLI command: **configure manager add**.

This command configures the device to initiate a connection to a Firewall Management Center.

Tip
Configure remote management on the device, only for the active Firewall Management Center. When you establish high availability, the devices are automatically registered to the standby Firewall Management Center.

d. Log in to the active Firewall Management Center and register the device.

Step 3 If the standby Firewall Management Center is behind NAT, edit the hostname of the standby Firewall Management Center:

a. Access the Firewall Threat Defense shell and use the `show managers` command to get the standby Firewall Management Center entry identifier value.

b. In the Firewall Threat Defense shell, edit the standby Firewall Management Center hostname to the public IP address. Execute the `configure manager edit <standby_uuid> hostname <standby_ip>` command using the entry identifier value and the host IP address.

Switch peer roles in the Firewall Management Center high availability pair

Switching peer roles in a high availability pair allows you to promote the standby Firewall Management Center to active status if the current active appliance fails. This ensures that system functionality restricted to the active peer remains available.

Use this task when the active Firewall Management Center in a high availability deployment becomes unavailable or you need to manually switch roles between peers. The system restricts some functionality to the active appliance, so switching roles maintains operational continuity.

Procedure

Step 1 Log into one of the Firewall Management Centers that you paired using high availability.

Step 2 Choose **Integration** > **Other Integrations**, and then choose **High Availability**.

Step 3 Choose **Switch Peer Roles** to change the local role between Active and Standby. This switches the active and standby roles between the two peers, and the Primary and Secondary designations stay the same. In this context, *role* means the active or standby status of the management center, not the primary or secondary designation set during HA setup.

Pause communication between paired Firewall Management Centers

If you want to temporarily disable high availability, you can disable the communications channel between the Firewall Management Centers. You can pause synchronization from an active or standby peer.

Procedure

Step 1 Log into one of the Firewall Management Centers that you paired using high availability.

Step 2 Choose **Integration** > **Other Integrations**, and then choose **High Availability**.

Step 3 Choose **Pause Synchronization**.

Restart communication between paired Firewall Management Centers

Restarting communication between paired Firewall Management Centers allows you to resume high availability synchronization after it has been temporarily disabled.

If you temporarily disable high availability, you can restart high availability by enabling the communications channel between the Firewall Management Centers. You can resume synchronization from an active or standby peer.

Procedure

Step 1 Log into one of the Firewall Management Centers that you paired using high availability.

Step 2 Choose **Integration** > **Other Integrations**, and then choose **High Availability**.

Step 3 Choose **Resume Synchronization**.

Change the IP address of the Firewall Management Center in a high availability pair

Changing the IP address of one peer in a Firewall Management Center high availability pair does not automatically update the other peer's information. To maintain proper synchronization and connectivity between peers, you must manually update the IP address on both appliances.

Procedure

Step 1 Log in to the peer Firewall Management Center where you want to manually modify the IP address of the other peer manager.

Step 2 Choose **Integration** > **Other Integrations**.

Step 3 Choose **High Availability**.

Step 4 Choose **Peer Manager**.

Step 5 Choose **Edit** ( ).

Step 6 Enter the display name of the appliance, which is used only within the context of the system.

Entering a different display name does not change the host name for the appliance.

Step 7 Enter the fully qualified domain name, that is the name that resolves through the local DNS to a valid IP address (the host name), or the host IP address.

Step 8 Click **Save**.

Disable Firewall Management Center high availability

Disable high availability for Firewall Management Center to break the HA pair and manage devices independently.

Breaking HA allows you to control device management from either console or stop managing devices altogether.

High availability for Firepower Management Centers is not available in versions 6.0 or 6.0.1. This procedure is relevant when you need to break an existing HA pair and reassign device management.

Use this task when you want to manage devices from a single console or discontinue device management from both consoles.

After breaking HA, you may need to re-register device licenses. Access control policies remain on the standby Firewall Management Center.

Procedure

Step 1 Log into one of the Firewall Management Centers in the high availability pair.

Step 2 Choose **Integration** > **Other Integrations**, and then choose **High Availability**.

Step 3 Choose **Break High Availability**.

Step 4 Choose one of the options for handling managed devices:

- To control all managed devices with this Firewall Management Center, choose **Manage registered devices from this console**. All devices will be unregistered from the peer.
- To control all managed devices with the other Firewall Management Center, choose **Manage registered devices from peer console**. All devices will be unregistered from this Firewall Management Center.
- To stop managing devices altogether, choose **Stop managing registered devices from both consoles**. All devices will be unregistered from both Firewall Management Centers.

Note

- If you choose to manage the registered devices from the secondary Firewall Management Center, the devices will be unregistered from the primary Firewall Management Center. Devices will then be registered to the secondary Firewall Management Center. License applications on these devices are deregistered due to the high availability break operation. You must re-register (enable) licenses on the devices from the secondary Firewall Management Center. For more information see Assign licenses to devices, on page 294.
- The standby management center retains the access control policies after the HA break is complete.

Step 5 Click **OK**.

Replace Firewall Management Centers in a high availability pair

If you need to replace a failed unit in the Firewall Management Center high availability pair, you must follow one of these procedures. The table lists four possible failure scenarios and their corresponding replacement procedures.

Failure Status	Data Backup Status	Replacement Procedure
Primary Firewall Management Center failed	Data backup successful	Replace a failed primary Firewall Management Center (successful backup), on page 337
	Data backup not successful	Replace a failed primary Firewall Management Center (unsuccessful backup), on page 338
Secondary Firewall Management Center failed	Data backup successful	Replace a failed secondary Firewall Management Center (successful backup), on page 339
	Data backup not successful	Replace a failed secondary Firewall Management Center (unsuccessful backup), on page 340

Replace a failed primary Firewall Management Center (successful backup)

This task enables you to replace a failed primary Firewall Management Center in a high availability pair using a successful backup, ensuring minimal disruption and restoration of high availability operations.

Two Firewall Management Centers, *FMC1* and *FMC2*, are part of a high availability pair. *FMC1* is the primary and *FMC2* is the secondary. Use this task to replace a failed primary Firewall Management Center, *FMC1*, when data backup from the primary is successful.

Procedure

Step 1 Contact Support to request a replacement for a failed Firewall Management Center - *FMC1*.

Step 2 When the primary Firewall Management Center - *FMC1* fails, access the web interface of the secondary Firewall Management Center - *FMC2* and switch peers. For more information, refer to Switch peer roles in the Firewall Management Center high availability pair, on page 334.

This promotes the secondary Firewall Management Center - *FMC2* to active.

You can use *FMC2* as the active Firewall Management Center until the primary Firewall Management Center - *FMC1* is replaced.

Caution
Do not break Firewall Management Center high availability from *FMC2*, since licenses that were synced to *FMC2* from *FMC1* (before failure), will be removed from *FMC2* and you will be unable to perform any deploy actions from *FMC2*.

Step 3 Reimage the replacement Firewall Management Center with the same software version as *FMC1*.

Step 4 Restore the data backup retrieved from *FMC1* to the new Firewall Management Center.

Step 5 Install required Firewall Management Center patches, geolocation database (GeoDB) updates, vulnerability database (VDB) updates, and system software updates to match *FMC2*.

The new Firewall Management Center and *FMC2* will now both be active peers, resulting in a high availability split-brain.

Step 6 When the Firewall Management Center web interface prompts you to choose an active appliance, select *FMC2* as active.

This syncs the latest configuration from *FMC2* to the new Firewall Management Center - *FMC1*.

Step 7 When the configuration syncs successfully, access the web interface of the secondary Firewall Management Center - *FMC2* and switch roles to make the primary Firewall Management Center - *FMC1* active. For more information, refer to Switch peer roles in the Firewall Management Center high availability pair, on page 334.

What to do next

High availability has now been re-established and the primary and the secondary Firewall Management Centers will now work as expected.

Replace a failed primary Firewall Management Center (unsuccessful backup)

This task enables you to replace a failed primary Firewall Management Center in a high availability pair when data backup from the primary is unsuccessful.

This task restores high availability and ensures proper license management after a primary failure.

Two Firewall Management Centers, *FMC1* (primary) and *FMC2* (secondary), are part of a high availability pair. If the primary fails and backup is unsuccessful, you must perform these steps to restore the system.

Procedure

Step 1 Contact Support to request a replacement for a failed Firewall Management Center - *FMC1*.

Step 2 When the primary Firewall Management Center - *FMC1* fails, access the web interface of the secondary Firewall Management Center - *FMC2* and switch peers. For more details, refer to Switch peer roles in the Firewall Management Center high availability pair, on page 334.

This promotes the secondary Firewall Management Center - *FMC2* to active.

You can use *FMC2* as the active Firewall Management Center until the primary Firewall Management Center - *FMC1* is replaced.

Caution
Do not break Firewall Management Center High Availability from *FMC2*, since licenses that were synced to *FMC2* from *FMC1* (before failure), will be removed from *FMC2* and you will be unable to perform any deploy actions from *FMC2*.

Step 3 Reimage the replacement Firewall Management Center with the same software version as *FMC1*.

Step 4 Install required Firewall Management Center patches, geolocation database (GeoDB) updates, vulnerability database (VDB) updates and system software updates to match *FMC2*.

Step 5 Deregister one of the Firewall Management Centers - *FMC2* from the Cisco Smart Software Manager. For more information, see Deregister the Firewall Management Center, on page 296.

Deregistering Firewall Management Center from the Cisco Smart Software Manager removes the Management Center from your virtual account. All license entitlements associated with the Firewall Management Center release back to your virtual account. After deregistration, the Firewall Management Center enters Enforcement mode where no update or changes on licensed features are allowed.

Step 6 Access the web interface of the secondary Firewall Management Center - *FMC2* and break Firewall Management Center high availability. For more information, see Disable Firewall Management Center high availability, on page 336. When prompted to select an option for handling managed devices, choose **Manage registered devices from this console**.

As a result, licenses that were synced to the secondary Firewall Management Center- *FMC2* will be removed, and you cannot perform deployment activities from *FMC2*.

Step 7 Re-establish Firewall Management Center high availability, by setting up the Firewall Management Center - *FMC2* as the primary and Firewall Management Center - *FMC1* as the secondary. For more information, see Establish high availability for Firewall Management Center, on page 325.

Step 8 Register a Smart License to the primary Firewall Management Center - *FMC2*. For more information see Register the Firewall Management Center with the Smart Software Manager, on page 288.

What to do next

High availability has now been re-established and the primary and the secondary Firewall Management Centers will now work as expected.

Replace a failed secondary Firewall Management Center (successful backup)

This task enables you to replace a failed secondary Firewall Management Center when a successful backup exists, ensuring continued high availability and minimal disruption.

Two Firewall Management Centers—*FMC1* (primary) and *FMC2* (secondary)—form a high availability pair. If the secondary fails but its backup is successful, you can restore high availability by replacing it and restoring the backup.

Before you begin

Verify that the data backup from the failed secondary Firewall Management Center is successful.

Procedure

Step 1 Contact Support to request a replacement for a failed Firewall Management Center - *FMC2*.

Step 2 Continue to use the primary Firewall Management Center - *FMC1* as the active Firewall Management Center.

Step 3 Reimage the replacement Firewall Management Center with the same software version as *FMC2*.

Step 4 Restore the data backup from *FMC2* to the new Firewall Management Center.

Step 5 Install the required Firewall Management Center patches, geolocation database (GeoDB) updates, vulnerability database (VDB) updates, and system software updates to match *FMC1*.

Step 6 Resume data synchronization (if paused) from the web interface of the new Firewall Management Center - *FMC2*, to synchronize the latest configuration from the primary Firewall Management Center - *FMC1*. For more information, refer to Restart communication between paired Firewall Management Centers, on page 335.

Classic and Smart Licenses work seamlessly.

What to do next

High availability has now been re-established and the primary and the secondary Firewall Management Centers will now work as expected.

Replace a failed secondary Firewall Management Center (unsuccessful backup)

This task enables you to restore high availability by replacing a failed secondary Firewall Management Center when backup from the secondary is unsuccessful.

- Ensures the primary management center remains active during the replacement process.
- Synchronizes configuration and licensing between the primary and replacement secondary management centers.

Two Firewall Management Centers, *FMC1* (primary) and *FMC2* (secondary), form a high availability pair. If the secondary fails and backup is unsuccessful, you must replace it to restore high availability.

This procedure is relevant when the secondary management center cannot be recovered from backup and must be replaced to maintain system redundancy and continuity.

Applicable in high availability deployments of Secure Firewall Management Centers.

Before you begin

Ensure you have access to the primary Firewall Management Center and the replacement device for the secondary.

- Verify the replacement device is available and ready for reimaging.
- Obtain the required software version and updates for the replacement device.

Procedure

Step 1 Contact Support to request a replacement for a failed Firewall Management Center - *FMC2*.

Step 2 Continue to use the primary Firewall Management Center - *FMC1* as the active Firewall Management Center.

Step 3 Reimage the replacement Firewall Management Center with the same software version as *FMC2*.

Step 4 Install the required Firewall Management Center patches, geolocation database (GeoDB) updates, vulnerability database (VDB) updates and system software updates to match *FMC1*.

Step 5 Access the web interface of the primary Firewall Management Center - *FMC1* and break Firewall Management Center high availability. For more information, refer to Disable Firewall Management Center high availability,

on page 336. When prompted to select an option for handling managed devices, choose **Manage registered devices from this console**.

Step 6 Re-establish Firewall Management Center high availability, by setting up the Firewall Management Center - *FMC1* as the primary and Firewall Management Center - *FMC2* as the secondary. For more information, refer to Establish high availability for Firewall Management Center, on page 325.

- When high availability is successfully established, the latest configuration from the primary Firewall Management Center - *FMC1* is synchronized to the secondary Firewall Management Center - *FMC2*.
- Both Classic and Smart licenses function without issues.

What to do next

High availability has been re-established. The primary and secondary Firewall Management Centers now operate as expected.

Firewall Management Center high availability disaster recovery

A Firewall Management Center high availability disaster recovery is a failover mechanism that requires manual switchover between primary and secondary Firewall Management Center units.

Manual switchover in disaster recovery

Manual switchover is required when a disaster recovery situation occurs in a high availability pair.

- When the primary Firewall Management Center, FMC1, fails, use the web interface of the secondary Firewall Management Center, FMC2, to switch peers.
- This process is applicable if either the primary or secondary unit fails.

For more information, see Switch peer roles in the Firewall Management Center high availability pair, on page 334.

For restoring a failed unit, see Replace Firewall Management Centers in a high availability pair, on page 337.

Restore management center in a high-availability pair (no hardware failure)

To restore a Firewall Management Center high-availability pair when there is no hardware failure, use these procedures.

- Restore a backup on the primary Firewall Management Center, on page 342
- Restore a backup on the secondary Firewall Management Center, on page 342

Restore a backup on the primary Firewall Management Center

Use this task when you need to recover the primary Firewall Management Center from a backup, such as after a configuration error or data corruption, and there is no hardware failure or replacement involved.

Before you begin

- Confirm that there is no hardware failure or that the Firewall Management Center has not been replaced.
- You are familiar with the backup and restore process. See Backup/Restore, on page 471.

Procedure

Step 1 Check whether a backup of the primary Firewall Management Center exists on either local storage in /var/sf/backup/ or a remote network volume.

Step 2 Pause synchronization on the primary Firewall Management Center. Choose **Integration** > **Other Integrations**, and then choose **High Availability** tab to pause synchronization.

Step 3 Restore the backup on the primary Firewall Management Center. The Firewall Management Center reboots when the restoration is complete.

Step 4 Once the primary Firewall Management Center is active and its user interface is reachable, resume synchronization on the secondary Firewall Management Center. Choose **Integration** > **Other Integrations**, and then choose **High Availability** tab to resume synchronization.

Restore a backup on the secondary Firewall Management Center

This task is performed when you need to restore the secondary Firewall Management Center from a backup, such as after a configuration issue or to synchronize with the primary Firewall Management Center in a high availability setup.

Before you begin

- There is no hardware failure and replacement of the Firewall Management Center.
- You are familiar with the backup and restore process. See Backup/Restore, on page 471.

Procedure

Step 1 Verify if backup of the secondary Firewall Management Center is available—either a local storage in /var/sf/backup/, or a remote network volume.

Step 2 Pause synchronization on the primary Firewall Management Center. Choose **Integration** > **Other Integrations**, and then choose**High Availability** tab to pause synchronization.

Step 3 Restore the backup on the secondary Firewall Management Center. The Firewall Management Center reboots when the restoration is complete.

Step 4 Once the secondary Firewall Management Center is active and its user interface is reachable, resume synchronization on the primary Firewall Management Center. Choose **Integration** > **Other Integrations**, and then choose **High Availability** tab to resume synchronization.

Unified backups of management centers in high availability

A unified backup of management centers is a backup solution that

- creates a single backup file for both active and standby Firewall Management Centers
- removes redundant data and storage constraints, and
- enables recovery of a new RMA device in high availability environments.

Benefits of unified backups

Unified backups offer several advantages compared to normal backups.

- With unified backup, you do not need to take separate backups on active and standby Firewall Management Centers.
- A unified backup removes redundant data in backups and storage constraints.
- In a normal backup, if the primary unit fails and a backup of the secondary unit is not available, you must break the high availability pairing for the secondary RMA. Unified backup eliminates this requirement.
- You can use the unified backup to recover a new RMA device if an unanticipated incident occurs.
- Typically, you cannot schedule a backup of a standby unit. When you schedule a unified backup, the system backs up both active and standby units.
- While executing unified backup, you do not have to pause the high availability synchronization to perform backup on the standby unit.

Unified backup can only be used for configuration-only backups. If you need an eventing or TID backup, you must take separate backups for the active and standby Firewall Management Centers.

When you select configuration-only backup, by default, unified backup is applied. In a unified backup, if the active Firewall Management Center cannot retrieve a backup tar file from the standby Firewall Management Center, the system generates a normal backup file for the active unit. You can use this file for restoration.

Unified backup file identification

You can identify the unified backup file by its name. A prefix "Unified" is added to the unified backup file name. You can select the Firewall Management Center to restore and also select its State (Active/Standby).

Split-brain conflict prevention

Ensure that you select the appropriate state of the restored Firewall Management Center to prevent Split-Brain conflict.

Restore Management Center from unified backup

Restore the management center configuration using a unified backup file to recover system settings and ensure business continuity after a failure or migration.

Use this task when you need to recover or migrate your Firewall Management Center from a configuration-only unified backup. This is typically required after a system failure, hardware replacement, or when moving to a new deployment.

This procedure restores only the configuration and does not include event or log data.

Procedure

Step 1 Log into the Firewall Management Center you want to restore.

Step 2 Select **System** (⚙) > **Tools** > **Backup/Restore**.

The Backup Management page lists all locally and remotely stored backup files including the unified backup file (configuration-only).

If the unified backup file is not in the list and you have it saved on your local computer, click **Upload Backup**; refer to Manage backups and remote storage, on page 499.

Step 3 Select the unified backup file that you want to restore and click **Restore**.

Step 4 In the **Restore Backup** page, select which unit you want to restore. The unified backup stores the backup configuration of both primary and secondary units. You need to choose which unit to restore.

Step 5 To select the state of the restored Firewall Management Center, click the **Active** or **Standby** radio button. Check the role and state of your active management center. Prevent both peers from sharing the same role and state configuration. If you select the wrong role or state, high availability failure can occur.

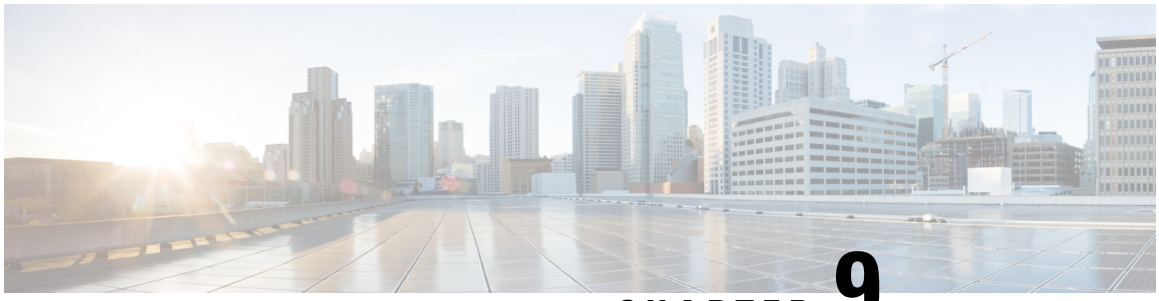
Step 6 Click **Restore**, and then **Confirm Restore** to begin the restoration.

History for Firewall Management Center high availability

The History table summarizes the introduction and evolution of the high availability in Cisco Minimum Firewall Management Center across releases, enabling quick tracking of feature availability and changes

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Single backup file for high availability Firewall Management Centers.	7.4.1 7.2.6	Any	When performing a configuration-only backup of the active Firewall Management Center in a high availability pair, the system now creates a single backup file which you can use to restore either unit.
Support for high availability on AWS and OCI.	7.1.0	Any	We now support high availability on Firewall Management Center Virtual for AWS and OCI.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Support for high availability on HyperFlex.	7.0.0	Any	We now support high availability on Firewall Management Center Virtual for HyperFlex.
Support for high availability on VMware.	6.7.0	Any	We now support high availability on Firewall Management Center Virtual for VMware.
Single sign-on.	6.7.0	Any	When configuring one or both members of a high availability pair for single sign-on, you must take into account special considerations.



CHAPTER 9

Security certifications compliance

These topics describe how to configure your system to comply with security certifications standards:

- Security certifications compliance modes, on page 347
- Security certifications compliance characteristics, on page 348
- Security certifications compliance recommendations, on page 349
- Enable security certifications compliance, on page 352

Security certifications compliance modes

A security certifications compliance mode refers to a set of security standards established by the U.S. Department of Defense and global certification organizations. Organizations must use only equipment and software that comply with these standards.

Security certification standards

Secure Firewall supports compliance with these security certifications standards:

- **Common Criteria (CC)**: A global standard established by the international Common Criteria Recognition Arrangement, defining properties for security products.
- **Unified Capabilities Approved Products List (UCAPL)**: A list of products meeting security requirements established by the U.S. Defense Information Systems Agency (DISA).

Note The U.S. Government has changed the name of the Unified Capabilities Approved Products List (UCAPL) to the Department of Defense Information Network Approved Products List (DODIN APL). References to UCAPL in this documentation and the Secure Firewall Management Center web interface can be interpreted as references to DODIN APL.

- **Federal Information Processing Standards (FIPS) 140**: A requirements specification for encryption modules used in secure systems.

Organizations can enable security certifications compliance in CC mode or UCAPL mode. Enabling these modes does not guarantee compliance with every requirement of the selected security mode. For more information on hardening procedures, refer to the product guidelines provided by the certifying entity.

Caution After you enable a security certifications compliance mode, you cannot disable it. To remove an appliance from CC or UCAPL mode, you must reimage the appliance.

Security certifications compliance characteristics

This table describes changes to system behaviors and compliance characteristics when CC (Common Criteria) or UCAPL (Unified Capabilities Approved Product List) mode is enabled. Restrictions on login accounts apply to command line access. These restrictions do not apply to web interface access.

System Change	Secure Firewall Management Center		Classic Managed Devices		Secure Firewall Threat Defense	
	CC Mode	UCAPL Mode	CC Mode	UCAPL Mode	CC Mode	UCAPL Mode
FIPS compliance is enabled.	Yes	Yes	Yes	Yes	Yes	Yes
The system does not allow remote storage for backups or reports.	Yes	Yes	—	—	—	—
The system starts an additional system audit daemon.	No	Yes	No	Yes	No	No
The system boot loader is secured.	No	Yes	No	Yes	No	No
The system applies additional security to login accounts.	No	Yes	No	Yes	No	No
The system disables the reboot key sequence Ctrl+Alt+Del.	No	Yes	No	Yes	No	No
The system enforces a maximum of ten simultaneous login sessions.	No	Yes	No	Yes	No	No
Passwords must be at least 15 characters long, and must consist of alphanumeric characters of mixed case and must include at least one numeric character.	No	Yes	No	Yes	No	No
The minimum required password length for the local `admin` user can be configured using the local device CLI.	No	No	No	No	Yes	Yes
Passwords cannot be a word that appears in a dictionary or include consecutive repeating characters.	No	Yes	No	Yes	No	No
The system locks out users other than `admin` after three failed login attempts in a row. In this case, the password must be reset by an administrator.	No	Yes	No	Yes	No	No
The system stores password history by default.	No	Yes	No	Yes	No	No

System Change	Secure Firewall Management Center		Classic Managed Devices		Secure Firewall Threat Defense	
	CC Mode	UCAPL Mode	CC Mode	UCAPL Mode	CC Mode	UCAPL Mode
The `admin` user can be locked out after a maximum number of failed login attempts configurable through the web interface.	Yes	Yes	Yes	Yes	—	—
The `admin` user can be locked out after a maximum number of failed login attempts configurable through the local appliance CLI.	No	No	Yes, regardless of security certifications compliance enablement.	Yes, regardless of security certifications compliance enablement.	Yes	Yes
The system automatically rekeys an SSH session with an appliance: • After a key has been in use for one hour of session activity • After a key has been used to transmit 1 GB of data over the connection	Yes	Yes	Yes	Yes	Yes	Yes
The system performs a file system integrity check (FSIC) at boot-time. If the FSIC fails, Secure Firewall software does not start, remote SSH access is disabled, and you can access the appliance only via local console. If this happens, contact Cisco TAC.	Yes	Yes	Yes	Yes	Yes	Yes

Security certifications compliance recommendations

Cisco recommends that you observe the following best practices when using a system with security certifications compliance enabled:

- Enable security certifications compliance first on the Firewall Management Center, then enable it in the same mode on all managed devices.
- The Firewall Management Center will not receive event data from a managed device unless both are operating in the same security certifications compliance mode.
- For all users, enable password strength checking and set the minimum password length to the value required by the certifying agency.
- To use Firewall Management Centers in a high-availability configuration, configure them both to use the same security certifications compliance mode before forming the high availability pair.

- When you configure Firewall Threat Defense on a Firepower 4100/9300 to operate in CC or UCAPL mode, you should also configure the Firepower 4100/9300 to operate in CC mode. For more information, see the *Cisco Firepower 4100/9300 FXOS Chassis Manager Configuration Guide*.

Do not configure the system to use any features:

- Email reports, alerts, or data pruning notifications.
- Nmap Scan, Cisco IOS Null Route, Set Attribute Value, or ISE EPS remediations.
- Remote storage for backups or reports.
- Third-party client access to the system database.
- External notifications or alerts transmitted via email (SMTP), SNMP trap, or syslog.
- Audit log messages transmitted to an HTTP server or to a syslog server without using SSL certificates to secure the channel between the appliance and the server.

Configure authentication and access settings appropriately for your security certifications compliance mode.

- Do not enable external authentication using LDAP or RADIUS in deployments using CC mode.
- Do not enable CACs in deployments using CC mode.
- Disable access to the Firewall Management Center and managed devices via the Secure Firewall REST API in deployments using CC or UCAPL mode.
- Enable CACs in deployments using UCAPL mode.
- Do not configure SSO in deployments using CC mode.

The system does not support CC or UCAPL mode for:

- Firewall Threat Defense devices in clusters
- Firewall Threat Defense container instances on the Firepower 4100/9300
- Exporting event data to an external client using eStreamer.

Appliance hardening

To further harden your system, see the latest versions of the *Cisco Secure Firewall Management Center Hardening Guide* and the *Cisco Secure Firewall Threat Defense Hardening Guide*.

Refer to these topics for configuration steps:

Identity and access management

- Users, on page 139
- Custom user roles for the web interface, on page 226
- Add or edit an internal user, on page 148
- Logging into the Management Center, on page 29
- Session timeouts, on page 120

Monitoring, Auditing, and Compliance

- Audit and Syslog, on page 423
- Audit logs, on page 53
- Audit log certificates, on page 57
- Security certifications compliance, on page 347

System and Network Configuration

- Licenses, on page 271
- Time synchronization, on page 123
- DNS caches, on page 66
- Access lists, on page 51
- HTTPS certificates, on page 74
- Scheduling, on page 505
- Configure SSH for remote storage, on page 118
- *SNMP* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *Site-to-Site VPNs for Firewall Threat Defense* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *SMTP Server* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *Configure SNMP for FXOS* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *Configure Dynamic DNS* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *Remote Access VPN* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *FlexConfig Policies* in the *Cisco Secure Firewall Management Center Device Configuration Guide*

Alerting and Device-Level Security

- Create an email alert response, on page 567
- Configure email alerting for intrusion events, on page 579
- Create an SNMP alert response, on page 563
- *Configure NTP Time Synchronization for Threat Defense* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *About Configuring Syslog* in the *Cisco Secure Firewall Management Center Device Configuration Guide*

Protect your network

This section provides links to features that you can configure to enhance and protect your network security. Refer to each feature's documentation for detailed configuration guidance.

- *Access Control Policies* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.
- *Security Intelligence* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *Getting Started with Intrusion Policies* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *Tuning Intrusion Policies Using Rules* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *Custom Intrusion Rules* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- Intrusion rule updates, on page 260
- *Transport and Network Layer Preprocessors* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *Specific Threat Detection* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *Application Layer Preprocessors* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- Audit and Syslog, on page 423
- Intrusion Events, on page 787
- Event Search, on page 707
- Workflows, on page 667
- *Device Management* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- Login banners, on page 86
- Updates, on page 255

Enable security certifications compliance

Enable security certifications compliance to meet regulatory requirements for Common Criteria (CC) or Unified Capabilities Approved Products List (UCAPL) certification.

This configuration applies to either a Secure Firewall Management Center or managed device:

- For the Secure Firewall Management Center, this configuration is part of the system configuration.
- For a managed device, you apply this configuration from the Firewall Management Center as part of a platform settings policy.

In either case, the configuration does not take effect until you save your system configuration changes or deploy the shared platform settings policy.

Caution After you enable this setting, you cannot disable it. To disable CC or UCAPL mode, you must reimage the appliance.

Before you begin

- We recommend registering all devices that will be part of your deployment to the Firewall Management Center before enabling security certifications compliance on any appliances.
- Secure Firewall Threat Defense devices cannot use an evaluation license; your Smart Software Manager account must be enabled for export-controlled features.
- Secure Firewall Threat Defense devices must be deployed in routed mode.
- You must have administrative privileges to perform this task.

Procedure

Step 1 Depending on whether you are configuring a Firewall Management Center or a managed device:

- Firewall Management Center: Choose **System** (⚙) > **Configuration**.
- Firewall Threat Defense device: Choose **Devices** > **Platform Settings** and create or edit a Secure Firewall Threat Defense policy.

Step 2 Click **UCAPL/CC Compliance**.

Note
Appliances reboot when you enable UCAPL or CC compliance. The Firewall Management Center reboots when you save the system configuration; managed devices reboot when you deploy configuration changes.

Step 3 To *permanently* enable security certifications compliance on the appliance, you have two choices:

- To enable security certifications compliance in Common Criteria mode, choose **CC** from the drop-down list.
- To enable security certifications compliance in Unified Capabilities Approved Products List mode, choose **UCAPL** from the drop-down list.

Step 4 Click **Save**.

What to do next

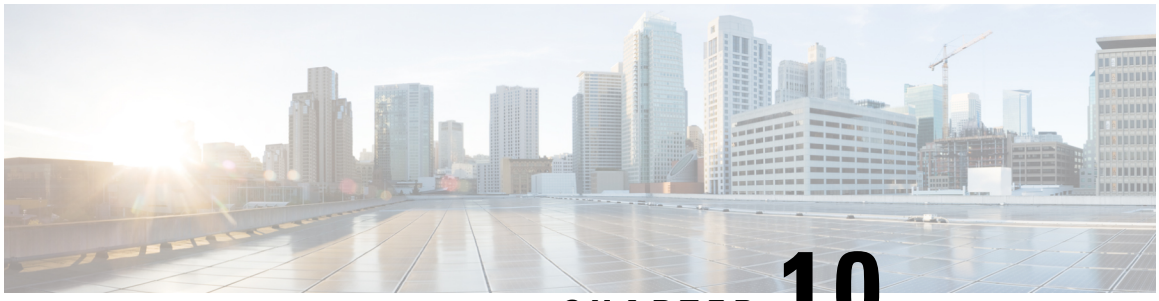
- Establish additional configuration changes as described in the guidelines for this product provided by the certifying entity.
- Deploy configuration changes; see the *Cisco Secure Firewall Management Center Device Configuration Guide*.



PART III

Health and Monitoring

- Dashboards, on page 357
- Health, on page 379
- Audit and Syslog, on page 423
- Statistics, on page 433
- Troubleshooting, on page 443



CHAPTER 10

Dashboards

The following topics describe how to use dashboards:

- Dashboards, on page 357
- Best practices for dashboard usage, on page 358
- Dashboard widgets, on page 358
- Widget availability by appliance, user role, and domain, on page 359
- Predefined dashboard widgets, on page 360
- Manage dashboards, on page 371

Dashboards

A dashboard is a customizable monitoring interface that provides at-a-glance views of the current system status, including data about events collected and generated by the firewall deployment. This interface also shows information on the status and overall health of the appliances in your deployment, and

The information that the dashboard provides depends on how you license, configure, and deploy the system.

Dashboard components and features

Firewall Management Center includes several predefined dashboards that you can use as-is or modify to suit your needs. By default, if your user role has dashboard access, your home page is the predefined **Summary Dashboard**. You can configure a different default home page or change the default dashboard to another predefined or custom dashboard.

Dashboard components include:

- Tabs: Each dashboard uses tabs to organize and display widgets grouped by related information. Tabs help you navigate different aspects of system monitoring within the same dashboard.
- Widgets: Small, self-contained components within tabs that provide insights into specific system aspects. For example, the predefined **Appliance Information** widget displays the appliance name, model, and currently running software version. Widgets display data constrained by the dashboard's time range, which you can adjust from as short as the last hour to as long as the last year.

User role requirements for accessing Dashboards

Access to dashboards depends on your user role. Roles with dashboard access include:

- **Administrator**
- **Maintenance User**
- **Security Analyst**
- **Security Analyst (Read Only)**
- Custom roles with the **Dashboards** permission

If your role does not have dashboard access, your default home page will be relevant to your role. For example, a **Discovery Admin** sees the **Network Discovery** page instead of a dashboard.

Best practices for dashboard usage

Consider the following best practices and guidelines for accessing and using dashboards in Firewall Management Center.

- Use dashboards to monitor your network with customizable views and detailed information. For a quick, visual summary of your monitored network, use **Context Explorer**.
- Create custom dashboards based on predefined ones. These custom dashboards can be shared with others, or kept private. When creating or modifying dashboards, use predefined dashboards as a base to maintain consistency, and ease of use.
- If you have **Administrator** access, you can view or edit private dashboards created by other users.
- Some drill-down pages and event tables include a dashboard toolbar link that opens a relevant predefined dashboard. Deleting a predefined dashboard or tab disables these toolbar links.
- In multidomain environments, dashboards from ancestor domains are not accessible. Create copies of higher-level dashboards if needed.
- To view the correlated device metrics on the dashboard, ensure that you have enabled the REST API (**System** (⚙) > **Configuration** > **REST API Preferences**).

Dashboard widgets

Widgets are modular dashboard components that display data such as events, operations, or general information, organized in a three-column layout.

The dashboard widgets that you can view depend on:

- the type of appliance you are using,
- your user role, and
- your current domain in a multidomain deployment.

In addition, each dashboard has a set of preferences that determines its behavior.

You can minimize and maximize widgets, add and remove widgets from tabs, as well as rearrange the widgets on a tab.

Note For widgets that display event counts over a time range, the total number of events may not reflect the number of events for which detailed data is available in the tables on pages under the **Analysis** menu. This occurs because Firewall Management Center sometimes prunes older event details to manage disk space usage. To minimize the occurrence of event detail pruning, you can fine-tune event logging to log only those events most important to your deployment.

Widget availability by appliance, user role, and domain

The dashboard widgets that you can view depend on the type of appliance you are using, your user role, and your current domain (in a multidomain deployment).

Under these conditions, a widget becomes unavailable for you:

- Widget is invalid because you are using the wrong type of appliance.
- You are not authorized to view the widget because your user account does not have the necessary privileges.

  For example, the **Appliance Status** widget is available only to users with **Administrator**, **Maintenance User**, **Security Analyst**, or **Security Analyst (Read Only)** account privileges.

Although you cannot add an unauthorized or invalid widget to a dashboard, an imported dashboard may contain unauthorized or invalid widgets. For example, such widgets can be present if the imported dashboard:

- was created by a user with different access privileges, or
- belongs to an ancestor domain.

Unavailable widgets are disabled and display error messages that indicate why you cannot view them. Individual widgets also display error messages when those widgets have timed out or are otherwise experiencing problems.

Note You can delete or minimize unauthorized and invalid widgets, as well as widgets that display no data, keeping in mind that modifying a widget on a shared dashboard modifies it for all users of the appliance.

In a multidomain deployment, if you do not see a widget that you expect to see, switch to the Global domain. Refer to Switch domains on Secure Firewall Management Center, on page 22.

Dashboard widget availability by user role

This table lists the user account privileges required to view each widget. Only user accounts with Administrator, Maintenance User, Security Analyst, or Security Analyst (Read Only) access can use dashboards.

Users with custom roles may have access to any combination of widgets, or none at all, as their user roles permit.

Table 18: User roles and dashboard widget availability

Widget	Administrator	Maintenance User	Security Analyst	Security Analyst (Read Only)
Appliance Information	yes	yes	yes	yes
Appliance Status	yes	yes	yes	no
Correlation Events	yes	no	yes	yes
Current Interface Status	yes	yes	yes	yes
Current Sessions	yes	no	no	no
Custom Analysis	yes	no	yes	yes
Disk Usage	yes	yes	yes	yes
Interface Traffic	yes	yes	yes	yes
Intrusion Events	yes	no	yes	yes
Network Compliance	yes	no	yes	yes
Product Licensing	yes	yes	no	no
Product Updates	yes	yes	no	no
RSS Feed	yes	yes	yes	yes
System Load	yes	yes	yes	yes
System Time	yes	yes	yes	yes
Allow List Events	yes	no	yes	yes

Predefined dashboard widgets

Predefined dashboard widgets are system-delivered, ready-to-use widgets that

- provide at-a-glance views of current system status and operational data
- display information about events that the system collects and generates
- display the status and overall health of the appliances in your deployment, and
- can be further customized through preferences that control display options and update frequency.

Widget categories

Predefined widgets are grouped into these categories based on the type of information they present:

- **Analysis & Reporting** widgets display data about the events collected and generated by the system.

- **Miscellaneous** widgets display neither event data nor operations data. Currently, the only widget in this category displays an RSS feed.
- **Operations** widgets display information about the status and overall health of the system.

Some widgets may not be visible to you, depending on appliance type, your user role, or domain restrictions.

The Allow list events widget

The **Allow List Events** widget shows the average events per second by priority, over the dashboard time range.

You can configure the widget to display allow list events of different priorities by modifying the widget preferences. In the widget preferences, you can:

- Check one or more **Priorities** check boxes to display separate graphs for events of specific priorities, including events that do not have a priority.
- Check the **Show All** check box to display an additional graph for all allow list events, regardless of priority.
- Choose **Linear** (incremental) or **Logarithmic** (factor of ten) scale from the **Vertical Scale** drop-down list.
- Specify how often the widget updates.

You can click a graph to view allow list events of a specific priority, or click the **All** graph to view all allow list events. In either case, the events are constrained by the dashboard time range; accessing allow list events via the dashboard changes the events (or global) time window for the Firewall Management Center.

The appliance information widget

The Appliance Information widget provides a snapshot of the appliance. It appears by default on the Status tabs of the **Detailed Dashboard** and the **Summary Dashboard**.

The widget provides:

- The name, IPv4 address, IPv6 address, and model of the appliance.
- The versions of the system software, operating system, Snort, rule update, rule pack, module pack, vulnerability database (VDB), and geolocation update installed on the appliances with dashboards, except for Firewall Management Center Virtual.
- For managed appliances, the name and status of the communications link with the managing appliance.
- For Firewall Management Centers in a high availability pair, the name, model, and system software and operating system versions of the peer Firewall Management Center, as well as how recently the Firewall Management Centers made contact.

You can configure the widget to display more or less information by modifying the widget preferences to display a simple or an advanced view; the preferences also control how often the widget updates.

Appliance status widget

The Appliance Status widget indicates the health of the appliance and of any appliances it is managing. Note that because the Firewall Management Center does not automatically apply a health policy to managed devices, you must manually apply a health policy to devices or their status appears as **Disabled**. This widget appears by default on the **Status** tabs of the **Detailed Dashboard** and the **Summary Dashboard**.

You can configure the widget to display appliance status as a pie chart or in a table by modifying the widget preferences.

The preferences also control how often the widget updates.

Click a section on the pie chart or one of the numbers on the appliance status table to go to the health monitoring page and view the compiled health status of the appliance and of any appliances it is managing.

The correlation events widget

The Correlation Events widget shows the average number of correlation events per second, by priority, over the dashboard time range. It appears by default on the Correlation tab of the Detailed Dashboard.

You can configure the widget to display correlation events of different priorities by modifying the widget preferences, as well as to choose a linear (incremental) or logarithmic (factor of ten) scale.

Check one or more **Priorities** check boxes to display separate graphs for events of specific priorities, including events that do not have a priority. Choose **Show All** to display an additional graph for all correlation events, regardless of priority. The preferences also control how often the widget updates.

Click a graph to view correlation events of a specific priority, or click the **All** graph to view all correlation events. In either case, the events are constrained by the dashboard time range; accessing correlation events via the dashboard changes the events (or global) time window for the appliance.

Current interface status widget

The Current Interface Status widget shows the status of all interfaces on the appliance, enabled or unused. On the Firewall Management Center, you can display the management (`eth0`, `eth1`, and so on) interfaces. On a managed device, you can choose to show only sensing (`s1p1` and so on) interfaces or both management and sensing interfaces. Interfaces are grouped by type: management, inline, passive, switched, routed, and unused.

For each interface, the widget provides:

- the name of the interface
- the link state of the interface
- the link mode (for example, 100Mb full duplex, or 10Mb half duplex) of the interface
- the type of interface, that is, copper or fiber
- the amount of data received (Rx) and transmitted (Tx) by the interface

The color of the ball representing link state indicates the current status, as follows:

- green: link is up and at full speed
- yellow: link is up but not at full speed
- red: link is not up

- gray: link is administratively disabled
- blue: link state information is not available (for example, ASA)

The widget preferences control how often the widget updates.

Current sessions widget

The **Current Sessions** widget shows which users are currently logged into the appliance, the IP address associated with the machine where the session originated, and the last time each user accessed a page on the appliance (based on the local time for the appliance). The user that represents you, that is, the user currently viewing the widget, is marked with a **User icon** and rendered in bold type. Sessions are pruned from this widget's data within one hour of logoff or inactivity. This widget appears by default on the Status tabs of the **Detailed Dashboard** and the **Summary Dashboard**.

On the **Current Sessions** widget, you can:

- Click any user name to manage user accounts on the **User Management** page.
- Click the **Host icon** or **Compromised Host icon** next to any IP address to view the host profile for the associated machine.
- Click any IP address or access time to view the audit log constrained by that IP address and by the time that the user associated with that IP address logged on to the web interface.

The widget preferences control how often the widget updates.

The custom analysis widget

The **Custom Analysis** widget is a highly customizable widget that allows you to display detailed information on the events collected and generated by the system.

The widget is delivered with multiple presets that provide quick access to information about your deployment. The predefined dashboards make extensive use of these presets. You can use these presets or create a custom configuration. At a minimum, a custom configuration specifies the data you are interested in (table and field), and an aggregation method for that data. You can also set other display-related preferences, including whether you want to show events as relative occurrences (bar graph) or over time (line graph).

The widget displays the last time it updated, based on local time. The widget updates with a frequency that depends on the dashboard time range. For example, if you set the dashboard time range to an hour, the widget updates every five minutes. On the other hand, if you set the dashboard time range to a year, the widget updates once a week. To determine when the dashboard will update next, hover your pointer over the **Last updated** notice in the bottom left corner of the widget.



Note A red-shaded Custom Analysis widget indicates that its use is harming system performance. If the widget continues to stay red over time, remove the widget. You can also disable all Custom Analysis widgets from the Dashboard settings in your system configuration (**System** (⚙) > **Configuration** > **Dashboard**)

Displaying relative occurrences of events (bar graphs)

For bar graphs in the Custom Analysis widget, the colored bars in the widget background show the relative number of occurrences of each event. Read the bars from right to left.

The **Direction icon** indicates and controls the sort order of the display. A downward-pointing icon indicates descending order; an upward-pointing icon indicates ascending order. To change the sort order, click the icon.

Next to each event, the widget can display one of three icons to indicate any changes from the most recent results:

- The new event icon **Add** (+) signifies that the event is new to the results.
- The **Up Arrow icon** indicates that the event has moved up in the standings since the last time the widget updated. A number indicating how many places the event has moved up appears next to the icon.
- The **Down Arrow icon** indicates that the event has moved down in the standings since the last time the widget updated. A number indicating how many places the event has moved down appears next to the icon.

Displaying events over time (line graphs)

If you want information on events or other collected data over time, you can configure the Custom Analysis widget to display a line graph, such as one that displays the total number of intrusion events generated in your deployment over time.

Limitations to the custom analysis widget

A Custom Analysis widget may indicate that you are unauthorized to view the data that is configured to display. For example, Maintenance Users are not authorized to view discovery events. As another example, the widget does not display information related to unlicensed features. However, you (and any other users who share the dashboard) can modify the widget preferences to display data that you can see, or even delete the widget. If you want to make sure that this does not happen, save the dashboard as private.

When viewing user data, the system displays only authoritative users.

When viewing URL category information, the system does not display uncategorized URLs.

When viewing intrusion events aggregated by **Count**, the count includes reviewed events for intrusion events; if you view the count in tables on pages under the Analysis menus, the count will not include reviewed events.



Note In a multidomain deployment, the system builds a separate network map for each leaf domain. As a result, a leaf domain can contain an IP address that is unique within its network, but identical to an IP address in another leaf domain. When you view Custom Analysis widgets in an ancestor domain, multiple instances of that repeated IP address can be displayed. At first glance, they might appear to be duplicate entries. However, if you drill down to the host profile information for each IP address, the system shows that they belong to different leaf domains.

How to create dashboard widgets for a device

Any widgets that show events from devices can be configured to use a filter that limits the display of events for a given device or a set of devices.

1. Create and save a search: Go to **Analysis** > **Search** and enter the search parameters to match the specific device names.

Note You must provide exact text match as there is no drop-down listing the deployed device names.

2. Go to **Overview** > **Dashboards heading** > **Dashboard**, then click **Add Widgets** to create a **Custom Analysis** widget.
3. Return to **Overview** > **Dashboards heading** > **Dashboard** and modify the new widget to customize with the scope of search.

Example: configuration of custom analysis widget

You can configure the Custom Analysis widget to display a list of recent intrusion events by configuring the widget to display data from the **Intrusion Events** table. Choosing the **Classification** field and aggregating this data by **Count** displays the number of events that were generated for each type.

On the other hand, aggregating by **Unique Events** displays the number of unique intrusion events of each type (for example, how many detections of network trojans, potential violations of corporate policy, attempted denial-of-service attacks, and so on).

You can further customize the widget using a saved search, either one of the predefined searches delivered with your appliance or a custom search that you created. For example, constraining the first example (intrusion events using the **Classification** field, aggregated by **Count**) using the **Dropped Events** search displays the number of intrusion events that were dropped for each type.

Related Topics

Modify dashboard time settings, on page 376

Custom analysis widget preferences

This table describes the preferences you can set in the Custom Analysis widget.

Different preferences appear depending on how you configure the widget. For example, a different set of preferences appears if you configure the widget to show relative occurrences of events (a bar graph) vs a graph over time (a line graph). Some preferences, such as **Filter**, only appear if you choose a specific table from which to display data.

Table 19: Custom analysis widget preferences

Preference	Details
Title	If you do not specify a title for the widget, the system uses the configured event type as the title.
Preset	Custom Analysis presets provide quick access to information about your deployment. The predefined dashboards make extensive use of these presets. You can use these presets or you can create a custom configuration.
Table (required)	The table of events or assets that contains the data the widget displays.
Field (required)	The specific field of the event type you want to display. To show data over time (line graphs), choose **Time**. To show relative occurrences of events (bar graphs), choose another option.

Preference	Details
Aggregate (required)	The aggregation method configures how the widget groups the data it displays. For most event types, the default option is **Count**.
Filter	You can use application filters to constrain data from the Application Statistics and Intrusion Event Statistics by Application tables.
Search	You can use a saved search to constrain the data that the widget displays. You do not have to specify a search, although some presets use predefined searches. Only you can access searches that you have saved as private. If you configure the widget on a shared dashboard and constrain its events using a private search, the widget resets to not using the search when another user logs in. This affects your view of the widget as well. If you want to make sure that this does not happen, save the dashboard as private. Only fields that constrain connection summaries can constrain Custom Analysis dashboard widgets based on connection events. Invalid saved searches are dimmed. If you constrain a Custom Analysis widget using a saved search, then edit the search, the widget does not reflect your changes until the next time it updates.
Show	Choose whether you want to display the most (**Top**) or the least (**Bottom**) frequently occurring events.
Results	Choose the number of result rows to display.
Show Movers	Choose whether you want to display the icons that indicate changes from the most recent results.
Time Zone	Choose the time zone you want to use to display results.
Color	You can change the color of the bars in the widget's bar graph.

Related Topics

Configure widget preferences, on page 374

View associated events from the custom analysis widget

From a Custom Analysis widget, you can invoke an event view (workflow) that provides detailed information about the events displayed in the widget. The events appear in the default workflow for that event type, constrained by the dashboard time range. This also changes the appropriate time window on the Firewall Management Center, depending on how many time windows you configured and on the event type.

For example:

- If you configure multiple time windows, then access health events from a Custom Analysis widget, the events appear in the default health events workflow, and the health monitoring time window changes to the dashboard time range.
- If you configure a single time window and then access any type of event from the Custom Analysis widget, the events appear in the default workflow for that event type, and the global time window changes to the dashboard time range.

Procedure

Step 1 View the dashboard you want to modify. For more information on viewing dashboard, refer to View dashboards, on page 378.

Step 2 Choose one of the following options:

- On any custom analysis widget, click **View** ( ) in the lower right corner of the widget to view all associated events, constrained by the widget preferences.
- On a custom analysis widget showing relative occurrences of events (bar graph), click any event to view associated events constrained by the widget preferences, as well as by that event.

Disk usage widget categories and features

The disk usage widget displays the percentage of space used on the hard drive, organized by disk usage categories. It also shows the usage and capacity of each partition of the appliance's hard drive. The disk usage widget displays the same information for the malware storage pack if installed in the device, or if the Firewall Management Center manages a device containing a malware storage pack. This widget appears by default on the Status tabs of the Default Dashboard and the Summary Dashboard.

The **By Category** stacked bar displays each disk usage category as a proportion of the total available disk space used. The table describes the available categories.

Table 20: Disk usage categories

Disk Usage Category	Description
Events	All events logged by the system
Files	All files stored by the system
Backups	All backup files
Updates	All files related to updates, such as rule updates and system updates
Other	System troubleshooting files and other miscellaneous files
Free	Free space remaining on the appliance

Hover your pointer over a disk usage category in the stacked bar to view the percentage of available disk space used by that category, the actual storage space on the disk, and the total disk space available for that category. If your device has a malware storage pack, the Files category shows the space on the malware storage pack.

You can configure the widget to display only the **By Category** stacked bar, or you can show the stacked bar plus the admin (`/`), `/Volume`, and `/boot` partition usage, as well as the `/var/storage` partition if the malware storage pack is installed, by modifying the widget preferences.

Use widget preferences to control how often the widget updates and whether it displays current disk usage or collected disk usage statistics over the dashboard time range.

The interface traffic widget

The Interface Traffic widget shows the rate of traffic received (Rx) and transmitted (Tx) on the appliance's management interface. The widget does not appear by default on any of the predefined dashboards.

The widget preferences control how often the widget updates.

This widget displays the input and output rates from Snort. Note that the traffic rates in the **Interface** dashboard (**System** (⚙) > **Health** > **Monitor** page) may differ, as those values are collected from Lina.

Devices with Malware licenses enabled periodically attempt to connect to the AMP cloud even if you have not configured dynamic analysis. Because of this, these devices show transmitted traffic; this is expected behavior.

The intrusion events widget

The **Intrusion Events** widget shows the intrusion events that occurred over the dashboard time range, organized by priority. This includes statistics on intrusion events with dropped packets and different impacts. This widget appears by default on the **Intrusion Events** tab of the **Summary Dashboard**.

Widget preferences

In the widget preferences, you can choose:

- **Event Flags** to display separate graphs for events with dropped packets, would have dropped packets, or specific impacts. Choose **All** to display an additional graph for all intrusion events, regardless of impact or rule state.

  For explanations of the icons, see Intrusion Events, on page 787. The arrow (if any) that appears above the impact level numbers describes the inline result and is defined as follows:

  Table 21: Inline Result Field Contents in Workflow and Table Views

This Icon	Indicates
	The system dropped the packet that triggered the rule.
	IPS would have dropped the packet if you enabled the **Drop when Inline** intrusion policy option (in an inline deployment), or if a Drop and Generate rule generated the event while the system was pruning.
	IPS may have transmitted or delivered the packet to the destination, but the connection that contained this packet is now blocked.
No icon (blank)	The triggered rule was not set to Drop and Generate Events

  In a passive deployment, the system does not drop packets, including when an inline interface is in tap mode, regardless of the rule state or the inline drop behavior of the intrusion policy.

- **Show** to specify **Average Events Per Second** (EPS) or **Total Events**.
- **Vertical Scale** to specify **Linear** (incremental) or **Logarithmic** (factor of ten) scale.
- How often the widget updates.

Widget actions

On the widget, you can:

- Click a graph corresponding to dropped packets, to would have dropped packets, or to a specific impact to view intrusion events of that type.
- Click the graph corresponding to dropped events to view dropped events.
- Click the graph corresponding to would have dropped events to view would have dropped events.
- Click the **All** graph to view all intrusion events.

The resulting event view is constrained by the dashboard time range; accessing intrusion events via the dashboard changes the events (or global) time window for the appliance. Note that packets in a passive deployment are not dropped, regardless of intrusion rule state or the inline drop behavior of the intrusion policy.

Network compliance widget

The Network Compliance widget summarizes your hosts' compliance with the allow lists you configured. By default, the widget displays a pie chart that shows the number of hosts that are compliant, non-compliant, and that have not been evaluated, for all compliance allow lists in active correlation policies. This widget appears by default on the Correlation tab of the Detailed Dashboard.

You can configure the widget to display network compliance either for all allow lists or for a specific allow list by modifying the widget preferences.

If you choose to display network compliance for all allow lists, the widget considers a host to be non-compliant if it is not compliant with any allow list in an active correlation policy.

The preferences control how often the widget updates. You can check the **Show Not Evaluated** box to hide events which have not been evaluated.

You can also use the widget preferences to specify which of these three different styles you want to use to display network compliance:

- The **Network Compliance** style (the default) displays a pie chart that shows the number of hosts that are compliant, non-compliant, and that have not been evaluated. You can click the pie chart to view the host violation count, which lists the hosts that violate at least one allow list.
- The **Network Compliance over Time (%)** style displays a stacked area graph showing the relative proportion of hosts that are compliant, non-compliant, and that have not yet been evaluated, over the dashboard time range.
- The **Network Compliance over Time** style displays a line graph that shows the number of hosts that are compliant, non-compliant, and that have not yet been evaluated, over the dashboard time range.

The product licensing widget

The Product Licensing widget shows the device and feature licenses currently installed on the Firewall Management Center. It also indicates the number of items licensed and the number of remaining licensed items allowed. It does not appear by default on any of the predefined dashboards.

The top section of the widget displays all device and feature licenses installed on the Firewall Management Center, including temporary licenses, while the Expiring Licenses section displays only temporary and expired licenses.

The bars in the widget background show the percentage of each type of license that is being used; you should read the bars from right to left. Expired licenses are marked with a strikethrough.

You can configure the widget to display either the features that are currently licensed, or all the features that you can license, by modifying the widget preferences. The preferences also control how often the widget updates.

You can click any of the license types to go to the License page of the local configuration and add or delete feature licenses.

Product updates widget

The Product Updates widget provides you with a summary of the software currently installed on the appliance as well as information on updates that you have downloaded, but not yet installed. This widget appears by default on the Status tabs of the **Detailed Dashboard** and the **Summary Dashboard**.

Because the widget uses scheduled tasks to determine the latest version, it displays `Unknown` until you configure a scheduled task to download, push or install updates.

You can configure the widget to hide the latest versions by modifying the widget preferences. The preferences also control how often the widget updates.

The widget also provides you with links to pages where you can update the software. You can:

- Manually update an appliance by clicking the current version.
- Create a scheduled task to download an update by clicking the latest version.

RSS feed widget

The RSS Feed widget adds an RSS feed to a dashboard. By default, the widget shows a feed of Cisco security news. It appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard.

You can also configure the widget to display a preconfigured feed of company news, the Snort.org blog, or the Cisco Threat Research blog, or you can create a custom connection to any other RSS feed by specifying its URL in the widget preferences. The Firewall Management Center can display encrypted RSS feeds only if they use trusted server certificates signed by a certificate authority (CA) that the Firewall Management Center recognizes. If you configure the RSS Feed widget to display an encrypted RSS feed that uses a CA the Firewall Management Center does not recognize, or that uses a self-signed certificate, the verification fails and the widget does not display the feed.

Feeds update every 24 hours (although you can manually update the feed), and the widget displays the last time the feed was updated based on the local time of the appliance. Keep in mind that the appliance must have access to the web site (for the two preconfigured feeds) or to any custom feed you configure.

When you configure the widget, you can also choose how many stories from the feed you want to show in the widget, as well as whether you want to show descriptions of the stories along with the headlines; keep in mind that not all RSS feeds use descriptions.

On the RSS Feed widget, you can:

- Click one of the stories in the feed to view the story.

- Click the **more** link to go to the feed's web site.
- Click **Refresh** ( ) to manually update the feed.

The system load widget

The System Load widget shows the CPU usage (for each CPU), memory (RAM) usage, and system load (also called the load average, measured by the number of processes waiting to execute) on the appliance, both currently and over the dashboard time range. It appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard.

You can configure the widget to show or hide the load average by modifying the widget preferences. The preferences also control how often the widget updates.

System time widget

The **System Time** widget shows the local system time, uptime, and boot time for the appliance. It appears by default on the Status tabs of the **Detailed Dashboard** and the **Summary Dashboard**.

You can configure the widget to hide the boot time by modifying the widget preferences. The preferences also control how often the widget synchronizes with the appliance's clock.

Manage dashboards

Manage dashboards to customize your monitoring interface by creating, deleting, and modifying dashboards and their widgets to meet your specific operational requirements.

Procedure

Step 1 Choose **Overview** > **Dashboards** > **Management**, and then manage your dashboards:

- Create Dashboards: Create a custom dashboard; refer to Create custom dashboards, on page 374.
- Delete Dashboards: To delete a dashboard, click **Delete** ( ) next to the dashboard you want to delete. If you delete your default dashboard, you must define a new default or the appliance prompts you to choose a dashboard every time you attempt to view a dashboard.
- Edit Options: Edit custom dashboard options; refer to Edit dashboard options, on page 375.
- Modify Time Constraints: Modify the time display or pause/unpause the dashboard as described in Modify dashboard time settings, on page 376.

Step 2 To modify a single dashboard, choose **Overview** > **Dashboards heading** > **Dashboard**, and then go to the dashboard you want to modify.

Step 3 Add, delete, or rename dashboard tabs.

- For more information about adding a dashboard tab, refer to Add a dashboard, on page 372.
- For more information about renaming a dashboard tab, refer to Rename a dashboard, on page 377.
- To delete a dashboard tab, click **Close** ( ).

Note
You cannot change the order of dashboards.

Step 4 Manage dashboard widgets:

- Add Widgets: Add widgets to a dashboard; refer to Adding widgets to a dashboard, on page 372.
- Configure Preferences: Configure widget preferences; refer to Configure widget preferences, on page 374.
- Customize Display: Customize the widget display; refer to Customize the widget display, on page 373.
- View Events: View associated events from the Custom Analysis Widget; refer to View associated events from the custom analysis widget, on page 366.

Note
Every configuration of the Custom Analysis widget in the Cisco predefined dashboards corresponds to a system preset for that widget. If you change or delete one of these widgets, you can restore it by creating a new Custom Analysis widget based on the appropriate preset.

Add a dashboard

Add a dashboard to customize your security monitoring interface and organize relevant data displays for your specific operational needs.

Use this procedure when you need to create a new dashboard to group related monitoring widgets and data visualizations for your security environment.

Procedure

Step 1 View the dashboard you want to modify; see View dashboards, on page 378.

Step 2 Click **Add** (+).

Step 3 Enter a name.

Step 4 Click **OK**.

The new dashboard is created and becomes available in your dashboard list for configuration and use.

Adding widgets to a dashboard

Add widgets to dashboard tabs to customize your monitoring and analysis capabilities. You can add up to 15 widgets per tab to create personalized views of system data and metrics.

Each tab can display one or more widgets in a three-column layout. When adding a widget to a dashboard, you choose the tab to which you want to add the widget. The system automatically adds it to the column with the fewest widgets. If all columns have an equal number of widgets, Firewall Management Center adds the new widget to the left-most column. You can add a maximum of 15 widgets to a dashboard tab.

Tip After you add widgets, you can move them to any location on the tab. You cannot, however, move widgets from tab to tab.

The dashboard widgets you see depend on your appliance type, your user role, and your current domain (in a multidomain deployment). If you have fewer permissions than the dashboard creator, you may not be able to use all of the widgets on the dashboard. Firewall Management Center displays unauthorized widgets as disabled.

Procedure

Step 1 View the dashboard where you want to add a widget. For instructions, refer to View dashboards, on page 378.

Step 2 Click the tab where you want to add the widget.

Step 3 Click **Add Widgets**. You can view the widgets in each category by clicking on the category name, or you can view all widgets by clicking **All Categories**.

Step 4 Click **Add** next to the widgets you want to add. The **Add Widgets** page shows how many widgets of each type are on the tab, including the one you want to add.

To add multiple widgets of the same type, click **Add** again.

Step 5 When you are finished adding widgets, click **Done** to return to the dashboard.

What to do next

- If you added a **Custom Analysis** widget, configure the widget preferences. For more information, refer to Configure widget preferences, on page 374.

Related Topics

Widget availability by appliance, user role, and domain, on page 359

Customize the widget display

Customize widget display to optimize dashboard layout and focus on relevant information. You can minimize and maximize widgets, as well as rearrange the widgets on a tab.

Procedure

Step 1 View a dashboard. For more information, refer to View dashboards, on page 378.

Step 2 Customize the widget display:

- To rearrange a widget on a tab, click the title bar of the widget you want to move, then drag it to its new location.

Note

You cannot move widgets from tab to tab. If you want a widget to appear on a different tab, you must delete it from the existing tab and add it to the new tab.

- To minimize or maximize a widget on the dashboard, click **Minimize** ( ) or **Maximize** ( ) in a widget's title bar.
- To delete a widget if you no longer want to view it on a tab, click **Close** ( ) in the title bar of the widget.

Configure widget preferences

Configure widget preferences to customize the behavior and functionality of individual widgets according to your specific requirements.

Procedure

Step 1 On the title bar of the widget whose preferences you want to change, click **Show Preferences** ( ).

Step 2 Make changes as needed.

Step 3 On the widget title bar, click **Hide Preferences** ( ) to hide the preferences section.

Create custom dashboards

Custom dashboards allow you to organize and display specific data visualizations that meet your monitoring requirements and provide tailored views of your network security information.

Create new dashboards to customize your monitoring environment. Alternatively, you can export dashboards from other appliances and import them to save time. You can then edit the imported dashboard to suit your needs.

Procedure

Step 1 Choose **Overview** > **Dashboards** > **Management**.

Step 2 Click **Create Dashboard**.

Step 3 Modify the custom dashboard options as described in Custom dashboard options, on page 374.

Step 4 Click **Save**.

Custom dashboard options

This table describes options you can use when creating or editing custom dashboards.

Table 22: Custom dashboard options

Option	Description
Copy Dashboard	When you create a custom dashboard, you can choose to base it on any existing dashboard, whether user-created or system-defined. This option makes a copy of the preexisting dashboard, which you can modify to suit your needs. Optionally, you can create a blank new dashboard by choosing **None**. This option is available only when you create a new dashboard. In a multidomain deployment, you can copy any non-private dashboards from ancestor domains.
Name	A unique name for the custom dashboard.
Description	A brief description of the custom dashboard.
Change Tabs Every	Specifies (in minutes) how often the dashboard should cycle through its tabs. Unless you pause the dashboard or your dashboard has only one tab, this setting advances your view to the next tab at the interval you specify. To disable tab cycling, enter `0` in the **Change Tabs Every** field.
Refresh Page Every	Determines how often the entire dashboard page automatically refreshes. Refreshing the entire dashboard allows you to see any preference or layout changes that were made to a shared dashboard by another user, or that you made to a private dashboard on another computer, since the last time the dashboard refreshed. A frequent refresh can be useful, for example, in a networks operations center (NOC) where a dashboard is displayed at all times. If you make changes to the dashboard at a local computer, the dashboard in the NOC automatically refreshes at the interval you specify, and no manual refresh is required. This refresh does not update the data, and you do not need to refresh the entire dashboard to see data updates; individual widgets update according to their preferences. This value must be greater than the **Change Tabs Every** setting. Unless you pause the dashboard, this setting will refresh the entire dashboard at the interval you specify. To disable the periodic page refresh, enter `0` in the **Refresh Page Every** field. **Note** This setting is separate from the update interval available on many individual widgets; although refreshing the dashboard page resets the update interval on individual widgets, widgets will update according to their individual preferences even if you disable the **Refresh Page Every** setting.
Save As Private	Determines whether the custom dashboard can be viewed and modified by all users of the appliance or is associated with your user account and reserved solely for your own use. Keep in mind that any user with dashboard access, regardless of role, can modify shared dashboards. If you want to make sure that only you can modify a particular dashboard, save it as private.

Edit dashboard options

Edit dashboard options to modify the configuration and settings of an existing dashboard to better meet your monitoring and analysis requirements.

Procedure

Step 1 View the dashboard you want to edit. For more information, refer to View dashboards, on page 378.

Step 2 Click **Edit** ( ).

Step 3 Change the options as described in Custom dashboard options, on page 374.

Step 4 Click **Save**.

Modify dashboard time settings

You can change the time range to reflect a period as short as the last hour (the default) or as long as the last year. When you change the time range, the widgets that can be constrained by time automatically update to reflect the new time range.

The maximum number of data points in any graph is 300, and the time setting determines how much time is summarized within each data point. This is the number of data points, and the time span covered, in the dashboards for each time range:

- 1 hour = 12 data points, 5 minutes each
- 6 hours = 72 data points, 5 minutes each
- 1 day = 288 data points, 5 minutes each
- 1 week = 300 data points, 33.6 minutes each
- 2 weeks = 300 data points, 67.2 minutes each
- 30 days = 300 data points, 144 minutes each
- 90 days = 300 data points, 432 minutes each
- 180 days = 300 data points, 864 minutes each
- 1 year = 300 data points, 1752 minutes each

Note that not all widgets can be constrained by time. For example, the dashboard time range has no effect on the Appliance Information widget, which provides information that includes the appliance name, model, and current version of the software. For more information about pre-defined tables that cannot be constrained by time, refer to Time constraints in dashboard tables, on page 377.

Keep in mind that for enterprise deployments of the Secure Firewall System, changing the time range to a long period may not be useful for widgets like the Custom Analysis widget, depending on how often newer events replace older events.

You can also pause a dashboard, which allows you to examine the data provided by the widgets without the display changing and interrupting your analysis. Pausing a dashboard has these effects:

- Individual widgets stop updating, regardless of any **Update Every** widget preference.
- Dashboard tabs stop cycling, regardless of the **Cycle Tabs Every** setting in the dashboard properties.
- Dashboard pages stop refreshing, regardless of the **Refresh Page Every** setting in the dashboard properties.

- Changing the time range has no effect.

When you are finished with your analysis, you can unpause the dashboard. Unpausing the dashboard causes all appropriate widgets on the page to update to reflect the current time range. In addition, dashboard tabs resume cycling and the dashboard page resumes refreshing according to the settings you specified in the dashboard properties.

If you experience connectivity problems or other issues that interrupt the flow of system information to the dashboard, the dashboard automatically pauses and an error notice appears until the problem is resolved.

Note Your session normally logs you out after 1 hour of inactivity (or another configured interval), regardless of whether the dashboard is paused. If you plan to passively monitor the dashboard for long periods of time, consider exempting some users from session timeout, or changing the system timeout settings.

Follow these steps to modify dashboard time settings:

Procedure

Step 1 View the dashboard where you want to add a widget; refer to View dashboards, on page 378.

Step 2 Optionally, to change the dashboard time range, choose a time range from the **Show the Last** drop-down list.

Step 3 Optionally, pause or unpause the dashboard on the time range control, using **Pause** ( ) or **Play** ( ).

Time constraints in dashboard tables

In Firewall Management Center, information about your network is collected and stored in structured data tables. When you view network information through a dashboard widget, Firewall Management Center retrieves the relevant data from these tables.

Tables can be classified into these categories based on ability to constrain by time:

- Tables that are event-based, such as **Connection Events** or **User Activity**, include a time column and can be constrained by event time, allowing you to filter data based on specific time ranges.
- Tables that are state-based, such as **Hosts**, do not have an associated event time and therefore cannot be constrained by time filters. These tables represent the current or recent state of the system rather than discrete events.

Consider this distinction while you are customizing your dashboard and analyzing network events. Ensure that you apply time-based filtering only when it is supported by the displayed data.

Rename a dashboard

Change the name of an existing dashboard to improve clarity or organizational structure.

Procedure

Step 1 View the dashboard you want to modify. For more information on viewing dashboard, refer to View dashboards, on page 378.

Step 2 Click the dashboard title you want to rename.

Step 3 Enter a new name.

Step 4 Click **OK**.

View dashboards

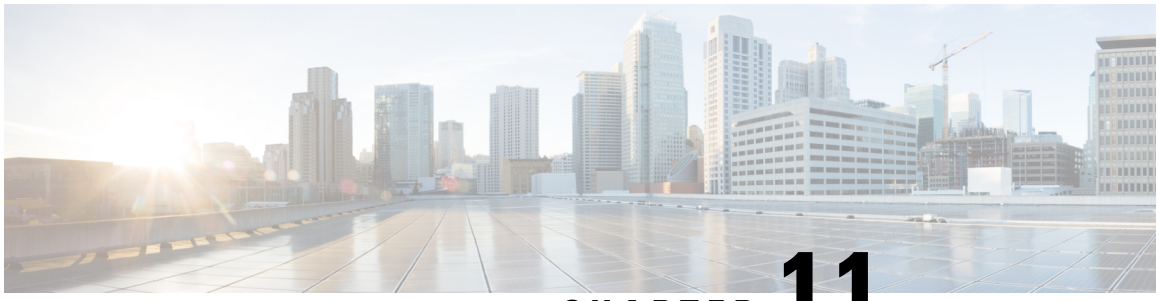
View dashboards to monitor and analyze data relevant to your appliance configuration and performance.

By default, the home page for your appliance displays the default dashboard. If you do not have a default dashboard defined, the home page shows the **Dashboard Management** page, where you can choose a dashboard to view.

Procedure

At any time, you can do one of the following:

- To view the default dashboard for your appliance, choose **Overview** > **Dashboards heading** > **Dashboard**.
- To view a specific dashboard, choose **Overview** > **Dashboards heading** > **Dashboard**, and choose the dashboard from the menu.
- To view all available dashboards, choose **Overview** > **Dashboards** > **Management**. You can then choose **View** ( ) next to an individual dashboard to view it.



CHAPTER 11

Health

The following topics describe how to use health monitoring:

- Requirements and prerequisites for health monitoring, on page 379
- Health monitoring, on page 379
- Health policies, on page 392
- Device exclusion in health monitoring, on page 396
- Health monitor alerts, on page 400
- About the health monitor, on page 402
- Health event views, on page 412
- History for health monitoring, on page 415

Requirements and prerequisites for health monitoring

This reference provides the requirements and prerequisites needed to implement health monitoring functionality.

Model support

Any

Supported domains

Any

User roles

Admin: Full permissions (create, view, edit, delete, and deploy) for health policies, exclusions, alerts, and related health settings.

Maintenance User and Security Analyst: View-only permissions for health monitoring data (dashboards, status, and events); no create, edit, delete, or deploy permissions.

Health monitoring

Health monitoring is a system capability that

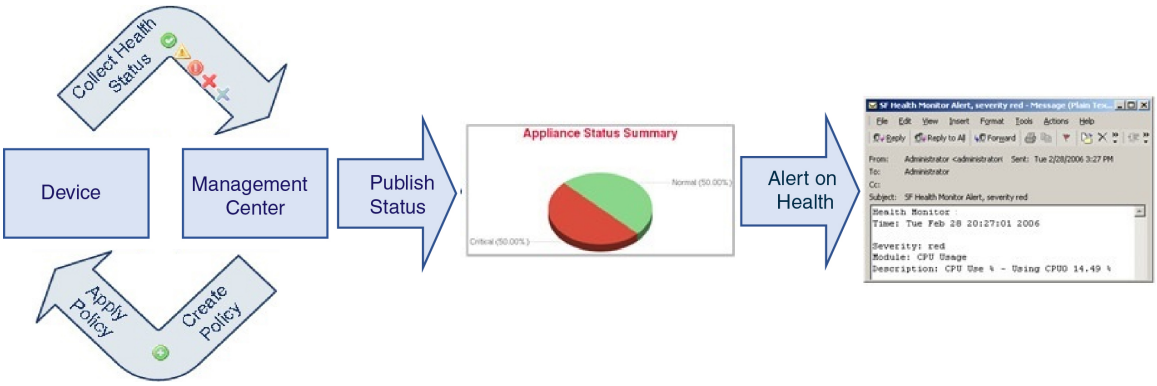
- tracks various health indicators to ensure hardware and software are working correctly,

- checks the status of critical functionality across your deployment, and
- enables configuration of test frequency for alerting and time series data collection.

Health monitoring capabilities

Health monitoring includes two key capabilities: health status monitoring through health policies and health modules, and time series data collection for metrics analysis.

With time series data collection, you can configure how frequently metric data is collected for the device and its health modules. The device monitor reports these metrics in several predefined health monitor dashboards by default. This metric data is collected for analysis. Some alerts are generated from metric values when they cross the thresholds specified in the health policy.



You can use the health monitor to create a collection of tests, referred to as a *health policy*, and apply the health policy to one or more appliances. The tests, referred to as *health modules*, are scripts that test for the criteria you specify. You can modify a health policy by enabling or disabling tests or by changing test settings, and you can delete health policies that you no longer need. You can also suppress messages from selected appliances by excluding them.

The health monitoring system run the tests in a health policy at the configured intervals. You can also run all tests, or a specific test, on demand. The health monitor collects health events based on the test conditions configured.

Note All appliances automatically report their hardware status via the Hardware Alarms health module. The Firewall Management Center also automatically reports status using the modules configured in the default health policy. Some health modules, such as the Appliance Heartbeat module, run on the Firewall Management Center and report the status of the Firewall Management Center's managed devices. For the health modules to provide managed device status, you must deploy all health policies to the device.

You can use the health monitor to access health status information for the entire system and for a particular appliance. In a multi-domain deployment, you can view the health status summary for a device in the domain where the device is located.

A hexagon-shaped widget and status tables on the **Health Status** page provide a visual summary of the status of all appliances on your network, including the Firewall Management Center. Individual appliance health monitors let you drill down into health details for a specific appliance.

Fully customizable event views allow you to quickly and easily analyze the health status events gathered by the health monitor. These event views allow you to search and view event data and to access other information that may be related to the events you are investigating. For example, if you want to see all the occurrences of CPU usage with a certain percentage, you can search for the CPU usage module and enter the percentage value.

You can also configure alerting in response to health events. A *health alert* is an association between a standard alert and a health status level. For example, if you want to make sure an appliance never fails due to hardware overload, you can set up an email alert. You can then create a health alert that triggers an email alert whenever CPU, disk, or memory usage reaches the Warning level you configure in the health policy applied to that appliance. You can set alerting thresholds to minimize the number of repeating alerts you receive.



Note The health monitoring can take 5–6 minutes from the occurrence of the health event to generate the health alert.

You can also generate troubleshooting files for an appliance if you are asked to do so by Support.

Only users with administrator user role privileges can access system health data.

High availability pair

In a Firewall Management Center high availability deployment, the active peer creates a health monitor page that uses REST APIs to show detailed metric-based information. The standby peer creates the health monitor page that shows the alert information and provide a visual summary of the status of all appliances on your network using pie charts and status tables. The standby peer does not display the metric-based information.

Health modules

Describes health modules that monitor system conditions and provide alerts or metrics for managing device health and performance.

Health modules, or *health tests,* test for the conditions that you specify in a health policy.

Health monitoring includes both health alerts or tests and metric or time series data collection. The two types of health modules are alert modules and metrics modules. Alerts modules (sometimes called *legacy* modules) monitor system infrastructure and report only the health status. When the conditions specified in the health policy for these monitored systems are met, these modules raise health alerts. Metrics modules (sometimes called *telegraf* modules) collect statistics (sometimes called *time series data*) that you can view on the health monitoring dashboard. Some alerts are also generated from metric values when they cross the thresholds specified in the health policy. You can create custom dashboards with your preferred health metrics, allowing you to monitor statistics or troubleshoot appliance health issues.

Table 23: Device Health Modules

Module	Type	Description
AMP Connection Status	Metrics	The module alerts if the device cannot connect to the AMP cloud or Cisco AMP Private Cloud after an initial successful connection, or if the private cloud cannot contact the public AMP cloud. Disabled by default.

Module	Type	Description
AMP Threat Grid Connectivity	Metrics	The module alerts if the device cannot connect to the AMP Threat Grid cloud after an initial successful connection.
ASP Drop	Metrics	Monitors the connections dropped by the data plane accelerated security path.
Automatic Application Bypass	Alert	Monitors bypassed detection applications.
Chassis Environment Status	Alert	Monitors chassis parameters such as fan speed and chassis temperature, and enables you to set a warning threshold and critical threshold for temperature. The **Critical Chassis Temperature (Celsius)** default value is `85`. The **Warning Chassis Temperature (Celsius)** default value is `75`.
Cluster/HA Failover Status	Alert	For threat defense clusters, alerts when a unit joins, leaves, or is elected primary.
Configuration Resource Utilization	Alert	Alerts if the size of your deployed configurations puts a device at risk of running out of memory. The alert shows you how much memory your configurations require, and by how much this exceeds the available memory. If this happens, reevaluate your configurations. You may be able to reduce the number or complexity of access control rules or intrusion policies.
Connection Statistics	Metrics	Monitors connection statistics and NAT translation counts. For standby Firewall Threat Defense devices in a high availability pair, this widget reflects only the statistics of connections replicated from the active device.
CPU Usage (per core)	Metrics	Alerts when CPU core use exceeds a configurable threshold.
Critical Process Statistics	Metrics	Monitors the state of critical processes, their resource consumption, and the restart counts.
CPU Usage Date Plane	Metrics	Alerts when data plane CPU use exceeds a configurable threshold.
Memory Usage Data Plane	Metrics	Alerts when data plane memory use exceeds a configurable threshold.
Deployed Configuration Statistics	Metrics	Monitors statistics about the deployed configuration, such as the number of ACEs and IPS rules.
Disk Status	Alert	Alerts if there is an issue with the hard disk or RAID controller. If this module alerts, contact Cisco TAC. This will prevent upgrade.

Module	Type	Description
Disk Usage	Metrics	This module compares disk usage on the appliance's hard drive to the limits configured for the module and alerts when usage exceeds the thresholds configured for the module. This module also alerts when the system excessively deletes files in monitored disk usage categories, or when disk usage excluding those categories reaches excessive levels, based on module thresholds. For more information, refer to Disk usage and drain of events health monitor alerts, on page 452 for information about troubleshooting scenarios for Disk Usage alerts. The Disk Usage module sends a health alert if the size of device configuration history files exceeds the allowed limit. For more information, refer to Disk Usage for Device Configuration History Files Health Monitoring Alert for information about troubleshooting scenarios for the disk usage alerts. This health alert is not supported on Secure Firewall Management Center Versions 7.2.0-7.2.5, 7.3.x, and 7.4.0. Use the Disk Usage health status module to monitor disk usage for the `/` and `/volume` partitions on the appliance and track draining frequency. Although the disk usage module lists the `/boot` partition as a monitored partition, the size of the partition is static so the module does not alert on the boot partition. **Attention** If you receive alerts for high unmanaged disk usage for the partition `/volume` although the usage is below the critical or warning threshold specified in the health policy, this could indicate that there are files which must be deleted manually from the system. Contact TAC if you receive these alerts.
File System Integrity Check	Alert	This module performs a file system integrity check and runs if the system has CC mode or UCAPL mode enabled, or if the system runs an image signed with a DEV key.
Firewall Threat Defense HA	Alert	Alerts if a threat defense high availability pair is split brain.
Firewall Threat Defense Platform Faults	Alert	Monitors Secure Firewall 1000 /2100 /3100 platform faults and generate health alerts for the faults. A platform fault represents a failure in the Firewall Threat Defense instance or an alarm threshold that has been raised. During the lifecycle of a platform fault, it can change from one state or severity to another. Each fault includes information about the operational state of the affected object at the time the fault was raised. If the fault is transitional and the failure is resolved, then the object transitions to a functional state. For more information, refer to *Cisco Firepower 1000/2100 FXOS Faults and Error Messages Guide*.
Flow Offload Statistics	Metrics	Monitors hardware flow offload.
Hardware Alarms	Alert	This module determines if hardware needs to be replaced on a physical managed device and alerts based on the hardware status. It also reports on the status of hardware-related daemons.
Inline Link Mismatch Alarms	Alert	Alerts if inline pair interfaces negotiate different speeds.

Module	Type	Description
Interface Status	Alert	Determines if the device currently collects traffic and alerts based on the traffic status of physical interfaces and aggregate interfaces. For physical interfaces, the information includes interface name, link state, and bandwidth. For aggregate interfaces, the information includes interface name, number of active links, and total aggregate bandwidth. For subinterfaces, this module reports health alerts in aggregate rather than generating separate alerts for each affected subinterface. If multiple subinterfaces on the same parent interface experience health issues, Firewall Management Center generates a single summarized alert indicating the number of affected subinterfaces for that parent interface. To view the specific status of individual subinterfaces, refer to the **Interfaces** tab. **Note** • This module also monitors the high availability standby device traffic flow. Though it is known that the standby device would not be receiving any traffic yet, the Firewall Management Center alerts that the interface is not receiving any traffic. The same alerting principle is applied when traffic is not received by some of the subinterfaces on a port channel. • This module displays the traffic rates according to the values from Lina. However, if you use the **show interface** CLI command to know the interface statistics of your device, the input and output rates in the CLI command result can be different from the traffic rates that appear in the **Interface** widget. The sampling intervals of Lina and the Firewall Management Center interface statistics are different. Due to the difference in sampling interval, throughput values in the Firewall Management Center GUI can be different from the throughput values appears in the device CLI result. • Note that traffic rates in the **Interface Traffic Rate** widget (**Overview** > **Dashboards heading** > **Dashboard** page) can be different as it displays the input and output rates from Snort.
Intrusion and File Event Rate	Alert	Alerts if intrusion events per second exceed a configurable threshold. We recommend a warning threshold of 1.5 times your average intrusion event rate, and a critical threshold of 2.5 times. For example, for an average event rate on network segment of 20 events per second, we recommend a warning value of 30 and a critical value of 50. The critical limit must be lower than 1000, and higher than the warning limit. Event rates for your devices are available on **System** (⚙) > **Monitoring** > **Statistics**. If the rate is zero, the Snort process may be down or the device may not be sending events.
Link State Propagation	Alert	For the ISA 3000, alerts when an interface in a inline set fails.

Module	Type	Description
Memory Usage	Alert	Alerts when memory use exceeds configurable thresholds. For appliances with more than 4 GB of memory, the preset alert thresholds are based on a formula that accounts for proportions of available memory likely to cause system problems. On >4 GB appliances, because the interval between Warning and Critical thresholds may be very narrow, it is recommended that you manually set the **Warning Threshold %** value to `50`. This will further ensure that you receive memory alerts for your appliance in time to address the issue. For more information, refer to Memory usage thresholds for health monitor alerts, on page 451 for additional information about how thresholds are calculated. Complex access control policies and rules can command significant resources and negatively affect performance.
Network Card Reset	Alert	Alerts when a network card restarts due to hardware failure.
NTP Statistics	Metrics	Monitors NTP synchronization status. Disabled by default.
Firewall Management Center Access Configuration Changes	Alert	Monitors configuration changes made on the Firewall Management Center directly using the **configure network management-data-interface** command. This module alerts when there is a conflict between the existing Firewall Management Center configuration and the out of band configuration changes made.
Process Status	Alert	Alerts when processes on the appliance exit or terminate outside of the process manager. If a process is deliberately exited outside of the process manager, the module status changes to Warning and the health event message indicates which process exited, until the module runs again and the process has restarted. If a process terminates abnormally or crashes outside of the process manager, the module status changes to Critical and the health event message indicates the terminated process, until the module runs again and the process has restarted.
Routing Statistics	Metrics	Monitors the current state of routing table.
Snort 3 Statistics	Metrics	Collects Snort 3 statistics for events, flows, and packets.
CPU Usage Snort	Metrics	This module checks that the average CPU usage of the Snort processes on the device is not overloaded and alerts when CPU usage exceeds the percentages configured for the module. The **Warning Threshold %** default value is `80`. The **Critical Threshold %** default value is `90`.
Snort Identity Memory Usage	Alert	Enables you to set a warning threshold for Snort identity processing and alerts when memory usage exceeds the level configured for the module. The **Critical Threshold %** default value is `80`. This health module specifically keeps track of the total space used for the user identity information in Snort. It displays the current memory usage details, the total number of user-to-IP bindings, and user-group mapping details. Snort records these details in a file. If the memory usage file is not available, the Health Alert for this module displays *Waiting for data*. This could happen during a Snort restart due to a new install or a major update, switch from Snort 2 to Snort 3 or back, or major policy deployment. Depending on the health monitoring cycle, and when the file is available, the warning disappears, and the health monitor displays the details for this module with its status turned Green.

Module	Type	Description
Memory Usage Snort	Metrics	This module checks the percentage of allocated memory used by the Snort process and alerts when memory usage exceeds the percentages configured for the module. The **Warning Threshold %** default value is `80`. The **Critical Threshold %** default value is `90`.
Snort Reconfiguring Detection	Metrics	Alerts if a device reconfiguration has failed. This module detects reconfiguration failure for both Snort 2 and Snort 3 instances.
Snort Statistics	Metrics	Monitors Snort statistics for events, flows, and packets.
SSE Connection Status	Metrics	The module alerts if the device cannot connect to the security services exchange cloud after an initial successful connection. Disabled by default.
CPU Usage System	Metrics	This module checks that the average CPU usage of all system processes on the device is not overloaded and alerts when CPU usage exceeds the percentages configured for the module. The **Warning Threshold %** default value is `80`. The **Critical Threshold %** default value is `90`.

Module	Type	Description
Threat Data Updates on Devices	Alert	Certain intelligence data and configurations that devices use to detect threats are updated on the Firewall Management Center from the cloud every 30 minutes. This module alerts you if this information has not been updated on the devices within the time period you have specified. Monitored updates include: • Local URL category and reputation data • Security Intelligence URL lists and feeds, including global Block and Do Not Block lists and URLs from Threat Intelligence Director • Security Intelligence network lists and feeds (IP addresses), including global Block and Do Not Block lists and IP addresses from Threat Intelligence Director • Security Intelligence DNS lists and feeds, including global Block and Do Not Block lists and domains from Threat Intelligence Director • Local malware analysis signatures (from ClamAV) • SHA lists from Threat Intelligence Director, as listed on the **Objects > Object Management > Security Intelligence > Network Lists and Feeds** page • Dynamic analysis settings configured on the **Integration** > **Dynamic Attributes Connector** page • Threat Configuration settings related to expiration of cached URLs, including the Cached URLs Expire setting on the **Integration** > **Other Integrations** > **Cloud Services** page. (Updates to the URL cache are not monitored by this module.) • Communication issues with the Cisco cloud for sending events. Refer to **Cisco Cloud** box on the **Integration** > **Other Integrations > Cloud Services** page. **Note** Threat Intelligence Director updates are included only if TID is configured on your system and you have feeds. By default, this module sends a warning after 1 hour and a critical alert after 24 hours. If this module indicates failure on the Firewall Management Center or on any devices, verify that the Firewall Management Center can reach the devices.
VPN Statistics	Metrics	Monitors site-to-site and remote access VPN tunnels between Firewall Threat Defense devices.
XTLS Counters	Metrics	Monitors XTLS/SSL flows, memory and cache effectiveness. Disabled by default.

Table 24: Management Center Health Modules

Module	Type	Description
Certificate Monitoring	Alert	Alerts when service authentication certificates are near expiration or have expired, based on a configurable threshold (in days). To avoid service disruption, renew certificates before they expire.
CPU Usage (per core)	Metrics	This module checks that the CPU usage on all the cores is not overloaded and alerts when CPU usage exceeds the thresholds configured for the module. The **Warning Threshold %** default value is `80`. The **Critical Threshold %** default value is `90`.
Critical Process Statistics	Metrics	Monitors the state of critical processes, their resource consumption, and the restart counts.
Database Size	Alert	Alerts if the configuration database size is too big. If this module alerts, contact Cisco TAC. This will prevent upgrade.
Discovery Host Limit	Alert	This module determines if the number of hosts the Firewall Management Center can monitor is approaching the limit and alerts based on the warning level configured for the module. For more information, refer to Host limits.
Disk Status	Alert	This module examines the performance of the hard disk and storage pack (if installed) on the appliance. The disk status health module generates a critical alert (red) if either of the following conditions is detected: • One or more disks may have failed or are experiencing errors. • An installed storage pack cannot be detected or unrecognized. All other disk health conditions generate a warning alert (yellow). For further diagnosis of warning conditions, contact Cisco Technical Assistance Center.

Module	Type	Description
Disk Usage	Metrics	This module compares disk usage on the appliance's hard drive and malware storage pack to the limits configured for the module and alerts when usage exceeds the thresholds configured for the module. This module also alerts when the system excessively deletes files in monitored disk usage categories, or when disk usage excluding those categories reaches excessive levels, based on module thresholds. Refer to Disk usage and drain of events health monitor alerts, on page 452 for information about troubleshooting scenarios for Disk Usage alerts. The Disk Usage module sends a health alert if the size of device configuration history files exceeds the allowed limit. Refer to Disk Usage for Device Configuration History Files Health Monitoring Alert for information about troubleshooting scenarios for the disk usage alerts. This health alert is not supported on Secure Firewall Management Center Versions 7.2.0-7.2.5, 7.3.x, and 7.4.0. Use the Disk Usage health status module to monitor disk usage for the `/` and `/volume` partitions on the appliance and track draining frequency. Although the disk usage module lists the `/boot` partition as a monitored partition, the size of the partition is static so the module does not alert on the boot partition. **Attention** If you receive alerts for high unmanaged disk usage for the partition `/volume` although the usage is below the critical or warning threshold specified in the health policy, this could indicate that there are files which must be deleted manually from the system. Contact TAC if you receive these alerts.
eStream Status	Alert	Monitors connections to third-party client applications that use the Event Streamer on the Firewall Management Center.
Event Backlog Status	Alert	Alerts if the backlog of event data awaiting transmission from the device to the Firewall Management Center has grown continuously for more than 30 minutes. To reduce the backlog, evaluate your bandwidth and consider logging fewer events.
Event Monitor	Metrics	This module monitors overall incoming event rate to Firewall Management Center.
File System Integrity Check	Alert	This module performs a file system integrity check and runs if the system has CC mode or UCAPL mode enabled, or if the system runs an image signed with a DEV key. This module is enabled by default.
Firewall Management Center HA Status	Alert	Monitors Firewall Management Center high availability. This module generates alerts if the HA pairs are not synchronized and if there is a discrepancy in the number of managed devices between the active and standby units.
Hardware Statistics	Metrics	Monitors Firewall Management Center hardware: fan speed, temperature, and power supply. Alerts when values exceed configurable thresholds.
License Monitor	Alert	This module monitors expiration of Classic licenses.

Module	Type	Description
Memory Usage	Alert	This module compares memory usage on the appliance to the limits configured for the module and alerts when usage exceeds the levels configured for the module. For appliances with more than 4 GB of memory, the preset alert thresholds are based on a formula that accounts for proportions of available memory likely to cause system problems. On >4 GB appliances, because the interval between Warning and Critical thresholds may be very narrow, it is recommended that you manually set the **Warning Threshold %** value to `50`. This will further ensure that you receive memory alerts for your appliance in time to address the issue. For additional information about how thresholds are calculated, refer to Memory usage thresholds for health monitor alerts, on page 451. Beginning with Version 6.6.0, the minimum required RAM for Firewall Management Center Virtual upgrades to Version 6.6.0+ is 28 GB, and the recommended RAM for Firewall Management Center Virtual deployments is 32 GB. We recommend you do not decrease the default settings: 32 GB RAM for most Firewall Management Center Virtual instances, 64 GB for the Firewall Management Center Virtual 300 (VMware only). **Attention** A critical alert is generated by the health monitor when insufficient RAM is allocated to a Firewall Management Center Virtual deployment. Complex access control policies and rules can command significant resources and negatively affect performance.
MySQL Statistics	Metrics	Monitors the status of the MySQL database, including the database size, number of active connections, and memory use.
MonetDB Statistics	Metrics	MonetDB is the database for firewall events and related data, such as connection summaries. This health module monitors the status of the MonetDB, including its size, the number of active connections, and memory usage. Additionally, it enables you to monitor the number of data requests being processed, and identify any slow-running requests. You can access these metrics and create custom dashboard to monitor MonetDB's health.
RabbitMQ Status	Metrics	Monitors and collects RabbitMQ statistics.

Module	Type	Description
Realm	Alert	Allows you to set a warning threshold for realm or user mismatches, which are: • User mismatch: A user is reported to the Secure Firewall Management Center without being downloaded. A typical reason for a user mismatch is that the user belongs to a group you have excluded from being downloaded to the Secure Firewall Management Center. Review the information discussed in *Cisco Secure Firewall Management Center Device Configuration Guide*. • Realm mismatch: A user logs into a domain that corresponds to a realm not known to the Firewall Management Center. For more information, *Cisco Secure Firewall Management Center Device Configuration Guide*. This module also displays health alerts when you try to download more users than the maximum number of downloaded users supported per realm. The maximum number of downloaded users for a single realm depends on your management center model. For more information, refer to *User Limit* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.
RRD Server Process	Alert	Alerts if the round robin data (RRD) server that stores time series data has restarted since the last time it updated. You can configure additional warning and critical thresholds for consecutive restarts.
Sybase Statistics	Metrics	Monitors the status of the Sybase database, including the database size, number of active connections, and memory use.
Time Series Data (RRD) Monitor	Alert	This module tracks the presence of corrupt files in the directory where time series data (such as correlation event counts) are stored and alerts when files are flagged as corrupt and removed.
Time Server Status	Alert	This module monitors the configuration of the NTP servers and alerts when the NTP server is unavailable or if the NTP server configuration is invalid. If you receive critical alert from this module, choose **System** (⚙) > **Configuration** > **Time Synchronization** and check the configuration of the NTP server specified in the alert. Requires Version 7.2.6.
VPN Tunnel Status	Alert	Alerts when VPN tunnels are down. Supported for both remote access and site-to-site VPN .
Web Server Connection Statistics	Metric	Alerts you when a single IP address exceeds the configurable limit of concurrent HTTP or HTTPS connections to the management center web server. It warns you when maximum number of browser tabs are opened to connect to management center from single IP address, ensuring optimal performance.

Configure health monitoring

Health monitoring enables proactive management of network appliances by tracking their operational status using customizable policies and alert mechanisms. You can configure specific monitoring policies for different appliance types, enabling only the relevant tests for each type.

Procedure

Step 1 Determine which health modules you want to monitor as discussed in Health modules, on page 381.

You can set up specific policies for each kind of appliance, enabling only the appropriate tests for that appliance.

Tip
To quickly enable health monitoring without customizing the monitoring behavior, you can apply the default policy provided for that purpose.

Step 2 Apply a health policy to each appliance where you want to track health status as discussed in Create a health policy, on page 393.

Step 3 (Optional.) Configure health monitor alerts as discussed in Create a health monitor alert, on page 401.

You can set up email, syslog, or SNMP alerts that trigger when the health status level reaches a particular severity level for specific health modules.

Health policies

A health policy is a configuration framework that

- enables you to define test criteria for monitoring the health of appliances,
- allows selective enabling of health modules and critical alert settings, and
- provides flexibility to apply, customize, or use default policies for each appliance.

Health policy configuration options

When you configure a health policy, you decide whether to enable each health module for that policy. You also select the criteria that control which health status each enabled module reports each time it assesses the health of a process.

You can create one health policy that can be applied to every appliance in your system, customize each health policy to the specific appliance where you plan to apply it, or use the default health policy provided for you.

Note When you register an appliance, the Firewall Management Center automatically assigns it the default health policy. To disassociate a health policy from an appliance, you must first associate a different health policy with it. An appliance must have at least one health policy assigned.

Default health policy

The Firewall Management Center setup process creates and applies an initial health policy, in which most—but not all—available health modules are enabled. The system also applies this initial policy to devices added to the Firewall Management Center.

This initial health policy is based on a default health policy. You cannot view or edit the default health policy, but you can copy it when creating a custom health policy.

Upgrades and the default health policy

When you upgrade the Firewall Management Center, any new health modules are added to all health policies, including the initial health policy, default health policy, and any other custom health policies. Usually, new health modules are added in an enabled state.

Note For a new health module to begin monitoring and alerting, reapply health policies after upgrade.

Create a health policy

If you want to customize a health policy to use with your appliances, you can create a new policy. The settings in the policy initially populate with the settings from the health policy you choose as a basis for the new policy. You can edit the policy to specify your preferences, such as enable or disable modules within the policy, change the alerting criteria for each module as needed, and specify the run time intervals.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Policy**.

Step 2 Click **Create Policy**.

Step 3 Enter a name for the policy.

Note that the following names are reserved for the default policies, and you cannot create a health policy using these names:

- Default Device Policy
- Default Health Policy

Step 4 Choose the existing policy that you want to use as the basis for the new policy from the **Base Policy** drop-down list.

Step 5 Enter a description for the policy.

Step 6 Choose **Save**.

What to do next

- Apply the health policy on devices as described in Apply a health policy, on page 394.

- Edit the policy to specify the module-level policy settings as described in Edit a health policy, on page 394.

Apply a health policy

When you apply a health policy to an appliance, the health tests for all the modules you enabled in the policy automatically monitor the health of the processes and hardware on the appliance. Health tests then continue to run at the intervals you configured in the policy, collecting health data for the appliance and forwarding that data to the Firewall Management Center.

If you enable a module in a health policy and then apply the policy to an appliance that does not require that health test, the health monitor reports the status for that health module as disabled.

To stop health monitoring for an appliance, create a health policy with all modules disabled, assign it to the appliance, and remove the appliance's other health policy assignments. The appliance retains one assigned health policy, but no health modules are monitored.

When you apply a different policy to an appliance that already has a policy applied, expect some latency in the display of new data based on the newly applied tests.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Policy**.

Step 2 Click the **Deploy health policy** (📥) next to the policy you want to apply.

Step 3 Choose the appliances where you want to apply the health policy.

Note
An appliance must have at least one health policy assigned to it. To stop health monitoring for an appliance, create a health policy with all modules disabled, assign it to the appliance, and remove the appliance's other health policy assignments. The appliance retains one assigned health policy, but no health modules are monitored. To disassociate a health policy from an appliance, you must first associate a different health policy with it.

Step 4 Click **Apply** to apply the policy to the appliances you chose.

The health policy is successfully applied to the selected appliances, and monitoring begins automatically.

What to do next

- Optionally, monitor the task status; see View task messages, on page 450.

  Monitoring of the appliance starts when the policy is successfully applied.

Edit a health policy

Edit a health policy when you need to adjust which modules are monitored, change critical or warning thresholds, or reassign policies to different devices.

Before you begin

- Ensure you have access to Firewall Management Center.
- Identify the health policy you need to modify.

Follow these steps to edit a health policy:

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Policy**.

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

Step 3 To edit the policy name and its description, click the **Edit** (✎) icon provided against the policy name.

Step 4 The **Health Modules** tab displays all the device modules and its attributes. Configure your health modules using the following actions:

- Click the toggle button that is provided against the module and its attributes—turn on () or turn off () to enable or disable testing of health status respectively.
- To execute a bulk enable or disable testing on the health modules, click the **Select All** toggle button

Note

- The modules and attributes are flagged with the supporting appliances—Firewall Threat Defense, Firewall Management Center, or both.
- You cannot choose to include or exclude the individual attributes of CPU and Memory modules.

For information on the modules, refer to Health modules, on page 381.

Step 5 Where appropriate, set the **Critical** and **Warning** threshold percentages.

Step 6 In the **Run Time Intervals** tab, enter the relevant values in the fields:

- **Health Module Run Interval**: The frequency for running the health modules. The minimum interval is 5 minutes.
- **Metric Collection Interval**: The frequency of collecting the time series data on the device and its health modules. The device monitor reports these metrics in several predefined health monitor dashboards by default. For detailed information on the dashboard, refer to View the device health monitor, on page 408. The metric data is collected for analysis and hence no alerting is associated with it.

Step 7 To view and modify the devices to which the policy is assigned, do the following:

a) Click **Policy Assignments & Deploy**.

b) From the **Available Devices** list, click the + icon next to the device to which you want to assign the health policy.

c) Click **Apply**.

Note

The UI prompts you to save unsaved changes before continuing with assignment and deployment.

Alternatively, you can apply the health policy to your appliance as described in Apply a health policy, on page 394

Apply the health policy to each appliance where you want to track health status. When you apply the health policy to an appliance, all the modules you enabled in the policy monitor the health of the processes and hardware on the appliance, and forwards that data to the Firewall Management Center.

Step 8 Click **Save**.

Delete a health policy

You can delete health policies that you no longer need. However, an appliance must have at least one health policy assigned to it. If you delete a policy that is still applied to an appliance, the policy settings remain in effect until you apply a different policy. In addition, if you delete a health policy that is applied to a device, any health monitoring alerts in effect for the device remain active until you disable the underlying associated alert response.

Note To stop health monitoring for an appliance, create a health policy with all modules disabled, assign it to the appliance, and remove the appliance's other health policy assignments. The appliance retains one assigned health policy, but no health modules are monitored.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Policy**.

Step 2 Click **Delete** ( ) next to the policy that you want to delete.

Step 3 Click **Delete health policy** to delete it.

Device exclusion in health monitoring

A device exclusion is a health monitoring configuration that

- prevents deliberate device or module outages from affecting the overall health summary,
- allows administrators to temporarily disable health monitoring for selected devices, Firewall Management Center where supported, or specific health modules, and
- ensures health events generated during exclusion are marked as disabled and do not affect health calculations.

Device exclusion behavior

In the course of normal network maintenance, you may disable managed devices or make them temporarily unavailable. In Firewall Management Center, you can also exclude Firewall Management Center health monitoring where applicable. Because those outages or expected conditions are deliberate, you may not want their health status to affect the overall health summary.

When you disable health monitoring status, health events are still generated, but they have a disabled status and do not affect the health status for the health monitor. If you remove the device or module from the excluded list, the events that were generated during the exclusion continue to show a status of disabled.



Note If all health modules or a device are added to the exclusion list, the device's health status will be set to disabled. This status refers only to health monitoring and should not be interpreted as the device being offline or disconnected.

You can also disable an individual health module. For example, when you reach the host limit on the Firewall Management Center, you can disable Host Limit status messages. Excluding health modules for individual interfaces is not supported on devices operating in transparent mode.

Note that on the main Health Monitor page you can distinguish between devices that are excluded if you expand to view the list of devices with a particular status by clicking the arrow in that status row.



Note On Firewall Management Center, Health Monitor exclusion settings are local configuration settings. Therefore, if you exclude a device, then delete it and later re-register it with the Firewall Management Center, the exclusion settings remain persistent. The newly re-registered device remains excluded.

Temporary device exclusion scenario

You can use the health monitor exclude feature to disable health monitoring status reporting for a device, Firewall Management Center where applicable, or specific health modules. For example, if you know that a segment of your network will be unavailable, you can temporarily disable health monitoring for a managed device on that segment to prevent the health status on the Firewall Management Center from displaying a warning or critical state because of the lapsed connection to the device.

To temporarily disable health events from a managed device, or from Firewall Management Center where supported, go to the exclusion configuration page and add the target to the exclusion list. After the setting takes effect, the system no longer considers the excluded device when calculating the overall health status. The Health Monitor Appliance Status Summary lists the device as disabled.

Exclude devices from health monitoring

Exclude devices from health monitoring to disable their events and health status when needed for maintenance or troubleshooting purposes.

You can exclude managed devices individually or by group, model, associated health policy, or other supported scope. In Firewall Management Center exclusion can also apply to Firewall Management Center where supported.

If you need to set the events and health status for an individual managed device to disabled, you can exclude the device from health monitoring. Where supported, you can also exclude Firewall Management Center health monitoring. After the exclusion settings take effect, the device shows as disabled in the Health Monitor Appliance Module Summary, and health events for the device have a status of disabled.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Exclude**.

Step 2 Click **Add Device**.

Step 3 In the **Device Exclusion** dialog box, under **Available Devices**, click **Add** (⊕) against the device that you want to exclude from health monitoring.

Step 4 Click **Exclude**.

The selected device is displayed in the exclusion main page.

Step 5 To remove the device from the exclusion list, click **Delete** (🗑).

Step 6 Click **Apply**.

Exclude health policy modules on a device

Excluding health policy modules is useful when certain modules produce irrelevant or temporary health alerts, which you do not want to affect the overall device status. Exclusion settings indicate the number of modules excluded from health monitoring.

Note Excluding health modules for individual interfaces is not supported on devices operating in transparent mode.

Tip Make sure that you keep track of individually excluded modules so you can reactivate them when you need them. You may miss essential warning or critical messages if you accidentally leave a module disabled.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Exclude**.

Step 2 Click **Edit** (✎) next to the Firewall Threat Defense device you want to modify.

Step 3 In the **Exclude Health Modules** dialog box, by default, all the modules of the device are excluded from health monitoring.

Certain modules are applicable to a specific device only; for more information, refer to Health modules, on page 381.

Step 4 To choose modules to be excluded from health monitoring, click the **Enable Module Level Exclusion** link.

The **Exclude Health Modules** dialog box displays all the modules of the device. The modules that are not applicable for the associated health policies are disabled by default. To exclude a module, perform these steps:

a. Click the **Slider** ( ) button next to the desired module.

b. To specify the duration of the exclusion for the selected modules, from the **Exclude Period** drop-down list, select the duration.

Step 5 If you select an **Exclude Period** other than **Permanent**, for your exclusion configuration, you can choose to automatically delete the configuration when it expires. To enable this setting, check the **Auto-delete expired configurations** check box.

Step 6 Click **OK**.

Step 7 Click **Apply**.

Note
Interface exclusion differs from other module exclusions because you can exclude selected physical interfaces while still monitoring associated subinterfaces, depending on the UI options available in your version.

Exclude interface health monitoring

Use interface exclusion when you want to suppress health status updates for selected physical interfaces without suppressing health monitoring for the entire device. Interface exclusion is useful when interface state is expected during maintenance or when you want to focus on subinterface health instead of physical interface status.

When you configure interface exclusion, select the physical interfaces that you want to exclude from the list of available interfaces. Depending on the available UI options in your version, selecting a physical interface can also exclude its subinterfaces. If subinterface controls are available, you can continue monitoring subinterfaces while excluding only the physical interface.

Use this option to reduce unnecessary health alerts when traffic loss or interface state changes are expected on a physical interface.

Manage expired health monitor exclusions

When monitoring device or module health, some items may have temporary exclusions. Once these exclusion periods lapse, you need to decide whether to renew the exclusion or clear the device/module from exclusion so alerts resume.

Procedure

Step 1 Choose **System** ( ) > **Health** > **Exclude**.

The **Warning** ( ) icon is displayed against the device indicating the expiry of the duration of exclusion of the device or the modules from alerting.

Step 2 To renew the exclusion of the device, click **Edit** ( ) next to the device. In the **Exclude Health Modules** dialog box, click the **Renew** link. The exclusion period of the device is extended with the current value.

Step 3 To clear the device from being excluded, click **Delete** ( ) next to the device, click **Remove the device from exclusion**, and then click **Apply**.

Step 4 To renew or clear the modules from exclusion, click **Edit** ( ) next to the device. In the **Exclude Health Modules** dialog box, click the **Enable Module Level Exclusion** link, and then click the **Renew** or **Clear** link against the modules. When you click **Renew**, the exclusion period is extended on the module with the current value.

Health monitor alerts

A health monitor alert is a notification mechanism that

- provides email, SNMP, or syslog notifications when module status changes occur,
- enables association of alert responses with health event levels to trigger notifications based on event severity, and
- allows customization to monitor specific system conditions such as disk space thresholds.

Disk space monitoring alert example

If you are concerned that your appliances may run out of hard disk space, you can automatically send an email to a system administrator when the remaining disk space reaches the warning level. If the hard drive continues to fill, you can send a second email when the hard drive reaches the critical level.

Health monitor alert information

The alerts generated by the health monitor contain specific information that helps identify issues and their severity.

Health monitor alerts include these attributes:

- Severity: Indicates the severity level of the alert.
- Module: Specifies the health module whose test results triggered the alert.
- Description: Includes the health test results that triggered the alert.

The table describes these severity levels.

Table 25: Alert severities

Severity	Description
Critical	The health test results met the criteria to trigger a Critical alert status.
Warning	The health test results met the criteria to trigger a Warning alert status.

Severity	Description
Normal	The health test results met the criteria to trigger a Normal alert status.
Error	The health test did not run.
Recovered	The health test results met the criteria to return to a normal alert status, following a Critical or Warning alert status.

Create a health monitor alert

When you create a health monitor alert, you create an association between a severity level, a health module, and an alert response. You can use an existing alert or configure a new one specifically to report on system health. When the severity level occurs for the selected module, the alert triggers.

If you create or update a threshold in a way that duplicates an existing threshold, you are notified of the conflict. When duplicate thresholds exist, the health monitor uses the threshold that generates the fewest alerts and ignores the others. The timeout value for the threshold must be between 5 and 4,294,967,295 minutes.

Before you begin

- You must be an Admin user.
- Ensure that you have configured alert responses that govern the Firewall Management Center's communication with the SNMP, syslog, email server, or webhook where you send the health alert; For more information, refer to Configuring external alerts with alert responses, on page 561.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor Alerts**.

Step 2 Click **Add**.

Step 3 In the **Add Health Alert** dialog box, enter a name for the health alert in the **Health Alert Name** field.

Step 4 From the **Severity** drop-down list, choose the severity level you want to use to trigger the alert.

Step 5 From the **Alert** drop-down list, choose the alert response that you want to trigger when the specified severity level is reached. If you have not yet configured the alert responses, click **Alerts** to visit the **Alerts** page and set them.

Step 6 From the **Health Modules** list, choose the health policy modules for which you want the alert to apply.

Step 7 Optionally, in the **Threshold Timeout** field, enter the number of minutes that should elapse before each threshold period ends and the threshold count resets.

Even if the policy run time interval value is less than the threshold timeout value, the interval between two reported health events from a given module is always greater. For example, if you change the threshold timeout to 8 minutes and the policy run time interval is 5 minutes, there is a 10-minute interval (5 x 2) between reported events.

Step 8 Click **Save** to save the health alert.

Edit a health monitor alert

Administrators may need to update the alert parameters for health monitor alerts in order to adjust severity, change modules, or select different responses as requirements evolve.

Before you begin

You must be an Admin user to perform this task.

Follow these steps to edit health monitor alerts:

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor Alerts**.

Step 2 Click the **Edit** (✎) icon that is provided against the required health alert that you want to modify.

Step 3 In the **Edit Health Alert** dialog box, from the **Alert** drop-down list, select the required alert entry, or click **Alerts** link to configure a new alert entry.

Step 4 Click **Save**.

Delete health monitor alerts

Health monitor alerts notify you about system events, issues, or warnings. You may need to delete outdated or irrelevant alerts to streamline monitoring.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor Alerts**.

Step 2 Click **Delete** (🗑) next to the health alert you want to delete.

Step 3 Click **Delete health alert** to delete it.

What to do next

Disable or delete the underlying alert response to ensure that alerting does not continue. For more information, refer to Configuring external alerts with alert responses, on page 561.

About the health monitor

The health monitor provides the compiled health status for all devices managed by the Firewall Management Center, plus the Firewall Management Center itself.

The health monitor is composed of:

- The Health Status summary page — Provides you with an at-a-glance view of the health of the Firewall Management Center and all of the devices that the Firewall Management Center manages. In a multi-domain deployment, you can view the health status summary for a device in the domain where the device is located. Devices are listed individually, or grouped according to their geolocation, high availability, or cluster status where applicable.
  - View the health summary of the Firewall Management Center and any device when you hover on the hexagon that represents the device health.
  - The dot to the left of a device indicates its health:
    - Green — No alarms.
    - Orange — At least one health warning.
    - Red — At least one critical health alarm.
- The Monitoring navigation pane — Allows you to navigate the device hierarchy. You can view health monitors for individual devices from the navigation pane.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Step 2 View the status of the Firewall Management Center and its managed devices in the **Health Status** landing page.

a) Hover your pointer over a hexagon to view the health summary of a device. The popup window shows a truncated summary of the top five health alerts. Click on the popup to open a detailed view of the health alert summary.

b) In the device list, click **Expand**(>) and **Collapse** (⌄) to expand and collapse the list of health alerts for a device.

When you expand the row, all of the health alerts are listed, including the status, title, and details.

Note
Health alerts are sorted by their severity level.

Step 3 Use the Monitoring navigation pane to access device-specific health monitors.

a) In the device list, click **Expand** (>) and **Collapse** (⌄) to expand and collapse the list of managed devices.

When you expand the row, all of the devices are listed.

b) Click on a device to view a device-specific health monitor.

c) In the health monitor, hover over a graph to view all metrics and their respective values at that specific point on the graph. Click on the graph to pin the metrics statistics box, allowing you to explore the metrics and their values in detail. To close the statistics box and move to another point on the graph, simply click the close button.

You have accessed the health monitor and can view the compiled health status for all managed devices. You can interact with the Health Status landing page and use the Monitoring navigation pane to access device-specific health information.

What to do next

- See Device health monitors, on page 407 for information about the compiled health status and metrics for any device managed by the Firewall Management Center.
- See Using Health Monitor for information about the health status of the Firewall Management Center.

  To return to the Health Status landing page at any time, click **Home**.

View the Firewall Management Center health monitor dashboard

The Firewall Management Center monitor provides a detailed view of the health status of the Firewall Management Center. The health monitor is composed of:

- High Availability (if configured): The High Availability (HA) panel displays the current HA status, including the status of the Active and Standby units, the last sync time, and overall device health.
- Event Rate: The Event Rate panel shows the maximum event rate as a base line as well as the overall event rate received by the Firewall Management Center.
- Event Capacity: The Event Capacity panel shows the current consumption by event categories, including the retention time of events, the current vs. maximum event capacity, and a capacity overflow mechanism where you are alerted when events are stored beyond the configured maximum capacity of the Firewall Management Center.
- Process Health: The Process Health panel has an at-a-glance view of the critical processes as well as a tab that lets you see state of all processed, including the CPU and memory usage for each process.
- CPU: The CPU panel lets you toggle between the average CPU usage (default) and the CPU usage of all cores.
- Memory: The Memory panel shows the overall memory usage on the Firewall Management Center.
- Interface: The Interface panel shows average input and output rate of all interfaces.
- Disk Usage: The Disk Usage panel shows the use of entire disk, and the use of the critical partitions where Firewall Management Center data is stored.

Tip Your session normally logs you out after one hour of inactivity (or another configured interval). If you plan to passively monitor health status for long periods of time, consider exempting some users from session timeout, or changing the system timeout settings. For more information, refer to Add or edit an internal user, on page 148 and Configure session timeouts, on page 121.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Step 2 Use the **Monitoring** navigation pane to access the Firewall Management Center and device-specific health monitors.

- A standalone Firewall Management Center is shown as a single node; a high-availability Firewall Management Center is shown as a pair of nodes.
- The health monitor is available to both the active and standby Firewall Management Center in an HA pair, and the health alerts appear in both the units. However, you must resolve any health alerts from the active Firewall Management Center, as access to various pages can be restricted in the standby Firewall Management Center.

Step 3 Explore the Firewall Management Center dashboard.

The Firewall Management Center dashboard includes a summary view of the HA state of the Firewall Management Center (if configured), as well as at-a-glance views of Firewall Management Center processes and device metrics such as CPU, memory, and disk usage.

Run all health modules for an appliance

Health module tests run automatically at the policy run time interval you configure when you create a health policy. However, you can also run all health module tests on demand to collect up-to-date health information for the appliance.

Follow these steps to run all health modules for an appliance:

Procedure

Step 1 View the health monitor for the appliance.

Step 2 Click **Run All Modules**.
The status bar indicates the progress of the tests, then the Health Monitor Appliance page refreshes.

Note
When you manually run health modules, the first refresh that automatically occurs may not reflect the data from the manually run tests. If the value has not changed for a module that you just ran manually, wait a few seconds, then refresh the page by clicking the device name. You can also wait for the page to refresh again automatically.

Run a health module test

Health module tests run automatically at the policy run time interval you configure when you create a health policy. However, you can also run a health module test on demand to collect up-to-date health information for that module.

Before you begin

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

Follow these steps to run a health module test:

Procedure

Step 1 View the health monitor for the appliance.

Step 2 In the **Module Status Summary** graph, click the color for the health alert status category you want to view.

Step 3 In the **Alert Detail** row for the alert for which you want to view a list of events, click **Run**.

The status bar indicates the progress of the test, then the Health Monitor Appliance page refreshes.

Note
When you manually run health modules, the first refresh that automatically occurs may not reflect the data from the manually run tests. If the value has not changed for a module that you just manually ran, wait a few seconds, then refresh the page by clicking the device name. You can also wait for the page to refresh automatically again.

Generate health module alert graphs

Health module alert graphs help you track appliance performance and diagnose issues using historical alert data.

Before you begin

Confirm you have access to the desired appliance's health monitor.

Follow these steps to generate health module alert graphs.

Procedure

Step 1 View the health monitor for the appliance.

Step 2 In the **Module Status Summary** graph of the Health Monitor Appliance page, click the color for the health alert status category you want to view.

Step 3 In the **Alert Detail** row for the alert for which you want to view a list of events, click **Graph**.

Tip
If no events appear, you may need to adjust the time range.

Device health monitors

The device health monitor provides the compiled health status for any device managed by the Firewall Management Center. The device health monitor collects health metrics for Secure Firewall devices in order to predict and respond to system events.

The device health monitor is comprised of these components:

- System Details: Displays information about the managed device, including the installed Secure Firewall version and other deployment details.
- Troubleshooting & Links: Provides convenient links to frequently used troubleshooting topics and procedures.
- Health alerts: A health alert monitor provides an at-a-glance view of the health of the device.
- Time range: An adjustable time window to constrain the information that appears in the various device metrics windows.
- Device metrics: An array of key firewall device health metrics categorized across predefined dashboards, including:
  - CPU: CPU utilization, including the CPU usage by process and by physical cores. The Firewall Threat Defense CPU core allocation dashboard shows core assignments for these categories:
    - **Data Plane**: Handles basic network functions, including core packet forwarding and network data processing.
    - **Snort**: Manages intrusion detection and deep packet inspection features.



Note Multi-threaded Snort processes are allocated across multiple Snort cores to improve device performance. The Snort core allocation shown in the **CPU** dashboard reflects the total number of Snort process threads assigned to Snort cores.

    - **System**: Includes all other system processes. While some processes may have dedicated CPU cores allocated, their usage is combined and displayed under the **System** category in the dashboard.

  - Memory: Device memory utilization, including data plane and Snort memory usage.
  - Interfaces: Interface status and aggregate traffic statistics.
  - Connections: Connection statistics (such as elephant flows, active connections, peak connections, and so on) and NAT translation counts.
  - Snort: Statistics related to the Snort process.
  - Disk Usage: Device disk usage, including the disk size and disk utilization per partition.
  - Critical Processes: Statistics related to managed processes, including process restarts and other select health monitors such as CPU and memory utilization.

For a comprehensive list of the supported device metrics, refer to Cisco Secure Firewall Threat Defense Health Metrics.

View system details and troubleshoot devices

Use this task when you need to monitor device health or initiate troubleshooting procedures on specific devices through the management center.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Use the Monitoring navigation pane to access device-specific health monitors.

Step 2 In the device list, click **Expand**(>) and **Collapse** (⌄) to expand and collapse the list of managed devices.

Step 3 Click on a device to view a device-specific health monitor.

Step 4 Click the link for **View System & Troubleshoot Details …**

This panel is collapsed by default. Clicking on the link expands the collapsed section to see **System Details** and **Troubleshooting & Links** for the device. The system details include:

- **Version:** The Secure Firewall software version.
- **Model:** The device model.
- **Mode:** The firewall mode. The Firewall Threat Defense device supports two firewall modes for regular firewall interfaces: Routed mode and Transparent mode.
- **VDB:** The Cisco vulnerability database (VDB) version.
- **SRU:** The intrusion rule set version.
- **Snort:** The Snort version.

Step 5 You have these troubleshoot choices:

- Generate troubleshooting files; refer to Generate troubleshooting files for specific system functions, on page 457
- Generate and download advanced troubleshooting files; refer to Download advanced troubleshooting files, on page 458.
- Create and modify health policies; refer to Create a health policy, on page 393.
- Create and modify health monitor alerts; refer to Create a health monitor alert, on page 401.

View the device health monitor

The device health monitor compiles device metrics and displays health status and trends using an array of dashboards.

Follow these steps to view the device health monitor.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Use the Monitoring navigation pane to access device-specific health monitors.

Step 2 In the device list, click **Expand**(>) and **Collapse** (⌄) to expand and collapse the list of managed devices.

Step 3 View the **Health Alerts** for the device in the alert notification at the top of page, directly to the right of the device name.

Hover your pointer over the **Health Alerts** to view the health summary of the device. The pop-up window shows a truncated summary of the top five health alerts. Click the pop-up to open a detailed view of the health alert summary.

Step 4 Configure the time range from the drop-down in the upper-right corner.

The time range can reflect a period as short as the last hour (the default) or as long as two weeks. Select **Custom** from the drop-down to configure a custom start and end date.

Click the refresh icon to set auto refresh to 5 minutes or to toggle off auto refresh.

Step 5 Click the **Show Deployment Info** ( ) icon for a deployment overlay on the trend graph, with respect to the selected time range.

The **Show Deployment Info** ( ) icon indicates the number of deployments during the selected time-range. A vertical band indicates the deployment start and end time. In the case of multiple deployments, multiple bands/lines can appear. Click the icon on top of the dotted line to view the deployment details.

Step 6 View the device monitor reports health and performance metrics in several predefined dashboards by default.

The metrics dashboards include:

- Overview: Highlights key metrics from the other predefined dashboards, including CPU, memory, interfaces, connection statistics; plus disk usage and critical process information.
- CPU: CPU utilization, including the CPU usage by process and by physical cores.
- Memory: Device memory utilization, including data plane and Snort memory usage.
- Interfaces: Interface status and aggregate traffic statistics.
- Connections: Connection statistics (such as elephant flows, active connections, peak connections, and so on) and NAT translation counts.
- Snort: Statistics related to the Snort process.
- ASP Drops: Statistics related to the Accelerated Security Path (ASP) performance and behavior.

Note

- Multi-threaded Snort processes are allocated across multiple Snort cores to improve device performance. Note that the Snort core allocation shown in the **CPU**dashboard reflects the total number of Snort process threads assigned to Snort cores.
- The **Process Health** widget in Firewall Management Center displays the CPU% usage for each process. The CPU% metric reflects usage relative to the number of cores in the Firewall Threat Defense device,

where 100% corresponds to full utilization of one core. For example, on an 8-core device, 200% indicates two fully utilized cores, leaving six available for other processes. To monitor overall system CPU usage, use the **CPU** widget instead of the **Process Health** widget.

You can navigate through the various metrics dashboards by clicking on the labels. For a comprehensive list of the supported device metrics, refer to Cisco Secure Firewall Threat Defense Health Metrics.

Step 7 Click the **Add New Dashboard**(+) to create a custom correlation dashboard by building your own variable set from the available metric groups.

For more information, refer to Correlating device metrics, on page 410.

Correlating device metrics

The device health monitor includes an array of key Firewall Threat Defense device metrics that serve to predict and respond to system events. The health of any Firewall Threat Defense device can be determined by these reported metrics.

The device monitor reports these metrics in several predefined dashboards by default. These dashboards include:

- Overview: Highlights key metrics from the other predefined dashboards, including CPU, memory, interfaces, connection statistics; plus disk usage and critical process information.
- CPU: CPU utilization, including the CPU usage by process and by physical cores.
- Memory: Device memory utilization, including data plane and Snort memory usage.
- Interfaces: Interface status and aggregate traffic statistics.
- Connections: Connection statistics (such as elephant flows, active connections, peak connections, and so on) and NAT translation counts.
- Snort: Statistics related to the Snort process.
- ASP Drops: Statistics related to the Accelerated Security Path (ASP) performance and behavior.

You can add custom dashboards to correlate metrics that are interrelated. Select from predefined correlation groups, such as CPU and Snort; or create a custom correlation dashboard by building your own variable set from the available metric groups. For a comprehensive list of the supported device metrics, refer to Cisco Secure Firewall Threat Defense Health Metrics.

Before you begin

- To view and correlate the time series data (device metrics) in the health monitor dashboard, enable REST API (**System** (⚙) > **Configuration** > **REST API Preferences**).

Note Correlating device metrics is available only for Firewall Threat Defense 6.7 and later versions. Hence, for Firewall Threat Defense versions earlier than 6.7, the health monitor dashboard does not display these metrics even if you enable REST API.

Follow these steps to correlate device metrics:

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**. In the **Devices** list, click **Expand**(>) and **Collapse** (⌄) to expand and collapse the list of managed devices, and choose the device for which you want to modify the dashboard.

Use the Monitoring navigation pane to access device-specific health monitors.

Step 2 Click the **Add Dashboard** (+) icon in the upper right corner of the device monitor to add a new dashboard.

Step 3 From the **Select Correlation Group** drop-down, choose a predefined correlation group or to create a custom group. To create a dashboard from a predefined correlation group, choose the group and click **Add**.

Step 4 To create a custom correlation dashboard:

a) Choose **Custom**.
b) Enter a unique name in the **Dashboard Name** field or accept the default.
c) Choose a group from the **Select Metric Group** drop-down, then select corresponding metrics from the **Select Metrics** drop-down.

For a comprehensive list of the supported device metrics, refer to Cisco Secure Firewall Threat Defense Health Metrics.

Step 5 Click **Add Metrics** to add and select metrics from another group, and to remove an individual metric, click the Remove ✕ icon on the right side of the item. Click the delete icon to remove the entire group. Click **Add** to add the dashboard to the health monitor.

Step 6 You can **Edit** or **Delete** custom correlation dashboards.

Health monitor status categories

The table summarizes the available health monitor status categories, their severity levels, icons, associated colors in pie charts, and descriptions.

Table 26: Health status indicator

Status Level	Status Icon	Status Color in Pie Chart	Description
Error	**Error** (✕)	Black	Indicates that at least one health monitoring module has failed on the appliance and has not been successfully re-run since the failure occurred. Contact your technical support representative to obtain an update to the health monitoring module.
Critical	**Critical** (!)	Red	Indicates that the critical limits have been exceeded for at least one health module on the appliance and the problem has not been corrected.

Status Level	Status Icon	Status Color in Pie Chart	Description
Warning	**Warning** ()	Yellow	Indicates that warning limits have been exceeded for at least one health module on the appliance and the problem has not been corrected. This status also indicates a transitionary state, where, the required data is temporarily unavailable or could not be processed because of changes in the device configuration. Depending on the monitoring cycle, this transitionary state is auto-corrected.
Normal	**Normal** ()	Green	Indicates that all health modules on the appliance are running within the limits configured in the health policy applied to the appliance.
Recovered	**Recovered** ()	Green	Indicates that all health modules on the appliance are running within the limits configured in the health policy applied to the appliance, including modules that were in a Critical or Warning state.
Disabled	**Disabled** ()	Blue	Indicates that an appliance is disabled or excluded, that the appliance does not have a health policy applied to it, or that the appliance is currently unreachable.

Health event views

Health event views enable you to monitor health events logged by the health monitor in the Firewall Management Center. The views are fully customizable, allowing you to analyze health statuses efficiently and access related event data.

You can search event data to easily access other information that may be related to the events you are investigating. If you understand what conditions each health module tests for, you can more effectively configure alerting for health events.

You can perform many of the standard event view functions on the health event view pages.

View health events

Use this task when you need to monitor the health of your appliances and review any incidents or events recorded.

When you access health events from the Health Monitor page on your Firewall Management Center, you retrieve all health events for all managed appliances.

Tip You can bookmark this view to allow you to return to the page in the health events workflow containing the Health Events table of events. The bookmarked view retrieves events within the time range you are currently viewing, but you can then modify the time range to update the table with more recent information if needed.

Follow these steps to view health events.

Procedure

Choose **System** (⚙) > **Health** > **Events**.

Tip
If you are using a custom workflow that does not include the table view of health events, click **(switch workflow)**. On the **Select Workflow** page, click **Health Events**.

Note
If no events appear, you may need to adjust the time range.

View health events by module and appliance

This task helps you efficiently monitor appliance health by filtering health events based on module status and viewing detailed alerts for specified events.

Procedure

Step 1 View the health monitor for the appliance; refer to View the device health monitor, on page 408.

Step 2 In the **Module Status Summary** graph, click the color for the event status category you want to view.

The Alert Detail list toggles the display to show or hide events.

Step 3 In the **Alert Detail** row for the alert for which you want to view a list of events, click **Events**.

The Health Events page appears, containing results for a query with the name of the appliance and the name of the specified health alert module as constraints. If no events appear, you may need to adjust the time range.

Step 4 If you want to view all health events for the specified appliance, expand **Search Constraints**, and click the **Module Name** constraint to remove it.

View the health events table

The Health Events Table provides actionable insight into system health and allows users to organize, modify, and respond to health events.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Events**.

Step 2 You have the following choices:

- Bookmark: To bookmark the current page so that you can quickly return to it, click **Bookmark This Page**, provide a name for the bookmark, and click **Save**.
- Change Workflow: To choose another health events workflow, click **(switch workflow)**.
- Delete Events: To delete health events, check the check box next to the events you want to delete, and click **Delete**. To delete all the events in the current constrained view, click **Delete All**, then confirm you want to delete all the events.
- Generate Reports: Generate a report based on data in the table view — click **Report Designer**.
- Modify: Modify the time and date range for events listed in the Health table view. Note that events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.
- Navigate: Navigate through event view pages.
- Navigate Bookmark: To navigate to the bookmark management page, click **View Bookmarks** from any event view.
- Navigate Other: Navigate to other event tables to view associated events.
- Sort: Sort the events that appear, change what columns display in the table of events, or constrain the events that appear
- View All: To view event details for all events in the view, click **View All**.
- View Details: To view the details associated with a single health event, click the down arrow link on the left side of the event.
- View Multiple: To view event details for multiple health events, choose the check box next to the rows that correspond with the events you want to view details for and then click **View**.
- View Status: To view all events of a particular status, click status in the Status column for an event with that status.
- Search: Search for specific health events or health metrics recorded within your deployment. Customize your search query to analyze health events and other monitored data and save the search constraints for quick access. For more information, refer to Event searches, on page 707.

Health event fields

The health monitor modules you enable in your health policy run tests to determine appliance health status. When the health status meets specified criteria, a health event is generated. The health event table includes fields that are searchable and viewable for monitoring appliance health.

Table 27: Health event fields

Field	Description
Module Name	Specify the name of the module which generated the health events you want to view. For example, to view events that measure CPU performance, type **CPU**. The search should retrieve applicable CPU Usage and CPU temperature events.
Test Name (Search only)	The name of the health module that generated the event.
Time (Search only)	The timestamp for the health event.

Field	Description
Description	The description of the health module that generated the event. For example, health events generated when a process was unable to execute are labeled **Unable to Execute**.
Value	The value (number of units) of the result obtained by the health test that generated the event. For example, if the Firewall Management Center generates a health event whenever a device it is monitoring is using 80 percent or more of its CPU resources, the value could be a number from 80 to 100.
Units	The units descriptor for the result. You can use the asterisk (*) to create wildcard searches. For example, if the Firewall Management Center generates a health event when a device it is monitoring is using 80 percent or more of its CPU resources, the units descriptor is a percentage sign (%).
Status	The status (Error, Critical, Warning, Normal, Recovered, or Disabled) reported for the appliance.
Device	The appliance where the health event was reported.

History for health monitoring

Provides the history of health monitoring features and enhancements for each release of Secure Firewall Management Center.

Table 28: History for health monitoring

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Health alerts for NTP server sync issues.		Any	Introduced the **Time Server Status** module in the Secure Firewall Management Center Health Policy. When enabled, this module monitors the configuration of the NTP servers and alerts when the NTP server is unavailable or if the NTP server configuration is invalid. New/modified screens: **System** > **Health** > **Policy** > **Firewall Management Center Health Policy** > **Health Modules** > **Time Synchronization**.
Health alert for device configuration history files size	7.2.6	Any	The Disk Usage module sends health alert when the size of device configuration history files on the Firewall Management Center exceeds the allowed limit. This alert is enabled by default. Health alert for exceeding the configuration versions size is not supported on the Secure Firewall Management Center versions 7.3.0 and 7.4.0.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Health monitor usability enhancements.	7.1	Any	Following UI page were improved for better usability and presentation of data: • Policy • Exclude • Monitor Alerts New/modified screens: • **System** > **Health** > **Policy**. • **System** > **Health** > **Exclude**. • **System** > **Health** > **Monitor Alerts**.
Elephant flow detection.	7.1	Any	The health monitor includes the following enhancements: • The Connection statistics includes active elephant flows. • The Connection Group Metrics includes the number of active elephant flows. The Elephant Flow Detection feature is not supported on the Cisco Firepower 2100 series.
Discontinued `high unmanaged disk usage` alerts.	7.0.6	Any	The Disk Usage health module no longer alerts with `high unmanaged disk usage`. After upgrade, you may continue to see these alerts until you either deploy health policies to managed devices (stops the display of alerts) or upgrade the devices (stops the sending of alerts). **Note** Versions 7.0–7.0.5, 7.1.x, 7.2.0–7.2.3, and 7.3.x continue to support these alerts. If your Firewall Management Center is running any of these versions, you may also continue to see alerts.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
New health modules.	7.0	Any	

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details

			We added the following health modules:
			• AMP connectivity and cloud status
			  • AMP Connection Status: Monitors AMP cloud connectivity from the Firewall Threat Defense.
			  • AMP Threat Grid Status: Monitors AMP Threat Grid cloud connectivity from the Firewall Threat Defense.
			  • Security Services Exchange Connection Status: Monitors security services exchange cloud connectivity from the Firewall Threat Defense.
			• Firewall Threat Defense and Firewall Management Center status
			  • FMC Access Configuration Changes: Monitors access configuration changes made directly on the Firewall Management Center.
			  • FTD HA Status: Monitors the active and standby Firewall Threat Defense HA pair and the sync status between the devices.
			  • FMC HA Status: Monitors the active and standby Firewall Management Center and the sync status between the devices. Replaces the HA Status module.
			• Network and traffic monitoring
			  • ASP Drop: Monitors the connections dropped by the data plane accelerated security path.
			  • Flow Offload: Monitors hardware flow offload statistics on the Firepower 9300 and 4100 platforms.
			  • Routing Statistics: Monitors both IPv4 and IPv6 route information from the Firewall Threat Defense.
			  • VPN Statistics: Monitors site-to-site and remote access VPN tunnel statistics.
			• Snort and intrusion detection statistics
			  • Advanced Snort Statistics: Monitors Snort statistics related to packet performance, flow counters, and flow events.
			  • Hit Count: Monitors the number of times a particular rule is hit on the access control policy.
			  • xTLS Counters: Monitors xTLS/SSL flows, memory and cache effectiveness.

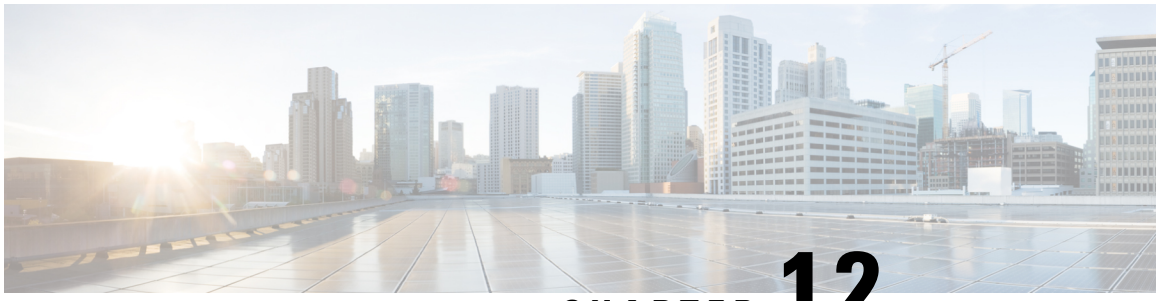
Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details

			• System and hardware health
			  • File System Integrity Check: Performs a file system integrity check if the system has CC mode or UCAPL mode enabled.
			  • MySQL Status: Monitors the status of the MySQL database.
			  • RabbitMQ Status: Monitors the status of the RabbitMQ messaging broker.
			  • Sybase Status: Monitors the status of the Sybase database.
			  • NTP Status FTD: Monitors the NTP clock synchronization status of the managed device.
			• Event and policy monitoring
			  • Event Stream Status: Monitors connections to third-party client applications that use the Event Streamer.
			  • Unresolved Groups Monitor: Monitors the unresolved groups used in access control policies.

Health monitor enhancements.	7.0	Any	The health monitor adds the following enhancements:
			• Enhanced Firewall Management Center dashboard with summary views of:
			  • High Availability
			  • Event Rate & Capacity
			  • Process Health
			  • CPU thresholds
			  • Memory
			  • Interface rates
			  • Disk Usage
			• Enhanced Firewall Threat Defense dashboard:
			  • Health alert for split brain scenario
			  • Additional health metrics available from new Health Modules

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
New health modules.	6.7	Any	The CPU Usage module is no longer used. Instead, see the following modules for CPU usage: • CPU Usage (per core): Monitors the CPU usage on all of the cores. • CPU Usage Data Plane: Monitors the average CPU usage of all data plane processes on the device. • CPU Usage Snort: Monitors the average CPU usage of the Snort processes on the device. • CPU Usage System: Monitors the average CPU usage of all system processes on the device. The following modules were added to track statistics: • Connection Statistics: Monitors the connection statistics and NAT translation counts. • Critical Process Statistics: Monitors the state of critical processes, their resource consumption, and the restart counts. • Deployed Configuration Statistics: Monitors statistics about the deployed configuration, such as the number of ACEs and IPS rules. • Snort Statistics: Monitors the Snort statistics for events, flows, and packets. The following modules were added to track memory usage: • Memory Usage Data Plane: Monitors the percentage of allocated memory used by the Data Plane processes. • Memory Usage Snort: Monitors the percentage of allocated memory used by the Snort process.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Health monitor enhancements.	6.7	Any	The health monitor adds the following enhancements: • Health Status summary page that provides an at-a-glance view of the health of the Firepower Management Center and all of the devices that the Firewall Management Center manages. • The Monitoring navigation pane allows you to navigate the device hierarchy. • Managed devices are listed individually, or grouped according to their geolocation, high availability, or cluster status where applicable. • You can view health monitors for individual devices from the navigation pane. • Custom dashboards to correlate interrelated metrics. Select from predefined correlation groups, such as CPU and Snort; or create a custom correlation dashboard by building your own variable set from the available metric groups.
Functionality moved to the Threat Data Updates on Devices module.	6.7	Any	The Local Malware Analysis module is no longer used. Instead, refer to the Threat Data Updates on Devices module for this information. Some information formerly provided by the Security Intelligence module and the URL Filtering Module is now provided by the Threat Data Updates on Devices module.
New health module: Configuration Memory Allocation.	7.0 6.6.3	Any	Version 6.6.3 improves device memory management and introduces a new health module: Configuration Memory Allocation. This module alerts when the size of your deployed configurations puts a device at risk of running out of memory. The alert shows you how much memory your configurations require, and by how much this exceeds the available memory. If this happens, re-evaluate your configurations. Most often you can reduce the number or complexity of access control rules or intrusion policies.
URL Filtering Monitor improvements.	6.5	Any	The URL Filtering Monitor module now alerts if the Firewall Management Center fails to register to the Cisco cloud.
URL Filtering Monitor improvements.	6.4	Any	You can now configure time thresholds for URL Filtering Monitor alerts.
New health module: Threat Data Updates on Devices.	6.3	Any	A new module, **Threat Data Updates on Devices**, was added. This module alerts you if certain intelligence data and configurations that devices use to detect threats has not been updated on the devices within the time period you specify.



CHAPTER 12

Audit and Syslog

The following topics describe how to audit activity on your system:

- System log, on page 423
- About system auditing, on page 425

System log

The System Log (syslog) page provides you with system log information for the appliance.

System audit capabilities

You can audit activity on your system in two ways. The appliances that are part of the system generate an audit record for each user interaction with the web interface, and also record system status messages in the system log.

The system log displays each message generated by the system. These items are listed in order:

- The date that the message was generated.
- The time that the message was generated.
- The host that generated the message.
- The message itself.

View the system log

View system log information that is local to your system to monitor and troubleshoot system events and activities.

System log information is local. For example, you **cannot** use the Firewall Management Center to view system status messages in the system logs on your managed devices.

You can filter messages using most syntax accepted by the UNIX file search utility Grep. This includes using Grep-compatible regular expressions for pattern matching.

Before you begin

To view system statistics, ensure you have Admin or Maintenance privileges in the Global domain.

Procedure

Step 1 Choose **System** (⚙) > **Monitoring** > **Syslog**.

Step 2 To search for specific message content in the system log:

a) Enter a word or query in the filter field as described in Syntax for system log filters, on page 424.

Only Grep-compatible search syntax is supported.

Examples:

- To search for all log entries that contain the user name "Admin," use `Admin`.
- To search for all log entries that are generated on November 27, use `Nov[[:space:]]*27` or `Nov.*27` (but not `Nov 27` or `Nov*27`).
- To search for all log entries that contain authorization debugging information on November 5, use `Nov[[:space:]]*5.*AUTH.*DEBUG`.

b) To make your search case-sensitive, select **Case-sensitive**.

By default, filters are not case-sensitive.

c) To search for all system log messages that do **not** meet the criteria you entered, select **Exclusion**.

d) Click **Go**.

Syntax for system log filters

This table shows the regular expression syntax you can use in System Log filters.

Table 29: System log filter syntax

Syntax Component	Description	Example
`.`	Matches any character or white space	`Admi.` matches `Admin`, `AdmiN`, `Admi1`, and `Admi`
`[[:alpha:]]`	Matches any alphabetic character	`[[:alpha:]]dmin` matches `Admin`, `bdmin`, and
`[[:upper:]]`	Matches any uppercase alphabetic character	`[[:upper:]]dmin` matches `Admin`, `Bdmin`, and
`[[:lower:]]`	Matches any lowercase alphabetic character	`[[:lower:]]dmin` matches `admin`, `bdmin`, and
`[[:digit:]]`	Matches any numeric character	`[[:digit:]]dmin` matches `0dmin`, `1dmin`, and
`[[:alnum:]]`	Matches any alphanumeric character	`[[:alnum:]]dmin` matches `1dmin`, `admin`, `2dmin`,
`[[:space:]]`	Matches any white space, including tabs	`Feb[[:space:]]29` matches logs from Februar
`*`	Matches zero or more instances of the character or expression it follows	`ab*` matches `a`, `ab`, `abb`, `ca`, `cab`, and `cabb` `[ab]*` matches anything
`?`	Matches zero or one instances	`ab?` matches `a` or `ab`

Syntax Component	Description	Example
\	Allows you to search for a character typically interpreted as regular expression syntax	`alert\?` matches `alert?`

About system auditing

The appliances that are part of the system generate an audit record for each user interaction with the web interface.

Related Topics

Standard reports, on page 531

Audit records

Firewall Management Center logs read-only auditing information for user activity. Audit logs appear in a standard event view that allows you to view, sort, and filter audit log messages based on any item in the audit view. You can easily delete audit information, generate reports, and review detailed records of the changes that users make.

Audit record storage and limitations

The audit log stores a maximum of 100,000 entries. When the number of audit log entries exceeds 100,000, the appliance prunes the oldest records from the database to reduce the number to 100,000.

The audit logs do not display the user or the source IP for login errors:

- When wrong password is used, the source IP is not displayed.
- When the user account does not exist, both source IP and the user are not displayed.
- If the attempt for an LDAP user fails, no audit log is triggered.

Related Topics

Guidelines for SSO configuration on Firewall Management Center, on page 170

View audit records

Monitor system activities and configuration changes by reviewing audit records.

Accessing audit records allows you to track events, system changes, and configuration activity for Firewall Management Center. You can use predefined workflows or create custom views that display only the information relevant to your needs.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

Ensure you have admin privileges.

Procedure

Step 1 Access the audit log from **System** (⚙) > **Monitoring** > **Audit**.

Step 2 If no events are visible, adjust the time range to include the desired events.

For more information, see Event time constraints, on page 692.

Note
Events that were generated outside the appliance's configured time window (whether global or event specific) may appear in an event view if you constrain the event view by time. This can occur even if you configured a sliding time window for the appliance.

Step 3 You have these choices:

The choices are applicable only based on the results of the search constraints. For example, when you search for **Health Events**, the resulting view page displays **Workflow** option. Similarly, only when you are in the **Vulnerabilities** table view, the option to view (**View** (👁)) specific vulnerability is displayed.

Action	Link to topic
To learn more about the contents of the columns in the table.	Refer to System log, on page 423.
To sort and constrain events on the current workflow page.	Refer to Use table view pages, on page 684.
To navigate between pages in the current workflow, keeping the current constraints, click the appropriate page link at the top left of the workflow page.	Refer to Use workflows, on page 676.
To drill down to the next page in the workflow.	Refer to Use drill-down pages, on page 684.
To constrain on a specific value, click a value within a row. If you click a value on a drill down page, you move to the next page and constrain on the value. Note that clicking a value within a row in a table view constrains the table view and does **not** drill down to the next page. **Tip** Each table view includes **Table View** in its page name.	See Event view constraints, on page 698 for more information.
To delete audit records, check the check boxes next to events you want to delete, then click **Delete**, or click **Delete All** to delete all events in the current constrained view.	-
To bookmark the current page, click **Bookmark This Page** so you can quickly return to it.	For more information, see Bookmarks, on page 703.

Action	Link to topic
To navigate to the bookmarked management page, click **View Bookmarks**.	For more information, see Bookmarks, on page 703.
To generate reports based on the data in the current view, click **Create Reports**.	For more information, see Create a report template from an event view, on page 536.
To compare system changes recorded in the audit log, click **Compare** next to applicable events in the **Message** column.	For more information, see Use audit log to examine changes, on page 428.

Related Topics

Event view constraints, on page 698

Audit log workflow fields

This table describes the audit log fields that can be viewed and searched.

Table 30: Audit log fields

Field	Description
Time	Time and date that the appliance generated the audit record.
User	User name of the user that triggered the audit event.
Subsystem	The full menu path the user followed to generate the audit record. For example, **System** (⚙) > **Monitoring** > **Audit** is the menu path to view the audit log. In a few cases where a menu path is not relevant, the Subsystem field displays only the event type. For example, **Login** classifies user login attempts.
Message	The action the user performed or the button the user clicked on the page. For example, `Page View` signifies that the user simply viewed the page indicated in the Subsystem, while `Save` means that the user clicked the **Save** button on the page. Changes made to the system appear with a **Compare icon** that you can click to see a summary of the changes.
Source IP	IP address associated with the host used by the user. Note: When searching this field you must type a specific IP address; you cannot use IP ranges when searching audit logs.
Domain	The current domain of the user when the audit event was triggered. This field is only present if you have ever configured the Firewall Management Center for multitenancy.
Configuration Change (search only)	Specifies whether to view audit records of configuration changes in the search results. (`yes` or `no`)

Field	Description
Count	The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows. This field is not searchable.

Related Topics

Event searches, on page 707

The audit events table view

You can customize the layout of the audit events table view or filter events by specific field value.

Use these options to customize the audit events table view:

- You can change the layout of the event view, or constrain the events in the view by a field value.
- To disable a column, click **Close** ( ) in the desired column heading. Click **Apply** in the pop-up window. When you disable a column, it is disabled for the duration of your session (unless you add it back later).

Note When you disable the first column, the Count column is added.

- To hide or show other columns, or to add a disabled column back to the view, select or clear the appropriate check boxes before you click **Apply**.
- Clicking a value within a row in a table view constrains the table view and does not drill down to the next page in the workflow.

Tip Table views always include "Table View" in the page name.

Related Topics

Use workflows, on page 676

Use audit log to examine changes

You can use the audit log to view detailed reports about changes made to your system. These reports show the current configuration and compare it to the most recent configuration before a supported change.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

You must have admin privileges to access and use the audit log features.

Procedure

Step 1 Choose **System** ( ) > **Monitoring** > **Audit**.

Step 2 In the audit log, find the event you want to review. Click **Compare** next to the log entry in the **Message** column.

Step 3 On the **Compare Configurations** page, review the highlighted differences between the previous and current configurations.

- Blue indicates that the highlighted setting is different in the two configurations. The difference is noted in red text.
- Green indicates that the highlighted setting appears in one configuration but not the other.
- The audit event type, timestamps, and user who made the change appear in the title bar for each configuration.

Tip
You can navigate through changes individually by clicking **Previous** or **Next** above the title bar. If the change summary is more than one page long, you can also use the scroll bar on the right to view additional changes.

Suppress audit records

Suppress audit records for specific types of user interactions that your auditing policy does not require monitoring.

If your auditing policy does not require that you audit specific types of user interactions with the system, you can prevent those interactions from generating audit records. For example, by default, each time a user views the online help, the system generates an audit record. If you do not need to keep a record of these interactions, you can automatically suppress them.

To configure audit event suppression, you must have access to an appliance's `admin` user account, and you must be able to either access the appliance's console or open a secure shell.



Caution Make sure that only authorized personnel have access to the appliance and to its `admin` account.

Before you begin

You must have admin privileges to perform this procedure.

Procedure

In the `/etc/sf` directory, create one or more `AuditBlock` files in the following form, where type is one of the types described in Audit block types, on page 430:

`AuditBlock.`*type*

Note

If you create an `AuditBlock.type` file for a specific type of audit message, but later decide that you no longer want to suppress them, you must delete the contents of the `AuditBlock.type` file, but leave the file itself on the system.

Audit block types

Each audit block type requires content in a specific format, as described in the following table. Ensure you use the correct capitalization for file names, because the contents are case sensitive.

Note When you add an `AuditBlock` file, an audit record with a subsystem of `Audit` and a message of `Audit Filter type Changed` is added to the audit events. For security reasons, this audit record **cannot** be suppressed.

Table 31: Audit block types

Type	Description
Address	Create a file named `AuditBlock.address` and include, one per line, each IP address that you want to suppress from the audit log. You can use partial IP addresses provided that they map from the beginning of the address. For example, the partial address `10.1.1` matches addresses from `10.1.1.0` through `10.1.1.255`.
Message	Create a file named `AuditBlock.message` and include, one per line, the message substrings that you want to suppress. Note that substrings are matched so that if you include `backup` in your file, all messages that include the word `backup` are suppressed.
Subsystem	Create a file named `AuditBlock.subsystem` and include, one per line, each subsystem that you want to suppress. Note that substrings are **not** matched. You must use exact strings. See Audited subsystems, on page 430 for a list of subsystems that are audited.
User	Create a file named `AuditBlock.user` and include, one per line, each user account that you want to suppress. You can use partial string matching provided that they map from the beginning of the username. For example, the partial username `IPSAnalyst` matches the user names `IPSAnalyst1` and `IPSAnalyst2`.

Audited subsystems

This table lists audited subsystems.

Table 32: Subsystem names

Name	Includes user interactions with...
Admin	Administrative features such as system and access configuration, time synchronization, backup and restore, device management, user account management, and scheduling

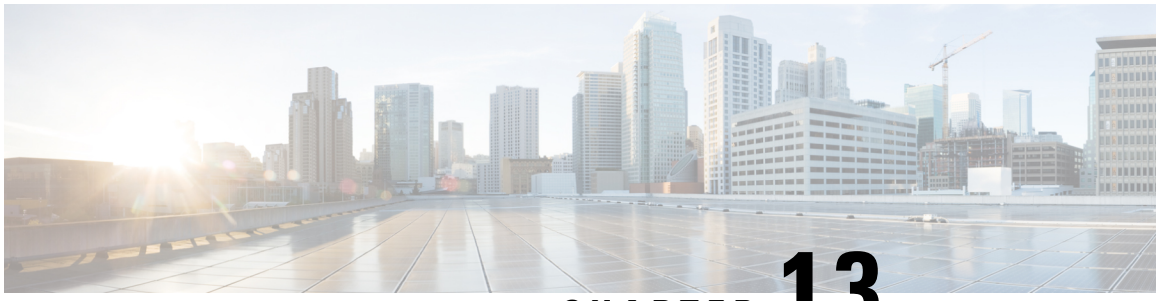
Name	Includes user interactions with...
Alerting	Alerting functions such as email, SNMP, and syslog alerting
Audit Log	Audit event views
Audit Log Search	Audit event searches
Command Line	Command line interface
Configuration	Email alerting
contextual cross-launch	External resources added to the system or accessed from dashboards and event views
COOP	Continuity of operations feature
Date	Date and time range for event views
Default Subsystem	Options that do not have assigned subsystems
Detection & Prevention Policy	Menu options for intrusion policies
Error	System-level errors
eStreamer	eStreamer configuration
EULA	Reviewing the end user license agreement
Events	Intrusion and discovery event views
Events Reviewed	Reviewed intrusion events
Events Search	Any event search
Failed to install rule update **rule_update_id**	Installing rule updates
Header	Initial presentation of the user interface after a user logs in
Health	Health monitoring
Health Events	Health monitoring event views
Help	Online help
High Availability	Establishing and handling Firewall Management Centers in high availability pairs
IDS Impact Flag	Impact flag configuration for intrusion events
IDS Policy	Intrusion policies
IDSRule sid:**sig_id** rev:**rev_num**	Intrusion rules by SID
Install	Installing updates

Name	Includes user interactions with...
Intrusion Events	Intrusion events
Login	Web interface login and logout functions
Logout	Web interface logout functions
Menu	Any menu option
Configuration export > **config_type** > **config_name**	Importing configurations of a specific type and name
Permission Escalation	User role escalation
Preferences	User preferences, such as the time zone for a user account and individual event preferences
Policy	Any policy, including intrusion policies
Register	Registering devices on a Firewall Management Center
RemoteStorageDevice	Configuring remote storage devices
Reports	Report listing and report designer features
Rules	Intrusion rules, including the intrusion rules editor and the rule importation process
Rule Update Import Log	Viewing the rule update import log
Rule Update Install	Installing rule updates
Session Expiration	Web interface session timeouts
Status	Syslog, as well as host and performance statistics
System	Various system-wide settings
Task Queue	Viewing background process status
Users	Creating and modifying user accounts and roles

About sending audit logs to an external location

To send audit logs to an external location from the Firewall Management Center, see:

- Audit logs, on page 53
- Audit log certificates, on page 57



CHAPTER 13

Statistics

The following topics describe how to monitor the system:

- System statistics, on page 433
- Host statistics section, on page 433
- Disk usage section, on page 434
- The processes section, on page 434
- SFDataCorrelator process statistics section, on page 439
- Intrusion event information section, on page 440
- View system statistics, on page 441

System statistics

The **Statistics** page lists the current status of general appliance statistics, including disk usage and system processes, data correlator statistics , and intrusion event information .

Host statistics section

The table lists the host statistics on the **Statistics** page.

Table 33: Host statistics

Category	Description
Time	The current time on the system.
Uptime	The number of days (if applicable), hours, and minutes since the system was last started.
Memory Usage	The percentage of system memory that is being used.
Load Average	The average number of processes in the CPU queue for the past 1 minute, 5 minutes, and 15 minutes.
Disk Usage	The percentage of the disk that is being used. Click the arrow to view more detailed host statistics.
Processes	A summary of the processes running on the system.

Disk usage section

The **Disk Usage** section of the **Statistics** page summarizes disk usage both by category and partition status. If a malware storage pack is installed on your device, you can also check its partition status within this section. Monitor this page regularly to ensure you have enough disk space for system processes and database operations.

Tip You can use the disk usage health monitor to monitor disk usage and receive alerts when the disk space is low.

The processes section

The **Processes** section of the **Statistics** page allows you to view the processes currently running on a managed appliance. This section provides both general and detailed information for every running process. You can use the Firewall Management Center's web interface to monitor the process status for any managed device.

There are two different types of processes that run on an appliance:

- Daemons: These processes always run in the background and are essential for system operations.
- Executable files: These processes run only when required, such as during specific tasks or operations.

Use the **Processes** section to quickly check which processes are active and understand their role in appliance management.

Process status fields

When you expand the **Processes** section of the **Statistics** page, you can also view the following:

CPU(s)

Lists the following CPU usage information:

- user process usage percentage
- system process usage percentage
- nice usage percentage (CPU usage of processes that have a negative nice value, indicating a higher priority). Nice values indicate the scheduled priority for system processes and can range between -20 (highest priority) and 19 (lowest priority).
- idle usage percentage

Mem

Lists the following memory usage information:

- total number of kilobytes in memory
- total number of used kilobytes in memory

- total number of free kilobytes in memory
- total number of buffered kilobytes in memory

Swap

Lists the following swap usage information:

- total number of kilobytes in swap
- total number of used kilobytes in swap
- total number of free kilobytes in swap
- total number of cached kilobytes in swap

The table describes each column that appears in the **Processes** section.

Table 34: Process list columns

Column	Description
Pid	The process ID number
Username	Name of the user or group running the process
Pri	The process priority
Nice	The value indicating the scheduling priority of a process. Values range from -20 (highest) to 19 (lowest)
Size	Memory size used by the process (in kilobytes, unless followed by `m` for megabytes)
Res	Amount of resident paging files in memory (in kilobytes, unless followed by `m` for megabytes)
State	The process state: • D: Uninterruptible sleep (usually Input/Output) • N: Positive nice value • R: Runnable (on queue to run) • S: Sleep mode • T: Being traced or stopped • W: Paging • X: Dead • Z: Defunct • <: Negative nice value
Time	Amount of time (in hours:minutes:seconds) that the process has been running
CPU	Percentage of CPU the process is using

Column	Description
Command	Executable name of the process

System daemons

Daemons continually run on an appliance. They ensure that services are available and spawn processes when required. The table lists daemons that you may see on the **Process Status** page and provides a brief description of their functionality.

Table 35: System daemons

Daemon	Description
crond	Manages the execution of scheduled commands (cron jobs)
dhclient	Manages dynamic host IP addressing
fpcollect	Manages the collection of client and server fingerprints
httpd	Manages the HTTP (Apache web server) process
httpsd	Manages the HTTPS (Apache web server with SSL) service; checks for SSL and valid ce authentication; runs in the background to provide secure web access to the appliance
keventd	Manages Linux kernel event notification messages
klogd	Manages the interception and logging of Linux kernel messages
kswapd	Manages Linux kernel swap memory
kupdated	Manages the Linux kernel update process, performing disk synchronization
mysqld	Manages database processes
ntpd	Manages the Network Time Protocol (NTP) process
pm	Manages all system processes, starts required processes, restarts any process that fails une
reportd	Manages reports
safe_mysqld	Manages safe mode operation of the database; restarts the database daemon if an error oc logs runtime information to a file
SFDataCorrelator	Manages data transmission
sfestreamer (Firewall Management Center only)	Manages connections to third-party client applications that use the Event Streamer
sfmgr	Provides the RPC service for remotely managing and configuring an appliance using an s connection to the appliance
SFRemediateD (Firewall Management Center only)	Manages remediation responses

Daemon	Description
sftimeserviced (Firewall Management Center only)	Forwards time synchronization messages to managed devices
sfmbservice	Provides access to the sfmb message broker process running on a remote appliance usi connection to the appliance. Used by health monitoring to send health events and alerts fr device to the Firewall Management Center.
sftroughd	Listens for connections on incoming sockets and invokes the correct executable (typic message broker, sfmb) to handle the request
sftunnel	Provides the secure communication channel for all processes requiring communication appliance
sshd	Manages the Secure Shell (SSH) process; runs in the background to provide SSH acce appliance
syslogd	Manages the system logging (syslog) process

Note This table does not include every possible process; additional daemons may be present depending on appliance configuration.

Executables and system utilities

There are a number of executables on the system that are launched by other processes or through user action. The table describes the executables and system utilities that may appear on the **Process Status** page.

Table 36: System executables and utilities

Executable	Description
awk	Executes programs written in the `awk` programming language
bash	GNU Bourne-Again Shell
cat	Reads files and writes content to standard output
chown	Changes user and group file permissions
chsh	Changes the default login shell
SFDataCorrelator (Firewall Management Center only)	Analyzes binary files created by the system to generate events, connection data, and network maps
cp	Copies files
df	Lists the amount of free space on the appliance
echo	Writes content to standard output

Executable	Description
egrep	Searches files and folders for specified input; supports extended set of regular expressions not supported in standard grep
find	Recursively searches directories for specified input
grep	Searches files and directories for specified input
halt	Stops the server
httpsdctl	Handles secure Apache Web processes
hwclock	Allows access to the hardware clock
ifconfig	Indicates the network configuration executable; ensures that the MAC address stays constant
iptables	Handles access restriction based on changes made to the Access Configuration page.
iptables-restore	Handles iptables file restoration
iptables-save	Handles saved changes to the iptables
kill	Ends a session and process
killall	Ends all sessions and processes
ksh	Public domain version of the Korn shell
logger	Provides a way to access the syslog daemon from the command line
md5sum	Prints checksums and block counts for specified files
mv	Moves (renames) files
myisamchk	Indicates database table checking and repairing
mysql	Indicates a database process; multiple instances may appear
openssl	Indicates authentication certificate creation
perl	Indicates a perl process
ps	Writes process information to standard output
sed	Edits one or more text files
sfheartbeat	Identifies a heartbeat broadcast, indicating that the appliance is active; used to maintain contact between a device and Firewall Management Center.
sfmb	Indicates a message broker process; handles communication between Firewall Management Centers and device.
sh	Public domain version of the Korn shell

Executable	Description
shutdown	Shuts down the appliance
sleep	Suspends a process for a specified number of seconds
smtpclient	Mail client that handles email transmission when email event notification functionality is enabled
snmptrap	Forwards SNMP trap data to the SNMP trap server specified when SNMP notification functionality is enabled
snort	Indicates that Snort is running
SSH	Indicates a Secure Shell (SSH) connection to the appliance
sudo	Indicates a sudo process, which allows users other than admin to run executables
top	Displays information about the top CPU processes **Note** The CPU usage output of this utility is a split-up of different types of usages of the CPU core. You must add both user and system processes usage to know the actual total CPU usage. For example, if the output of top command is: `%CPU(s): 76.6 us, 22.1 sy, 0.0 ni, 0.0 id, 0.0 wa, 0.0 hi, 1.3 si, 0.0 st` Here, 76.6% of CPU time is used by user processes, 22.1% of CPU time is used by system(kernel) processes. The total CPU usage is 98.7%. Thus, the CPU usage reported in this utility appears to be different from the Health Monitor dashboard. In addition, this utility uses a three seconds interval to calculate the CPU usage. Whereas, the management center health monitor uses one-second intervals.
touch	Changes the access and modification times of specified files
vim	Edits text files
wc	Performs line, word, and byte counts on specified files

SFDataCorrelator process statistics section

On the Firewall Management Center, you can view statistics about the data correlator and network discovery processes for the current day. The managed devices perform data acquisition, decoding, and analysis, and the network discovery process then correlates the data with the fingerprint and vulnerability databases. It then produces binary files, which are processed by the data correlator running on the Firewall Management Center. The data correlator analyzes the information from the binary files, generates events, and creates network maps.

The statistics that appear for network discovery and the data correlator are averages for the current day, using statistics gathered between 12:00 AM and 11:59 PM for each device.

The table describes the statistics displayed for the data correlator process.

Table 37: Data correlator process statistics

Category	Description
Events/Sec	Number of discovery events that the data correlator receives and processes per second.
Connections/Sec	Number of connections that the data correlator receives and processes per second.
CPU Usage — User (%)	Average percentage of CPU time spent on user processes for the current day.
CPU Usage — System (%)	Average percentage of CPU time spent on system processes for the current day.
VmSize (KB)	Average size of memory allocated to the data correlator for the current day, in kilobytes.
VmRSS (KB)	Average amount of memory used by the data correlator for the current day, in kilobytes.

These statistics help administrators monitor the performance and resource usage of the Firewall Management Center's data correlator and network discovery processes throughout the day.

Intrusion event information section

On both the Firewall Management Center and managed devices, you can view summary information about intrusion events on the **Statistics** page. This information includes the date and time of the last intrusion event, the total number of events that have occurred in the past hour and the past day, and the total number of events in the database.

Note The **Intrusion Event Information** section of the **Statistics** page reflects intrusion events stored locally on the managed device rather than those sent to the Firewall Management Center. If the managed device cannot store intrusion events, or if it is not configured to do so, no information appears in this section.

The table describes the statistics shown in the **Intrusion Event Information** section of the **Statistics** page.

Table 38: Intrusion event information

Statistic	Description
Last Alert Was	The date and time that the last event occurred.
Total Events Last Hour	The total number of events that occurred in the past hour.
Total Events Last Day	The total number of events that occurred in the past twenty-four hours.
Total Events in Database	The total number of events in the events database.

View system statistics

View system statistics to monitor the performance and resource utilization of the Firewall Management Center and its managed devices.

The display includes statistics for the Firewall Management Center and its managed devices.

Before you begin

- Make sure you have Admin or Maintenance user privileges.
- Ensure you are in the Global domain.

Follow these steps to view system statistics:

Procedure

Step 1 Choose **System** (⚙) > **Monitoring** > **Statistics**.

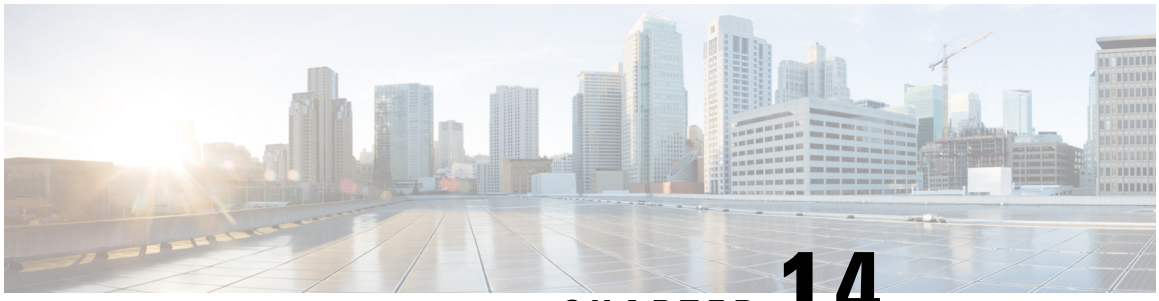
Step 2 Choose a device from the **Select Device(s)** list, and click **Select Devices**.

Step 3 View available statistics.

Step 4 In the Disk Usage section, you can:

- Hover your pointer over a disk usage category in the **By Category** stacked bar to view these items in this order:
  - the percentage of available disk space used by that category
  - the actual storage space on the disk
  - the total disk space available for that category
- Click the arrow next to **By Partition** to expand it. If a malware storage pack is installed, the **/var/storage** partition usage appears.

Step 5 (Optional) Click the arrow next to **Processes** to view the information described in View system statistics, on page 441.



CHAPTER 14

Troubleshooting

The following topics describe ways to diagnose problems you may encounter:

- Guidelines and best practices for troubleshooting, on page 443
- System messages, on page 444
- View basic system information, on page 446
- Manage system messages, on page 447
- Memory usage thresholds for health monitor alerts, on page 451
- Disk usage and drain of events health monitor alerts, on page 452
- Health monitor reports for troubleshooting, on page 456
- General troubleshooting, on page 458
- Connection-based troubleshooting, on page 458
- Advanced troubleshooting for the Secure Firewall Threat Defense device, on page 459
- Feature-specific troubleshooting, on page 467

Guidelines and best practices for troubleshooting

Follow these best practices to troubleshoot system issues effectively and prepare for potential support escalation:

- Before you make changes to try to fix a problem, generate a troubleshooting file to capture the original problem. See Health monitor reports for troubleshooting, on page 456 and its subsections.

  You may need this troubleshooting file if you need to contact Cisco TAC for support.

- Start your investigation by looking at error and warning messages in the Message Center. See System messages, on page 444

- Look for applicable Tech Notes and other troubleshooting resources under the "Troubleshoot and Alerts" heading on the product documentation page for your product.

- During the troubleshooting process, as several commands are executed simultaneously, the CPU usage becomes high. We recommend that you perform troubleshooting during periods of lower network traffic and fewer users.

- The time taken to troubleshoot the system depends on the configuration settings and the number of devices configured. This duration can range from nine to sixty minutes.

System messages

When you need to track down problems occurring in the system, the Message Center is the place to start your investigation. This feature allows you to view the messages that the system continually generates about system activities and status.

Message center access and display

To open the Message Center, click the System Status icon, next to the **Deploy** button on the home page. This icon changes based on the system status:

- **Error** ( ): Indicates one or more errors and any number of warnings are present on the system.
- **Warning** ( ): Indicates one or more warnings and no errors are present on the system.
- **Success** ( ): Indicates no warnings or errors are present on the system.

If a number is displayed with the icon, it indicates the total current number of error or warning messages.

To close the Message Center, click anywhere outside of it within the web interface.

The web interface also displays pop-up notifications in response to your actions and ongoing system activities. Some pop-up notifications automatically disappear after five seconds, while others are "sticky," meaning they display until you explicitly dismiss them by clicking **Dismiss** ( ). Click the **Dismiss** link at the top of the notifications list to dismiss all notifications at once.

Tip Hovering your cursor over a non-sticky pop-up notification causes it to be sticky.

The system determines which messages it displays to users in pop-up notifications and the Message Center based on their licenses, domains, and access roles.

Message types

The Message Center displays messages reporting system activities and status organized into three different tabs.

The Message Center organizes system status messages into these tabs:

- **Deployments**: Displays current status related to configuration deployment for each appliance in your system, grouped by domain.
- **Upgrades**: Displays the current status related to software upgrade tasks for the managed devices.
- **Health**: Displays current health status information for each appliance in your system, grouped by domain.
- **Tasks**: Displays the status of long-running tasks such as configuration backups, update installation, etc,.

Deployment status

The system reports these deployment statuses. You can get additional detail about the deployment jobs by clicking **Show History**.

- Running (**Spinning**): The configuration is in the process of deploying.
- **Success**: The configuration has successfully been deployed.
- **Warning** ( ) : Warning deployment statuses contribute to the message count displayed with the **Warning System Status icon**.
- **Failure**: The configuration has failed to deploy. Failed deployments contribute to the message count displayed with the **Error System Status icon**.

Upgrade status

The system reports these upgrade status values on the Upgrade tab:

- **In progress**: Indicates that the upgrade task is in progress.
- **Completed**: Indicates that the software upgrade task is completed successfully.
- **Failed**: Indicates that the software upgrade task has failed to complete.

Health status

This tab displays current health status information for each appliance in your system, grouped by domain. Health status is generated by health modules as described in Health monitoring, on page 379. The system reports the following health status values on this tab:

- **Warning** ( ): Indicates that warning limits have been exceeded for a health module on an appliance and the problem has not been corrected. The Health Monitoring page indicates these conditions with a **Yellow Triangle** ( ). Warning statuses contribute to the message count displayed with the **Warning System Status icon**.
- **Critical** ( ): Indicates that critical limits have been exceeded for a health module on an appliance and the problem has not been corrected. The Health Monitoring page indicates these conditions with a **Critical** ( ) icon. Critical statuses contribute to the message count displayed with the **Error System Status icon**.
- **Error** ( ): Indicates that a health monitoring module has failed on an appliance and has not been successfully re-run since the failure occurred. The Health Monitoring page indicates these conditions with an **Error icon**. Error statuses contribute to the message count displayed with the **Error System Status icon**.

You can click the links in the Health tab to view related detailed information on the Health Monitoring page. If there are no current health status conditions, the Health tab displays no messages.

Task status

Certain tasks (such as configuration backups or update installation) can require some time to complete. This tab displays the status of these long-running tasks, and can include tasks initiated by you or, if you have

appropriate access, other users of the system. The tab presents messages in reverse chronological order based on the most recent update time for each message. Some task status messages include links to more detailed information about the task in question. The system reports the following task status values on this tab:

- **Waiting**(): Indicates a task that is waiting to run until another in-progress task is complete. This message type displays an updating progress bar.
- **Running**: Indicates a task that is in-progress. This message type displays an updating progress bar.
- **Retrying**: Indicates a task that is automatically retrying. Note that not all tasks are permitted to try again. This message type displays an updating progress bar.
- **Success**: Indicates a task that has completed successfully.
- **Failure**: Indicates a task that did not complete successfully. Failed tasks contribute to the message count displayed with the **Error System Status icon**.
- **Stopped or Suspended**: Indicates a task that was interrupted due to a system update. Stopped tasks cannot be resumed. After normal operations are restored, start the task again.
- Skipped: A process in progress prevented the task from starting. Try again to start the task.

New messages appear in this tab as new tasks are started. As tasks complete (status success, failure, or stopped), this tab continues to display messages with final status indicated until you remove them. We recommend you remove messages to reduce clutter in the Tasks tab as well as the message database.

Message management

From the Message Center you can perform these tasks:

- Choose to display pop-up notifications.
- Display more task status messages from the system database (if any are available that have not been removed).
- Remove individual task status messages. (This affects all users who can view the removed messages.)
- Remove task status messages in bulk. (This affects all users who can view the removed messages.)

Tip Cisco recommends that you periodically remove accumulated task status messages from the Task tab to reduce clutter in both the display and the database. When the number of messages in the database approaches 100,000, the system automatically deletes task status messages that you have removed.

View basic system information

The About page displays information about your appliance, including the model, serial number, and version information for various components of the system. It also includes Cisco copyright information.

Procedure

Step 1 Click **Help** (?) in the toolbar at the top of the page.

Step 2 Choose **About**.

View appliance information

Procedure

Choose **System** (⚙) > **Configuration**.

Manage system messages

Use the Message Center to monitor and manage different types of system notifications including configuration deployments, device upgrades, system health alerts, and long-running tasks.

Procedure

Step 1 Click **Notifications** to display the Message Center.

Step 2 Choose from the available notification types:

- Click **Deployments** to view messages related to configuration deployments. See View deployment messages, on page 448. You must be an Admin user or have the **Deploy Configuration to Devices** permission to view these messages.
- Click **Upgrades** to view messages related to device upgrade tasks. See Viewing Upgrade Messages. See Viewing Upgrade Messages. You must be an Admin user or have **Updates** permission to view these messages.
- Click **Health** to view messages related to the health of your Firewall Management Center and the devices registered to it. See View health messages, on page 449. You must be an Admin user or have the **Health** permission to view these messages.

  You can navigate to the Health Monitor page by clicking the **Health monitor** link.

- Click **Tasks** to view or manage messages related to long-running tasks. See View task messages, on page 450 or Manage task messages, on page 450. Everyone can see their own tasks. To see the tasks of other users, you must be an Admin user or have the **View Other Users' Tasks** permission. You can remove the completed tasks from the notification by clicking the **Remove completed tasks** link.
- Click **Show Notifications** slider to enable or disable pop-up notification display.

View deployment messages

View deployment messages to monitor the status of configuration deployments and access detailed deployment history for troubleshooting.

If you are an Admin user or have the **Deploy Configuration to Devices** permission, you can view deployment messages and history.

Complete these steps to view deployment messages and deployment history:

Procedure

Step 1 Click **Notifications** to display the Message Center.

Step 2 Click **Deployments**.

Step 3 You have the following choices:

- Click **total** to view all current deployment statuses.
- Click a status value to view only messages with that deployment status.
- Hover your cursor over the time elapsed indicator for a message (for example, **1m 5s**) to view the elapsed time, and start and stop times for the deployment.

Step 4 Click **show deployment history** to view more detailed information about the deployment jobs.

The Deployment History table lists the deployment jobs in the left column in reverse chronological order.

a) Select a deployment job.

The table in the right column shows each device that was included in the job, and the deployment status per device.

b) To view responses from the device, and commands sent to the device during deployment, click download in the **Transcript** column for the device.

The transcript includes the following sections:

- **Snort Apply**: If there are any failures or responses from Snort-related policies, messages appear in this section. Normally, this section is empty.
- **CLI Apply**: This section covers features that are configured using commands sent to the Lina process.
- **Infrastructure Messages**: This section shows the status of different deployment modules.

In the **CLI Apply** section, the deployment transcript lists commands sent to the device as well as any responses returned from the device. These responses can be informative messages or error messages. If a deployment fails, review the error messages to understand which commands were problematic. If you use FlexConfig policies to configure features, these messages can help you identify and correct issues in your FlexConfig scripts.

Note
The transcript does not distinguish between commands sent for managed features and those generated from FlexConfig policies.

For example, this sequence shows that the Firewall Management Center sent commands to configure GigabitEthernet0/0 with the logical name outside. The device responded that it automatically set the security level to 0. The Firewall Threat Defense does not use the security level for anything.

```
========= CLI APPLY =========

FMC >> interface GigabitEthernet0/0
FMC >>  nameif outside
FTDv 192.168.0.152 >> [info] : INFO: Security level for "outside" set to 0 by default.
```

You can view deployment messages, monitor deployment statuses, and access detailed deployment history, including device transcripts to troubleshoot deployment issues.

View upgrade messages

View and monitor upgrade messages to track the status of current upgrade tasks and access detailed information about upgrade processes.

Only Admin users or users with the **Updates** permission can view these messages.

Procedure

Step 1 Click **Notifications** to display the Message Center.

Step 2 Click **Upgrades**.

Step 3 You can do the following:

- Click **total** to view all current upgrade tasks.
- Click on a status value to see messages with only that status.
- Click **Device Management** for more details on the upgrade task.

View health messages

Monitor system status, identify issues, and access detailed health information for troubleshooting and maintenance.

You must be an Admin user or have the **Health** permission to view health messages. Using the Message Center provides a streamlined way to monitor the system's health and filter messages by severity.

Complete these steps to view health messages in the Message Center:

Procedure

Step 1 Click **Notifications** to display the Message Center.

Step 2 Click **Health**.

Step 3 You have the following choices:

- Click **total** to view all current health statuses. You also see the severity levels: warning, critical, and error.
- Click on a status value to view only messages with that status.
- Hover your cursor over the relative time indicator for a message (for example, **3 day(s) ago**) to view the time of the most recent update for that message.
- To view detailed health status information for a particular message, click the message.
- To view complete health status on the Health Monitoring page, click **Health Monitor**.

You can now view, filter, and access detailed health status messages for monitoring and troubleshooting your system.

View task messages

View task status messages to monitor the progress and completion status of various system tasks.

You can view your tasks. But, to see the tasks of other users, you must be an Admin user or have the **View Other Users' Tasks** permission.

Procedure

Step 1 Click **Notifications** to display the Message Center.

Step 2 Click **Tasks**.

Step 3 You have these choices:

- Click **total** to view all current task statuses. To view tasks based on a specific status—such as waiting, running, retrying, success, or failure—click its corresponding label.
- Click a status value to view only messages for tasks with that status.

  Note
  Messages for stopped tasks appear only in the total list of task status messages. You cannot filter on stopped tasks.

- Hover your cursor over the relative time indicator for a message (e.g., **3 day(s) ago**) to view the time of the most recent update for that message.
- Click any link within a message to view more information about the task.
- If more task status messages are available for display, click **Fetch more messages** at the bottom of the message list to retrieve them.

Manage task messages

You can view your tasks. But, to see the tasks of other users, you must be an Admin user or have the **View Other Users' Tasks** permission.

Procedure

Step 1 Click System Status to display the Message Center.

Step 2 Click **Tasks**.

Step 3 Select one of the available management options:

- If more task status messages are available for display, click **Fetch more messages** at the bottom of the message list to retrieve them.
- To remove a single message for a completed task (status stopped, success, or failure), click **Remove** ( ) next to the message.
- To remove all messages for all tasks that have completed (status stopped, success, or failure), filter the messages on **total** and click **Remove all completed tasks**.
- To remove all messages for all tasks that have completed successfully, filter the messages on **success**, and click **Remove all successful tasks**.
- To remove all messages for all tasks that have failed, filter the messages on **failure**, and click **Remove all failed tasks**.

Memory usage thresholds for health monitor alerts

The memory usage health module compares memory usage on an appliance to the limits configured for the module and alerts when usage exceeds the levels. The module monitors data from managed devices and from the Firewall Management Center itself.

Two configurable thresholds for memory usage, critical and warning, can be set as a percentage of memory used. When these thresholds are exceeded, a health alarm is generated with the severity level specified. However, the health alarm system does not calculate these thresholds in an exact manner.

With high memory devices, certain processes are expected to use a larger percentage of total system memory than in a low memory footprint device. The design is to use as much of the physical memory as possible while leaving a small value of memory free for ancillary processes.

Compare two devices, one with 32 GB of memory and one with 4 GB of memory. In the device with 32 GB of memory, 5% of memory (1.6GB) is a much larger value of memory to leave for ancillary processes than in the device with 4 GB of memory (5% of 4GB = 200MB).

To account for the higher percentage use of system memory by certain processes, the Firewall Management Center calculates the total memory to include both total physical memory and total swap memory. Thus the enforced memory threshold for the user-configured threshold input can result in a Health Event where the **Value** column of the event does not match the value that was entered to determine the exceeded threshold.

The following table shows examples of user-input thresholds and the enforced thresholds, depending on the installed system memory.

Note The values in this table are examples. You can use this information to extrapolate thresholds for devices that do not match the installed RAM shown here, or you can contact Cisco TAC for more precise threshold calculations.

Table 39: Memory usage thresholds based on installed RAM

User-Input Threshold Value	Enforced Threshold Per Installed Memory (RAM)			
	4 GB	6 GB	32 GB	48 GB
10%	10%	34%	72%	81%
20%	20%	41%	75%	83%
30%	30%	48%	78%	85%
40%	40%	56%	81%	88%
50%	50%	63%	84%	90%
60%	60%	70%	88%	92%
70%	70%	78%	91%	94%
80%	80%	85%	94%	96%
90%	90%	93%	97%	98%
100%	100%	100%	100%	100%

Disk usage and drain of events health monitor alerts

This topic describes the symptoms and troubleshooting guidelines for two health alerts generated by the Disk Usage health module:

- Frequent Drain of Events
- Drain of Unprocessed Events

Disk usage overview

The disk usage health module compares disk usage on a managed device's hard drive and malware storage pack to the limits configured for the module and alerts when usage exceeds the percentages configured for the module. This module also alerts when the system excessively deletes files in monitored disk usage categories, or when disk usage excluding those categories reaches excessive levels, based on module thresholds.

The disk manager process manages the disk usage of a device. Each type of file monitored by the disk manager is assigned with a silo. Based on the amount of disk space available on the system the disk manager computes a High Water Mark (HWM) and a Low Water Mark (LWM) for each silo.

To display detailed disk usage information for each part of the system, including silos, LWMs, and HWMs, use the **show disk-manager** command.

Examples

The following is an example of the disk manager information:

```
> show disk-manager
Silo                                Used        Minimum     Maximum
Temporary Files                     0 KB        499.197 MB  1.950 GB
Action Queue Results                0 KB        499.197 MB  1.950 GB
User Identity Events                0 KB        499.197 MB  1.950 GB
UI Caches                           4 KB        1.462 GB    2.925 GB
Backups                             0 KB        3.900 GB    9.750 GB
Updates                             0 KB        5.850 GB    14.625 GB
Other Detection Engine              0 KB        2.925 GB    5.850 GB
Performance Statistics              33 KB       998.395 MB  11.700 GB
Other Events                        0 KB        1.950 GB    3.900 GB
IP Reputation & URL Filtering       0 KB        2.437 GB    4.875 GB
Archives & Cores & File Logs        0 KB        3.900 GB    19.500 GB
Unified Low Priority Events         1.329 MB    4.875 GB    24.375 GB
RNA Events                          0 KB        3.900 GB    15.600 GB
File Capture                        0 KB        9.750 GB    19.500 GB
Unified High Priority Events        0 KB        14.625 GB   34.125 GB
IPS Events                          0 KB        11.700 GB   29.250 GB
```

Health alert format

When the Health Monitor process on the Firewall Management Center runs (once every 5 minutes or when a manual run is triggered), the Disk Usage module looks into the diskmanager.log file and, if the correct conditions are met, the health alert is triggered.

The structures of these health alerts are as follows:

- Frequent drain of *SILO NAME*
- Drain of unprocessed events from *SILO NAME*

For example,

- Frequent drain of Low Priority Events
- Drain of unprocessed events from Low Priority Events

Its possible for any silo to generate a *Frequent drain of SILO NAME* health alert. However, the most commonly seen are the alerts related to events. Among the event silos, the *Low Priority Events* are often seen because device generates this type of events frequently.

A *Frequent drain of <SILO NAME>* event has a **Warning** severity level when seen in relation to an event-related silo, because events will be queued to be sent to the Firewall Management Center. For a non-event related silo, such as the *Backups* silo, the alert has a **Critical** severity level because this information is lost.

Note Only event silos generate a *Drain of unprocessed events from SILO NAME* health alert. This alert always has a **Critical** severity level.

Additional symptoms besides the alerts can include:

- Slowness on the Firewall Management Center user interface
- Loss of events

Common troubleshoot scenarios

A *Frequent drain of SILO NAME* event is caused by too much input into the silo for its size. In this case, the disk manager drains (purges) that file at least twice in the last 5-minute interval. In an event type silo, this is typically caused by excessive logging of that event type.

A *Drain of unprocessed events of SILO NAME* health alert is caused by a bottleneck in the event processing path.

There are three potential bottlenecks with respect to these Disk Usage alerts:

- Excessive logging — The EventHandler process on Firewall Threat Defense is oversubscribed (it reads slower than what Snort writes).
- Sftunnel bottleneck — The Eventing interface is unstable or oversubscribed.
- SFDataCorrelator bottleneck — The data transmission channel between the Firewall Management Center and the managed device is oversubscribed.

Excessive logging

One of the most common causes for the health alerts of this type is excessive input. The difference between the Low Water Mark (LWM) and High Water Mark (HWM) gathered from the **show disk-manager** command shows how much space there is available to take on that silo to go from LWM (freshly drained) to the HWM value. If there are frequent drain of events (with or without unprocessed events), review the logging configuration.

- Check for double logging — Double logging scenarios can be identified if you look at the correlator *perfstats* on the Firewall Management Center:

  ```
  admin@FMC:~$ sudo perfstats -Cq < /var/sf/rna/correlator-stats/now
  ```

- Check logging settings for the ACP — Review the logging settings of the Access Control Policy (ACP). If the logging setting includes both "Beginning" and "End" of connection, modify the setting to log only the end to reduce the number of events.

  Ensure that you follow the best practices described in Best practices for connection logging, on page 747.

Communications Bottleneck – Sftunnel

Sftunnel is responsible for encrypted communications between the Firewall Management Center and the managed device. Events are sent over the tunnel to the Firewall Management Center. Connectivity issues and/or instability in the communication channel (sftunnel) between the managed device and the Firewall Management Center can be due to:

- Sftunnel is down or is unstable (flaps).

  Ensure that the Firewall Management Center and the managed device have reachability between their management interfaces on TCP port 8305.

The sftunnel process should be stable and should not restart unexpectedly. Verify this by checking the `/var/log/message` file and search for messages that contain the **sftunneld** string.

- Sftunnel is oversubscribed.

  Review trend data from the Heath Monitor and look for signs of oversubscription of the Firewall Management Center's management interface, which can be a spike in management traffic or a constant oversubscription.

  Use as a secondary management interface for eventing. To use this interface, you must configure its IP address and other parameters at the Firewall Threat Defense CLI using the **configure network management-interface** command.

Communications Bottleneck – SFDataCorrelator

The SFDataCorrelator manages data transmission between the Firewall Management Center and the managed device; on the Firewall Management Center, it analyzes binary files created by the system to generate events, connection data, and network maps. The first step is to review the **diskmanager.log** file for important information to be gathered, such as:

- The frequency of the drain.
- The number of files with Unprocessed Events drained.
- The occurrence of the drain with Unprocessed Events.

Each time the disk manager process runs it generates an entry for each of the different silos on its own log file, which is located under **[/ngfw]/var/log/diskmanager.log**. Information gathered from the diskmanager.log (in CSV format) can be used to help narrow the search for a cause.

Additional troubleshooting steps:

- The command **stats_unified.pl** can help you to determine if the managed device does have some data which must be sent to Firewall Management Center. This condition can happen when the managed device and the Firewall Management Center experience a connectivity issue. The managed device stores the log data on to a hard drive.

  ```
  admin@FMC:~$ sudo stats_unified.pl
  ```

- The **manage_proc.pl** command can reconfigure the correlator on the Firewall Management Center side.

  ```
  root@FMC:~# manage_procs.pl
  ```

Before you contact Cisco TAC

It is highly recommended to collect these items before you contact Cisco TAC:

- Screenshots of the health alert seen.
- Troubleshoot file generated from the Firewall Management Center.
- Troubleshoot file generated from the affected managed device.
- Date and Time when the problem was first seen.
- Information about any recent changes done to the policies (if applicable).

- The output of the stats_unified.pl command as described in Communications Bottleneck — SFDataCorrelator, on page 455.

Disk usage for device configuration history files

The Disk Usage health module monitors the size of the device configuration history files on the Firewall Management Center and sends a health alert if the size exceeds the allowed limit. The maximum disk size allowed for storing the device configuration history files is 20 GB. In a Firewall Management Center high-availability deployment, this health alert appears on the standby Firewall Management Center only when high-availability synchronization is paused.

The size of device configuration history files exceeding the allowed limit can cause upgrade readiness failure while upgrading the Firewall Management Center. In a Firewall Management Center high-availability deployment, exceeding the device configuration history file size limit can reduce high-availability synchronization speed.

To resolve the device configuration history files size health alert, choose **Deploy** > and click **Deployment Setting**. In the **Deployment Setting** page, reduce the **Number of Versions to Retain**. Reducing the number of versions removes the oldest configuration versions to match the version size you have selected. The **Estimated Configuration Version Size** provides the approximate size for the configuration history files on the Firewall Management Center, based on the number of versions you choose to retain. Change the number of versions using the estimated value to decrease the size of configuration versions below the allowed limit.

For more information, refer to Set the Number of Configuration Versions in *Cisco Secure Firewall Management Center Device Configuration Guide*.

Health monitor reports for troubleshooting

In some cases, if you have a problem with your appliance, support may ask you to supply troubleshooting files to help them diagnose the problem. The system can produce troubleshooting files with information targeted to specific functional areas, as well as advanced troubleshooting files you retrieve in cooperation with Support. You can select any of the options listed in the table below to customize the contents of a troubleshooting file for a specific function.

Note Some options overlap in terms of the data they report, but the troubleshooting files will not contain redundant copies, regardless of what options you select.

Table 40: Selectable troubleshoot options

This option...	Reports...
Snort Performance and Configuration	data and configuration settings related to Snort on the appliance
Hardware Performance and Logs	data and logs related to the performance of the appliance hardware
System Configuration, Policy, and Logs	configuration settings, data, and logs related to the current system configuration of the appliance
Detection Configuration, Policy, and Logs	configuration settings, data, and logs related to detection on the appliance

This option...	Reports...
Interface and Network Related Data	configuration settings, data, and logs related to inline sets and network configuration of the appliance
Discovery, Awareness, VDB Data, and Logs	configuration settings, data, and logs related to the current discovery and awareness configuration on the appliance
Upgrade Data and Logs	data and logs related to prior upgrades of the appliance
All Database Data	all database-related data that is included in a troubleshoot report
All Log Data	all logs collected by the appliance database
Network Map Information	current network topology data

Generate troubleshooting files for specific system functions

You can generate and download customized troubleshooting files that you can send to support.

Before you begin

You must be an Admin, Maintenance, Security Analyst, or Security Analyst (Read Only) user to perform this task.

Procedure

Step 1 Perform the steps in View the device health monitor, on page 408.

Step 2 Choose **System** (⚙) > **Health** > **Monitor**, click the device in the left panel, then **View System & Troubleshoot Details**, and then click **Generate Troubleshooting Files**.

Note

- Troubleshooting files generated from the Firewall Management Center web interface are stored on the Firewall Management Center. Only the latest troubleshooting file for each appliance is stored.
- Troubleshooting files generated from the CLI are stored locally and are not overwritten.
- If you initiate the troubleshooting log generation while one is already in the queue, the new request is dropped.

Step 3 Choose **All Data** to generate all possible troubleshooting data, or check individual boxes as described in View task messages, on page 450.

Step 4 Click **Generate**.

Step 5 View task messages in the Message Center; see View task messages, on page 450.

Step 6 Find the task that corresponds to the troubleshooting files you generated.

Step 7 After the appliance generated the troubleshooting files and the task status changes to `Completed`, click **Click to retrieve generated files**.

Step 8 Follow your browser's prompts to download the file. (The troubleshooting files are downloaded in a single `.tar.gz` file.)

Step 9 Follow the directions from Support to send the troubleshooting files to Cisco.

Download advanced troubleshooting files

Download troubleshooting files and share them with Cisco support for advanced troubleshooting assistance when requested.

Before you begin

You must have **Admin**, **Maintenance**, **Security Analyst**, or **Security Analyst (Read Only)** user roles to perform this task.

Procedure

Step 1 View the health monitor for the appliance; refer to , View the device health monitor, on page 408.

Step 2 Choose **System** (⚙) > **Health** > **Monitor**.

Step 3 Click a device in the left pane for which you dowload the troubleshooting files.

Step 4 Click **View System & Troubleshoot Details**, and then click **Advanced Troubleshooting.**

Step 5 In **File Download** field, enter the name of the file to download. Cisco support provides this file name.

Step 6 Click **Download**.

Step 7 Follow your browser's prompts to download the file.

Note
For managed devices, the system renames the file by prepending the device name to the file name.

Step 8 Follow the directions from Cisco support to send the troubleshooting files to Cisco.

General troubleshooting

An internal power failure (hardware failure, power surge, and so on) or an external power failure (unplugged cord) can result in an ungraceful shutdown or reboot of the system. This can result in data corruption.

Connection-based troubleshooting

Connection-based troubleshooting or debugging provides uniform debugging across modules to collect appropriate logs for a specific connection. It also supports level-based debugging up to seven levels and enables uniform log collection mechanism across modules. Connection-based debugging supports the following:

- A common connection-based debugging subsystem to troubleshoot issues in Firewall Threat Defense
- Uniform format for debug messages across modules
- Persistent debug messages across reboots

- End-to-end debugging across modules based on an existing connection
- Debugging ongoing connections

Note Connection-based debugging is not supported on Firepower 2100 Series devices.

For more information about the troubleshooting connections, see Troubleshoot a connection, on page 459.

Troubleshoot a connection

This task enables you to identify and analyze connection issues by capturing and examining packet debug information for specific traffic flows.

Procedure

Step 1 Configure a filter to identify a connection using the **debug packet-condition** command.

Example:

```
Debug packet-condition match tcp 192.168.100.177 255.255.255.255 192.168.102.177
255.255.255.255
```

Step 2 Enable debugs for the interested modules and the corresponding levels. Enter the **debug packet** command.

Example:

```
Debug packet acl 5
```

Step 3 Start debugging the packets using the following command:

```
debug packet-start
```

Step 4 Fetch the debug messages from database to analyze the debug messages using the following command:

```
show packet-debugs
```

Step 5 Stop debugging the packets using the following command:

```
debug packet-stop
```

Advanced troubleshooting for the Secure Firewall Threat Defense device

You can use Packet Tracer and Packet Capture features to perform an in-depth troubleshooting analysis on a Secure Firewall Threat Defense device. A packet tracer allows a firewall administrator to inject a virtual packet into a security appliance and track the flow from ingress to egress. Along the way, the packet is evaluated against flow and route lookups, ACLs, protocol inspection, NAT, and intrusion detection. This utility is effective because it can simulate real-world traffic by specifying source and destination addresses with protocol

and port information. Packet capture is available with the trace option, which provides you with a verdict as to whether the packet is dropped or successful.

For more information about the troubleshooting files, see Download advanced troubleshooting files, on page 458.

Packet capture

The packet capture feature with trace option allows real packets that are captured on the ingress interface to be traced through the system. The trace information is displayed at a later stage. These packets are not dropped on the egress interface, as they are real data-path traffic. Packet capture for Firewall Threat Defense devices supports troubleshooting and analysis of data packets.

Snort verdict processing

Once the packet is acquired, Snort detects the tracing flag that is enabled in the packet. Snort writes tracer elements, through which the packet traverses. Snort verdict as a result of capturing packets can be one of the following:

Table 41: Snort verdicts

Verdict	Description
Pass	Allow analyzed packet.
Block	Packet not forwarded.
Replace	Packet modified.
AllowFlow	Flow passed without inspection.
BlockFlow	Flow was blocked.
Ignore	Flow was blocked; occurs only for sessions with flows blocked on passive interfaces.
Retry	Flow is stalled, waiting on a enamelware or URL category/reputation query. In the event of a timeout, processing continues with an unknown result: in the case of enamelware, the file is allowed; in the case of URL category/reputation, AC rule lookup continues with an uncategorized and unknown reputation.

Based on the Snort verdict, the packets are dropped or allowed. For example, the packet is dropped if the Snort verdict is **BlockFlow**, and the subsequent packets in the session are dropped before reaching Snort. When the Snort verdict is **Block** or **BlockFlow**, the **Drop Reason** can be one of these:

Table 42: Drop reasons

Blocked or Flow Blocked by...	Cause
Snort	Snort is unable to process the packet, erg., snort can't decode packet since it is corrupted or has invalid format.

Blocked or Flow Blocked by...	Cause
the App Id preprocessed	App Id module/preprocessed does not block packet by itself; but this may indicate that App Id detection causes other module (erg., firewall) to match a blocking rule.
the SSL preprocessed	There is a block/reset rule in SSL policy to match the traffic.
the firewall	There is a block/reset rule in firewall policy to match the traffic.
the captive portal preprocessed	There is a block/reset rule using the identity policy to match the traffic.
the safe search preprocessed	There is a block/reset rule using the safe-search feature in firewall policy to match the traffic.
the SI preprocessed	There is a block/reset rule a in Security Intelligence tab of AC Policy to block the traffic, erg., DNS or URL SI rule.
the filterer preprocessed	There is a block/reset rule in filterer tab of AC policy to match the traffic.
the stream preprocessed	There is an intrusion rule blocking/reset stream connection, erg., blocking when TCP normalization error.
the session preprocessed	This session was already blocked earlier by some other module, so session preprocessed is blocking further packets of the same session.
the fragmentation preprocessed	Blocking because earlier fragment of the data is blocked.
the snort response preprocessed	There is a react snort rule, erg., sending a response page on a particular HTTP traffic.
the snort response preprocessed	There is a snort rule to send custom response on packets matching conditions.
the reputation preprocessed	Packet matches a reputation rule, erg., blocking a given IP address.
the x-Link2State preprocessed	Blocking due to buffer overflow vulnerability detected in SMTP.
back orifice preprocessed	Blocking due to detection of back orifice data.
the SMB preprocessed	There is a snort rule to block SMB traffic.
the file process preprocessed	There is file policy that blocks a file, erg., enamelware blocking.

Blocked or Flow Blocked by...	Cause
the IPS preprocessed	There is a snort rule using IPS, erg., rate filtering.

The packet capture feature allows you to capture and download packets that are stored in the system memory. However, the buffer size is limited to 32 MB due to memory constraint. Systems capable of handling very high volume of packet captures exceed the maximum buffer size quickly and thereby the necessity of increasing the packet capture limit is required. It is achieved by using the secondary memory (by creating a file to write the capture data). The maximum supported file size is 10 GB.

When the **file-size** is configured, the captured data gets stored to the file and the file name is assigned based on the capture name **recapture**.

The **file-size** option is used when you need to capture packets with the size limit more than 32 MB.

For information, see the *Cisco Secure Firewall Threat Defense Command Reference*.

Use a packet capture trace

Packet capture is a utility that provides a live snapshot of network traffic passing the specified interface of a device based on a defined criteria. This process continues to capture the packets as long as it has not paused, or the allocated memory has not exhausted.

Packet capture data includes information from Snort and preprocessors about verdicts and actions the system takes while processing a packet. Multiple packet captures are possible at a time. You can configure the system to modify, delete, clear, and save captures.

Note Capturing packet data requires packet copy. This operation may cause delays while processing packets and may also degrade the packet throughput. We recommend that you use packet filters to capture specific traffic data.

Before you begin

To use the packet capture tool on Secure Firewall Threat Defense devices, you must be an Admin or Maintenance user.

Procedure

Step 1 In the Firewall Management Center, choose **Devices** > **Troubleshoot** > **Packet Capture**.

Step 2 Select a device.

Step 3 Click **Add Capture**.

Step 4 Enter the **Name** for capturing the trace.

Step 5 Select the **Interface** for the capturing the trace.

Step 6 Specify **Match Criteria** details:

a) Select the **Protocol**.

b) Enter the IP address for the **Source Host**.

c) Enter the IP address for the **Destination Host**.

d) Check **SGT number** check box, and enter a Security Group Tag (SGT).

Step 7 Specify **Buffer** details:

a) Enter a maximum **Packet Size**.
b) Enter a minimum **Buffer Size**.
c) Select either **Continuous Capture** if you want the traffic captured without interruption, or **Stop when full** if you want the capture to stop when the maximum buffer size is reached.

Note
If **Continues Capture** is enabled, and when the allocated memory is full, the oldest captured packets in the memory is overwritten by the new captured packets.

d) Check the check box of **Trace**, if you want to capture the details for each packet.
e) Enter the value in **Trace Count** field. Default value is 50. You can enter values in the range of 1-1000.

Step 8 Click **Save**.

The packet capture screen displays the packet capture details and its status. To have the packet capture page auto refreshed, check the **Enable Auto Refresh** check box and enter the auto refresh interval in seconds.

You can do the following on the packet capture:

- **Edit** ( ) to modify the capture criteria.
- **Delete** ( ) to delete the packet capture and the captured packets.
- **Clear** ( ) to erase all the captured packets from a Packet Capture. To erase the captured packets from all of the existing packet captures, click **Clear All Packets**.
- **Pause** ( ) to temporarily halt capturing packets.
- **Save** ( ) to save a copy of captured packets on a local machine in ASCII or PCAP format. Choose the required format option, and click **Save**. The saved packet capture is downloaded to your local machine.
- To view the details of the packets being captured, click the required capture row.

Packet tracer tool overview

A packet tracer tool is a network analysis utility. It simulates packet flows by modeling source and destination addresses and protocol characteristics. The tool validates network policies by performing lookups against access rules, NAT, routing, and rate-limiting policies. It also captures and logs detailed tracing information for each packet to assist with policy troubleshooting and debugging.

Packet tracer tool simulation and PCAP replay

Use the packet tracer tool to simulate the complete data path of a packet, including slow-path and fast-path modules. This simulation helps verify whether your configured policies allow or deny intended traffic. This tool helps you confirm your policy effectiveness and troubleshoot when your packets are blocked.

You can use packet capture (PCAP) files to replay complete traffic flows. Make sure your PCAP file contains a single Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) flow with up to 100 packets. On replay, the tool initializes client and server entities and synchronizes packet transmission, logging trace results for each packet in the flow.

PCAP Replay

Packet replay operates in sequence as captured in the PCAP file. Any interference during replay ends the session. Trace output is generated for all packets on the specified ingress and egress interfaces, supporting comprehensive flow analysis. Some features that dynamically alter packets during replay—such as IPsec, VPN, , HTTPS decryption, , and others—are not supported in PCAP replay.

Run a packet trace

Follow these steps to run a packet trace on your Firewall Threat Defense device.

Before you begin

- Ensure you have Admin or Maintenance user privileges.
- Have any required PCAP (.pcap or .pcapng) files ready if you intend to upload or replay captured traffic.
- Confirm the device supports packet tracing.

Procedure

Step 1 In Firewall Management Center, choose **Devices** > **Troubleshoot** > **Packet Tracer**.

Step 2 From the **Select Device** drop-down, choose the device where you want to run the trace.

Step 3 From the **Ingress Interface** drop-down, choose the ingress interface for the packet trace.

Note
Do not select VTI. VTI as ingress interface is not supported for packet tracer.

Step 4 To use a PCAP replay in the packet tracer, do the following:

a) Click **Select a PCAP File**.

b) To upload a new PCAP file, click **Upload a PCAP file**. To reuse a recently uploaded file, click the file name from the list.

Note
Only .pcap and .pcapng file formats are supported. The PCAP file can contain only a single TCP/UDP-based flow with a maximum of 100 packets. The maximum character limit for the PCAP file name (including the file formats) is 64.

c) In the **Upload PCAP** box, you can either drag a PCAP file or click to browse to the location where the file is stored, and select the file. When you select the file, the upload process starts automatically.

d) Go to this step

Step 5 To define the trace parameters, from the **Protocol** drop-down list, select the packet type for the trace, and specify the protocol characteristics:

- **ICMP** : Enter the ICMP type, ICMP code (0-255), and optionally, the ICMP identifier.
- **TCP** / **UDP** / **SCTP** : Enter the source and destination port numbers.
- **GRE** / **IPIP** : Enter the protocol number, 0-255.
- **ESP** : Enter the Security Parameter Index (SPI) value for the source. Valid range is 0-4294967295.
- **RAWIP** : Enter the port number. Valid range is 0-255.

Step 6 Select the **Source Type** for the packet trace, and enter the source IP address.

Source and destination types include IPv4, IPv6, and fully-qualified domain names (FQDN). You can specify IPv4 or IPv6 addresses and FQDN, if you use Cisco TrustSec.

Step 7 Select the **Source Port** for the packet trace.

Step 8 Select the **Destination** type for the packet trace, and enter the destination IP address.

Destination type options vary depending on the source type that you select.

Step 9 Select the **Destination Port** for the packet trace.

Step 10 Optionally, if you want to trace a packet where the Security Group Tag (SGT) value is embedded in the Layer 2 CMD header (TrustSec), enter a valid **SGT number** .

Step 11 If you want packet tracer to enter a parent interface, which is later redirected to a sub-interface, enter a **VLAN ID** .

This value is optional for non-sub-interfaces only, since all the interface types can be configured on a sub-interface.

Step 12 Specify a **Destination MAC Address** for the packet trace.

If the Secure Firewall Threat Defense device is running in transparent firewall mode, and the ingress interface is VTEP, **Destination MAC Address** is required if you enter a value in **VLAN ID**. Whereas if the interface is a bridge group member, **Destination MAC Address** is optional if you enter a **VLAN ID** value, but required if you do not enter a **VLAN ID** value.

If the Secure Firewall Threat Defense is running in routed firewall mode, **VLAN ID** and **Destination MAC Address** are optional if the input interface is a bridge group member.

Step 13 (Optional) If you want the packet-tracer to ignore the security checks on the simulated packet, click **Bypass all security checks for simulated packet** . This enables packet-tracer to continue with tracing of packet through the system which, otherwise would have been dropped.

Step 14 (Optional) To allow the packet to be sent out through the egress interface from the device, click **Allow simulated packet to transmit from device** .

Step 15 (Optional) If you want the packet-tracer to consider the injected packet as an IPsec/SSL VPN decrypted packet, click **Treat simulated packet as IPsec/SSL VPN decrypt** .

Step 16 Click **Trace**.

The **Trace Result** displays the results for each phase that the PCAP packets have traveled through the system. Click an individual packet to view the traces results for the packet. You can do the following:

- Copy (Copy) the trace results to the clipboard.
- Expand or collapse (Expand or collapse) the displayed results.
- Maximize (Maximize) the trace result window.

The time elapsed information, useful to gauge the processing efforts, is displayed for each phase. The results section also displays the total time taken for packets flowing from an ingress to an egress interface.

The **Trace History** pane displays the stored trace details for each PCAP trace. It can store up to 100 packet traces. You can select a saved trace and run the packet trace activity again. You can do the following:

- Search for a trace using any of the trace parameters.

- Disable saving of the trace to history using the Slider button.
- Delete specific trace results.
- Clear all the traces.

Use the Firewall Threat Defense diagnostic CLI from the web interface

Use the diagnostic CLI to execute selected commands for quick troubleshooting of your device from the web interface without requiring direct SSH access to the device.

You can execute the selected Firewall Threat Defense diagnostic CLI commands from the Firewall Management Center. The commands **ping** (except **ping system**), **traceroute**, and select **show** commands run in the diagnostic CLI rather than the regular CLI.

When you run the **show** commands, if the message `Unable to execute the command properly. Please see logs for more details` is displayed, it means that the command is not valid in the diagnostic CLI. For example, **show access-list** works, but this message will be displayed if you enter **show access-control-policy**. To use non-diagnostic commands, use SSH to log in to a device outside management center.

For more information on the Firewall Threat Defense CLI, see the *Cisco Secure Firewall Threat Defense Command Reference*.

Before you begin

- You must be an Admin, Maintenance, or Security Analyst to use the diagnostic CLI.
- The purpose of diagnostic CLI is to enable the quick use of a few commands that are useful in troubleshooting a device. For access to the full range of commands, open an SSH session directly with the device.
- In deployments using Firewall Management Center high availability,diagnostic CLI is available only in the active Firewall Management Center.

Procedure

Step 1 Choose **Devices** > **Troubleshoot** > **Threat Defense CLI**.

You can also access the CLI tool through the health monitor for the device (**System** (⚙) > **Health** > **Monitor**). From there, you can select the device, click the **View System and Troubleshoot Details** link, click **Advanced Troubleshooting**, then click **Threat Defense CLI** on that page.

Step 2 From the **Device** drop-down list, choose the device on which to execute the diagnostic command.

Step 3 From the **Command** drop-down list, choose the command that you want to execute.

Step 4 Enter the command parameters in the **Parameters** field.

See the *Cisco Secure Firewall Threat Defense Command Reference* for the valid parameters.

For example, to execute **show access-list** command, choose **show** from the **Command** drop-down list, then enter **access-list** in the **Parameters** field.

Note

Do not type the full command in the **Parameters** field. Type only the relevant keywords.

Step 5 Click **Execute** to view the command output.

If the message `Unable to execute the command properly. Please see logs for more details.` is displayed, examine the parameters closely. There might be syntax errors.

This message can also mean that the command you are trying to execute is not a valid command within the context of the diagnostic CLI (which you have accessed from the device using the **system support diagnostic-CLI** command). Log in to the device using SSH to use these commands.

Feature-specific troubleshooting

This reference provides a comprehensive list of troubleshooting resources organized by feature. Each entry includes links to detailed troubleshooting information, best practices, and configuration guidance for specific system features and components.

Table 43: Feature-specific troubleshooting topics

Feature	Relevant Troubleshooting Information
Application control	*Best Practices for Application Control* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
LDAP external authentication	Troubleshoot LDAP authentication connections, on page 232
Licensing	Troubleshoot Smart Licensing, on page 298 Troubleshoot Specific License Reservation, on page 309
Firewall Management Center high availability	Troubleshoot Firewall Management Center high availability, on page 319
User rule conditions	*Troubleshoot User Control* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
User identity sources	For troubleshooting information on ISE/ISE-PIC, TS Agent Identity Source, Captive Portal Identity Source, and Remote Access VPN Identity Source, see the corresponding sections in the *Cisco Secure Firewall Management Center Device Configuration Guide* Troubleshoot LDAP authentication connections, on page 232
URL filtering	*Troubleshoot URL Filtering* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Realms and user data downloads	*Troubleshoot Realms and User Downloads* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Network discovery	*Troubleshooting Your Network Discovery Strategy* in the *Cisco Secure Firewall Management Center Device Configuration Guide*

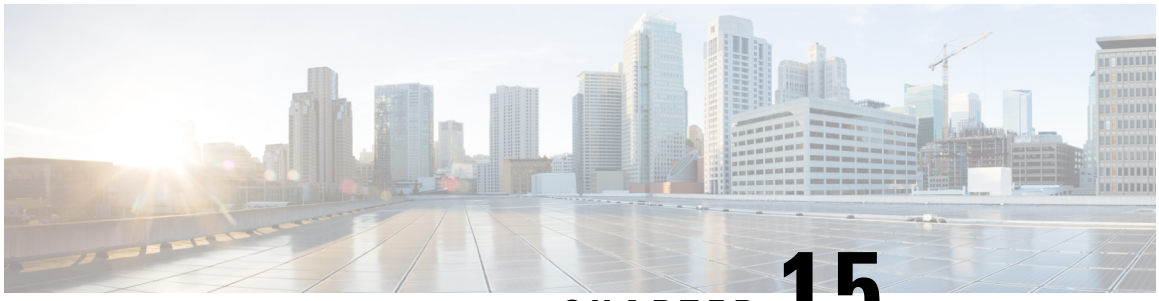
Feature	Relevant Troubleshooting Information
Custom Security Group Tag (SGT) rule conditions	*Custom SGT Rule Conditions* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
SSL rules	Chapter on SSL rules in the *Cisco Secure Firewall Device Manager Configuration Guide*
Cisco Threat Intelligence Director (TID)	*Troubleshoot Secure Firewall Threat Intelligence Director* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Secure Firewall Threat Defense syslog	*About Configuring Syslog* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Intrusion performance statistics	*Intrusion Performance Statistic Logging Configuration* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Connection-based Troubleshooting	Connection-based troubleshooting, on page 458



PART IV

Tools

- Backup/Restore, on page 471
- Scheduling, on page 505
- Import/Export, on page 513
- Data Purge and Storage, on page 521



CHAPTER 15

Backup/Restore

- Backup and restore, on page 471
- Requirements: backup and restore configuration, on page 473
- Guidelines and limitations for backup and restore, on page 474
- Best practices for backup and restore, on page 476
- Backing up Firewall Management Centers or managed devices, on page 478
- Restoring Firewall Management Centers and managed devices, on page 484
- Manage backups and remote storage, on page 499
- History for backup and restore, on page 502

Backup and restore

A backup and restore system is a disaster recovery mechanism that

- enables periodic backups to a secure remote location as part of system maintenance planning,
- provides the ability to recover from disasters, and
- supports both on-demand and scheduled backup operations.

Backup types and contents

Device backups are always configuration-only. Firewall Management Center backups vary by type.

Table 44: Firewall Management Center backups

Backup Type	Backed Up	Not Backed Up
Configurations	Most configurations are backed up. Configuration backups also include locally stored reports in Version 7.2.0–7.2.10. In a multidomain deployment, you must back up configurations. You cannot back up events or TID data only.	These configurations are not backed up and must be reconfigured after restore: • Remote storage settings. • Audit log server certificate settings. • Public and private AMP cloud connections, in Version 7.2.10-7.2.x. Locally stored reports are not backed up in Version 7.2.11-7.2.x. You should store reports in a secure remote location.
Events	All events in the Firewall Management Center database.	Intrusion event review status is not backed up. Restored intrusion events do not appear on Reviewed Events pages.
Threat Intelligence Director (TID) data.	For more information, see *About Backing Up and Restoring Threat Intelligence Director Data* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.	

Restoring configurations overwrites *all* backed-up configurations, with very few exceptions. On the Firewall Management Center, restoring events and TID data overwrites *all* existing events and TID data, with the exception of intrusion events.

Make sure you understand and plan for these restore considerations:

- You cannot restore what is not backed up, as described above.
- Restoring fails VPN certificates.

  The Firewall Threat Defense restore process removes VPN certificates and all VPN configurations from Firewall Threat Defense devices, including certificates added after the backup was taken. After you restore the Firewall Threat Defense device, you must re-add/re-enroll all VPN certificates, and redeploy the device.
- Restoring to a configured Firewall Management Center — instead of factory-fresh or reimaged — merges intrusion events and file lists.

  The Firewall Management Center event restore process does not overwrite intrusion events. Instead, the intrusion events in the backup are added to the database. To avoid duplicates, delete existing intrusion events before you restore.

  The Firewall Management Center configuration restore process does not overwrite clean and custom detection file lists used by Malware Defense. Instead, it merges existing file lists with the file lists in the backup. To replace file lists, delete existing file lists before you restore.

You can perform on-demand backups for the Firewall Management Center and many Firewall Threat Defense devices from the Firewall Management Center.

For more information, see Backing up Firewall Management Centers or managed devices, on page 478.

You can use the scheduler on Firewall Management Center to automate backups. You can also schedule remote device backups from the Firewall Management Center.

The Firewall Management Center setup process schedules weekly configuration-only backups, to be stored locally. This is not a substitute for full off-site backups—after initial setup finishes, you should review your scheduled tasks and adjust them to fit your organization's needs.

For more information, see Scheduling, on page 505.

You can store backups locally. However, we recommend you back up Firewall Management Centers and managed devices to a secure remote location by mounting an NFS, SMB, or SSHFS network volume as remote storage. After you do this, all subsequent backups are copied to that volume, but you can still use the Firewall Management Center to manage them.

For more information, see Remote storage devices, on page 114 and Manage backups and remote storage, on page 499.

You restore the Firewall Management Center from the Backup Management page. You must use the Firewall Threat Defense CLI to restore Firewall Threat Defense devices, except for the ISA 3000 zero-touch restore, which uses an SD card and the reset button.

For more information, see Restoring Firewall Management Centers and managed devices, on page 484.

Requirements: backup and restore configuration

Platform requirements

Ensure that you use supported platforms for backup and restore operations.

- Device backup is supported for both application and container instances.
- This table lists backup support by platform:

Table 45: Backup support by platform

Platform	Backup Supported?		
	Standalone	**High Availability**	**Clusters**
Management center, hardware and virtual	**YES**	**YES**	—
Threat defense hardware	**YES**	**YES**	—
Threat defense virtual, on-prem/private cloud	VMware HyperFlex Nutanix OpenStack	VMware	—
Threat defense virtual, public cloud	—	—	—

Restore platform requirements

Ensure that a replacement managed device meets these hardware requirements.

- A replacement managed device must be the same model as the one you are replacing, with the same number of network modules and same type and number of physical interfaces.
- For Firewall Management Centers, you can use backup and restore not only in an RMA scenario, but also to migrate configurations and events between Firewall Management Centers. For details, including supported target and destination models, see the *Cisco Secure Firewall Management Center Model Migration Guide*.

Version requirements

Ensure that both old and new appliances run compatible software versions before restoring backups.

- As the first step in any backup, note the patch level. To restore a backup, the old and the new appliance must be running the same software version, including patches. To restore to a Firepower 4100/9300 chassis, you must be running a compatible FXOS.
- For Firewall Management Center backups, you are *not* required to have the same VDB or SRU. Note, however, that restoring a backup replaces the existing VDB with the VDB in the backup file. If the restored SRU or the VDB version is older than the one available on the Cisco Support & Download site, we recommend you install the newer version.

License requirements

Address licensing or orphan entitlements concerns as described in the best practices and procedures. If you notice licensing conflicts, contact Cisco TAC.

Domain requirements

Ensure that you have appropriate domain permissions for backup and restore operations.

- To back up or restore the Firewall Management Center: Global only.
- To back up a device from the Firewall Management Center: Global only.
- To restore a device: None. Restore devices locally at the CLI.
- In a multidomain deployment you cannot back up only events/TID data. You must also back up configurations.

Guidelines and limitations for backup and restore

Backup and restore is for disaster recovery/RMA

Use backup and restore primarily for RMA scenarios. Before you begin the restore process of a faulty or failed physical appliance, contact Cisco TAC for replacement hardware.

You can also use backup and restore to migrate configurations and events between Firewall Management Centers. This makes it easier to replace Firewall Management Centers due to technical or business reasons such as a growing organization, migration from a physical to a virtual implementation, hardware refresh, and so on.

Restore on a reimaged management center

Always restore the Firewall Management Center on a freshly reimaged Firewall Management Center.

If you are restoring without reimaging and if you had registered Firewall Threat Defense on the Firewall Management Center after backing up, then when you register the device again on the restored Firewall Management Center, you will encounter device registration failure due to certification error. This error occurs due to the mismatch in the restored certificate database and the non-reimaged Firewall Management Center backup that is restored to.

Backup and restore is not configuration import/export

Do not use the backup and restore process to copy configurations between appliances or devices, or as a way to save configurations while testing new ones. Instead, use the import/export feature.

A backup file contains information that uniquely identifies an appliance, and cannot be shared. For example, Firewall Threat Defense device backups include the device's management IP address and all information the device needs to connect to its managing Firewall Management Center. Do not restore the Firewall Threat Defense backup to a device being managed by a different Firewall Management Center; the restored device will attempt to connect to the Firewall Management Center specified in the backup.

Restore is individual and local

Restore to Firewall Management Centers and managed devices individually and locally.

- You cannot batch-restore to high availability Firewall Management Centers or devices.
- You cannot use the Firewall Management Center to restore a device. For the Firewall Management Center, you can use the web interface to restore. For Firewall Threat Defense devices, you must use the Firewall Threat Defense CLI, except for the ISA 3000 zero-touch restore, which uses an SD card and the reset button.
- While restoring the Firewall Management Center from backup the health policy is also restored. However, any updates to the health monitoring settings are not deployed to devices. You must redeploy all health policies after a successful restore to avoid any discrepancy in health monitoring.
- You cannot use Firewall Management Center user accounts to log into and restore one of its managed devices. The Firewall Management Center and devices maintain their own user accounts.

Strong encryption in the restore device

Ensure that the same strong encryption configuration that was available during the backup is present in the Threat Defense unit that is being restored.

Otherwise, the HA peers may enter a split-brain state.

Recommendation: Configuration Import/Export Guidelines for Firepower 4100/9300

You can use the configuration export feature to export an XML file containing logical device and platform configuration settings for your Firepower 4100/9300 chassis to a remote server or your local computer. You can later import that configuration file to quickly apply the configuration settings to your Firepower 4100/9300 chassis to return to a known good configuration or to recover from a system failure.

- Do not modify the contents of the configuration file. If a configuration file is modified, configuration import using that file might fail.

- Application-specific configuration settings are not contained in the configuration file. You must use the configuration backup tools provided by the application to manage application-specific settings and configurations.
- When you import a configuration to the chassis, all existing configuration on the chassis (including any logical devices) are deleted and completely replaced by the configuration contained in the import file.
- Except in an RMA scenario, we recommend you only import a configuration file to the same chassis where the configuration was exported.
- The platform software version of the Firepower 4100/9300 chassis where you are importing should be the same version as when the export was taken. If not, the import operation is not guaranteed to be successful. We recommend you export a backup configuration whenever the chassis is upgraded or downgraded.
- The Firepower 4100/9300 chassis where you are importing must have the same Network Modules installed in the same slots as when the export was taken.
- The Firepower 4100/9300 chassis where you are importing must have the correct software application images installed for any logical devices defined in the export file that you are importing.
- To avoid overwriting existing backup files, change the file name in the backup operation or copy the existing file to another location.



Note You must backup the logicl APP separately as the FXOS import/export will backup only the FXOS configuration. The FXOS configuration import will cause logical device reboot and it rebuilds the device with the factory default configuration.

Best practices for backup and restore

When to back up

We recommend backing up during a maintenance window or other time of low use.

- Regular scheduled or on-demand backups. As part of your disaster recovery plan, we recommend that you perform periodic backups.
- After SLR changes. Back up the Firewall Management Center after you make changes to Specific Licensing Reservations (SLRs).
- Before upgrade or reimage. If an upgrade fails catastrophically, you may have to reimage and restore.
- During backup and restore process, some of major Management Center and Threat Defense processes which impact the connectivity to each other are not available. Therefore, ensure that parallel actions are not initiated when major processes are paused, stopped, or restarted during a backup and restore.
- After upgrade. Back up after you upgrade, so you have a snapshot of your freshly upgraded deployment.

Maintaining backup file security

We recommend you back up Firewall Management Centers and devices to a secure remote location and verify transfer success.

- Backups are stored as unencrypted archive (.tar) files.
- Private keys in PKI objects are decrypted before they are backed up and reencrypted with a randomly generated key when you restore the backup.

Backup and restoration in Firewall Management Center high-availability deployments

In Firewall Management Center high availability deployments, backing up one Firewall Management Center does not back up the other. You should regularly back up both peers.

Backup and restoration in Firewall Threat Defense high-availability deployments

In the Firewall Threat Defense high availability deployment, you should back up the device pair from the Firewall Management Center, but restore individually and locally from the Firewall Threat Defense CLI.

- Do *not* suspend or break high availability before you restore.
- Do *not* run the **restore** CLI command on both peers at the same time.
- When both peers fail, before the devices are decommissioned, ensure that unregister them both from the Firewall Management Center.

Backup and restoration for firepower 4100/9300 chassis

To restore Firewall Threat Defense software on a Firepower 4100/9300 chassis, the chassis must be running a compatible FXOS version. When you back up a Firepower 4100/9300 chassis, we strongly recommend you also back up FXOS configurations.

Prebackup tasks

Before you back up, you should update the VDB and SRU on the Firewall Management Center and check disk space.

- We always recommend you use the latest vulnerability database (VDB) and intrusion rules (SRU).
- Before you begin a backup, make sure you have enough disk space on the appliance or on your remote storage server.

Prerestoration tasks

Before restore, you should revert any licensing changes made since you took the backup, disconnect faulty appliances, avoid unregistering managed devices, and reimage when necessary.

- Revert any licensing changes made since you took the backup to avoid license conflicts.
- Disconnect the management interface, and for devices, the data interfaces to avoid IP conflicts.
- Do *not* unregister managed devices from the Firewall Management Center, even if you physically disconnect an appliance from the network.
- In an RMA scenario, we recommend you reimage if the replacement appliance is already configured.

Postrestoration tasks

After restoration, perform these tasks: reconfigure everything that was not restored, update the vulnerability database and intrusion rules, and deploy restored Management Center and its Threat Defense devices.

- Reconfigure licensing, remote storage, and audit log server certificate settings.
- Update the vulnerability database (VDB) and intrusion rules (SRU) on the Firewall Management Center.
- Whether you are restoring the Firewall Management Center or device, you must deploy. For a restored device, you may need to force deploy: see refer the "Redeploy Existing Configurations to a Device" section in the *Cisco Secure Firewall Management Center Device Configuration Guide*

Backing up Firewall Management Centers or managed devices

A backup is a data protection mechanism that

- allows you to perform on-demand or scheduled backups for supported appliances,
- can be performed without requiring a backup profile for devices from the Firewall Management Center, and
- requires backup profiles for Firewall Management Center backups.

Backup profile creation

The on-demand backup process allows you to create a new backup profile.

Back up the Firewall Management Center

Use this procedure to create an immediate backup of your Firewall Management Center to ensure system configurations, events, and threat intelligence data are preserved.

Use this procedure to perform an on-demand Firewall Management Center backup.

Before you begin

You must read and understand the requirements, guidelines, limitations, and best practices. You do not want to skip any steps or ignore security concerns. Careful planning and preparation can help you avoid missteps.

- Requirements: backup and restore configuration, on page 473
- Guidelines and limitations for backup and restore, on page 474
- Best practices for backup and restore, on page 476

Procedure

Step 1 Select **System** (⚙) > **Tools** > **Backup/Restore**.

The Backup Management page lists all locally and remotely stored backups. It also lists how much disk space you have available to store backups. Backups can fail if there is not enough space.

Step 2 Choose whether to use an existing backup profile or start fresh.

Firewall Management Center backups require that you use or create a backup profile.

- Click **Backup Profiles** to use an existing backup profile.

  Next to the profile you want to use, click the edit icon. You can then click **Start Backup** to begin the backup right now. Or, if you want to edit the profile, go on to the next step.

- Click **Firepower Management Backup** to start fresh and create a new backup profile.

  Enter a **Name** for the backup profile.

Step 3 Choose what to back up:

- **Back Up Configuration**. In high availability of Firewall Management Centers, if you choose to back up configuration only on an active Firewall Management Center, by default, both the active and standby Firewall Management Centers are backed up into a single unified backup file. For information of unified backup of Firewall Management Center in high availability, see Unified backups of management centers in high availability, on page 343.
- **Back Up Events**
- **Back Up Threat Intelligence Director**

In a multidomain deployment, you must back up configurations. You cannot back up events or TID data only. For details on what is and what is not backed up for each of these choices, see Backup and restore, on page 471.

Step 4 Note the **Storage Location** for Firewall Management Center backup files.

This will either be local storage in `/var/sf/backup/`, or a remote network volume. For more information, see Manage backups and remote storage, on page 499.

Step 5 (Optional) Enable **Copy when complete** to copy completed Firewall Management Center backups to a remote server.

Provide a hostname or IP address, the path to the remote directory, and a username and password. To use an SSH public key instead of a password, copy the contents of the **SSH Public Key** field to the specified user's `authorized_keys` file on the remote server. This option can be used to copy the same backup file to a secondary remote storage.

Note
This option is useful if you want to store backups locally and also SCP them to a remote location. If you configured SSH remote storage, do *not* copy backup files to the same directory using **Copy when complete**.

Step 6 (Optional) Enable **Email** and enter an email address to be notified when the backup completes.

To receive email notifications, you must configure the Firewall Management Center to connect to a mail server: Configure a mail relay host and notification address, on page 71.

Step 7 Click **Start Backup** to start the on-demand backup.

If you are not using an existing backup profile, the system automatically creates one and uses it. If you decide not to run the backup now, you can click **Save** or **Save As New** to save the profile. In either case, you can use the newly created profile to configure scheduled backups.

Note

If you configured remote storage and due to connectivity issues, backup to remote storage may fail. In such cases, if the local management center has atleast 30% of free space, the generated backup file is stored in the local, replacing the oldest local management center backup.

Step 8 Monitor progress in the Message Center.

While the system collects backup data, there may be a temporary pause in data correlation, and you may be prevented from changing configurations related to the backup. If you configured remote storage or enabled **Copy when complete**, the Firewall Management Center may write temporary files to the remote server. These files are cleaned up at the end of the backup process.

The backup process creates a backup file containing your selected system data and stores it in the configured location. You can monitor the backup completion through the Message Center notifications.

What to do next

If you configured remote storage or enabled **Copy when complete**, verify transfer success of the backup file.

Back up a Device from the Firewall Management Center

Use this procedure to perform an on-demand device backup.

Device backups protect your configuration and allow you to restore device settings if needed. You can save backups to local storage or to a remote network location.

Before you begin

You must read and understand the requirements, guidelines, limitations, and best practices. You do not want to skip any steps or ignore security concerns. Careful planning and preparation can help you avoid missteps.

- Requirements: backup and restore configuration, on page 473
- Guidelines and limitations for backup and restore, on page 474
- Best practices for backup and restore, on page 476

If you are backing up a Firepower 4100/9300 chassis, it is especially important that you also back up FXOS configurations: Export an FXOS configuration file, on page 481.

Follow these steps to back up a device from the management center:

Procedure

Step 1 Select **System** (⚙) > **Tools** > **Backup/Restore**, then click **Managed Device Backup**.

Step 2 Select one or more **Managed Devices**.

Step 3 Determine where you want to store the backup file. It can be in one of these locations:

- Local storage in Firewall Management Center at `/var/sf/remote-backup/`.
- Local storage in Firewall Threat Defense at `/var/sf/backup`. For Cluster devices, this option is not available.

- Remote location in a remote network volume. For the ISA 3000, if you have an SD card installed, a copy of the backup will also be made on the SD card at `/mnt/disk3/backup`. For more information, see Manage backups and remote storage, on page 499.

Step 4 If remote storage is configured and enabled. You can choose whether you want to save the backup to remote server or to the device using the **Retrieve to Management Center** check box:

- Enabled (default): Saves the backup to the configured location in the remote server.
- Disabled: Saves the backup to the device in `/var/sf/backup`.

Step 5 If you did not configure or did not enable remote storage, by default, the backup is saved to the local storage in the Firewall Management Center. Choose whether you want to save the backup to Firewall Management Center or to the device :

- Enabled (default): Saves the backup to the Firewall Management Center in `/var/sf/remote-backup/`.
- Disabled: Saves the backup to the device in `/var/sf/backup`.

Step 6 Click **Start Backup** to start the on-demand backup.

Step 7 Monitor progress in the Message Center.

The device backup completes and the backup file is saved to your selected location. You can monitor the backup progress in the Message Center.

What to do next

If you configured remote storage, verify if the transfer of the backup file was successful.

Export an FXOS configuration file

Export configuration files to create backups of logical device and platform settings for disaster recovery or system migration purposes.

Use the configuration export feature to export an XML file containing logical device and platform configuration settings for your Firepower 4100/9300 chassis to a remote server or your local computer.

Note This procedure explains how to use Secure Firewall Chassis Manager to export FXOS configurations when you back up Firewall Threat Defense. For the CLI procedure, see the appropriate version of the Cisco Firepower 4100/9300 FXOS CLI Configuration Guide.

Before you begin

Review the Guidelines and Restrictions.

Procedure

Step 1 Choose **System** > **Configuration** > **Export** on the Secure Firewall Chassis Manager.

Step 2 To export a configuration file to your local computer:

a. Click **Local**.

b. Click **Export**.

The configuration file is created and, depending on your browser, the file might be automatically downloaded to your default download location or you might be prompted to save the file.

Step 3 To export the configuration file to a remote server:

a. Click **Remote**.

b. Choose the protocol to use when communicating with the remote server. This can be one of the following: FTP, TFTP, SCP, or SFTP.

c. Enter the hostname or IP address of the location where the backup file should be stored. This can be a server, storage array, local drive, or any read/write media that the Firepower 4100/9300 chassis can access through the network.

If you use a hostname rather than an IP address, you must configure a DNS server.

d. If you are using a non-default port, enter the port number in the **Port** field.

e. Enter the username and password the system should use to log in to the remote server. This field does not apply if the protocol is TFTP.

Note
The password must not exceed 64 characters. If you enter a password more than 64 character, Firewall Chassis Manager will display an error stating that **property pwd of org-root/cfg-exp-policy-default is out of range**.

f. In the **Location** field, enter the full path to where you want the configuration file exported including the filename.

The configuration file is created and exported to the specified location.

The configuration file is exported to the specified location with the logical device and platform configuration settings for backup or migration purposes.

Create a backup profile

Create a backup profile to save backup preferences including what to back up, where to store the backup file, and notification settings. This allows for consistent backup configurations and enables scheduled backups.

A backup profile is a saved set of preferences—what to back up, where to store the backup file, and so on.

Firewall Management Center backups require backup profiles. Backup profiles are not required to back up a device from the Firewall Management Center.

When you perform an on-demand Firewall Management Center backup, if you do not pick an existing backup profile, the system automatically creates one and uses it. You can then use the newly created profile to configure scheduled backups.

Before you begin

Follow these steps to create a backup profile without performing an on-demand backup:

Procedure

Step 1 Select **System** (⚙) > **Tools** > **Backup/Restore**, then click **Backup Profiles**.

Step 2 Click **Create Profile** and enter a **Name**.

Step 3 Choose what to back up.

- **Back Up Configuration**
- **Back Up Events**
- **Back Up Threat Intelligence Director**

In a multidomain deployment, you must back up configurations. You cannot back up events or TID data only. For details on what is and what is not backed up for each of these choices, see Backup and restore, on page 471.

Step 4 Note the **Storage Location** for backup files.

This will either be local storage in `/var/sf/backup/`, or a remote network volume. For the ISA 3000, if you have an SD card installed, a copy of the backup will also be made on the SD card at `/mnt/disk3/backup`.For more information, see Manage backups and remote storage, on page 499.

Step 5 (Optional) Enable **Copy when complete** to copy completed Firewall Management Center backups to a remote server.

Provide a hostname or IP address, the path to the remote directory, and a username and password. To use an SSH public key instead of a password, copy the contents of the **SSH Public Key** field to the specified user's **authorized_keys** file on the remote server.

Note
This option is useful if you want to store backups locally and also SCP them to a remote location. If you configured SSHFS remote storage, do *not* copy backup files to the same directory using **Copy when complete**.

Step 6 (Optional) Enable **Email** and enter an email address to be notified when the backup completes.

To receive email notifications, you must configure the Firewall Management Center to connect to a mail server: Configure a mail relay host and notification address, on page 71.

Step 7 Click **Save**.

The backup profile is created and available for use in on-demand backups and scheduled backup configurations.

Restoring Firewall Management Centers and managed devices

Restoring Firewall Management Centers and managed devices is a recovery process that

- uses different interfaces depending on the device type
- requires the web interface for Firewall Management Center, and
- requires the CLI for Firewall Threat Defense devices.

Restoration methods

The restoration process varies by device type. You cannot use the Firewall Management Center to restore a device.

Restore Firewall Management Center from Backup

This task restores Firewall Management Center from backup files to recover system configurations, events, and TID data when needed.

When you restore Firewall Management Center backups, you can choose to restore any or all of the components included in the backup file (events, configurations, TID data).

Note Restoring configuration data overwrites *all* configurations, with very few exceptions. Replacing configuration-only data does not restore event data, security logs, or operational history. It also reboots the Firewall Management Center. Restoring events and TID data overwrites *all* existing events and TID data, with the exception of intrusion events. Make sure you are ready.

Use this procedure to restore the Firewall Management Center from backup. For more information on backup and restore in Firewall Management Center HA deployments, see Replace Firewall Management Centers in a high availability pair, on page 337.

Before you begin

You must read and understand the requirements, guidelines, limitations, and best practices. You do not want to skip any steps or ignore security concerns. Careful planning and preparation can help you avoid missteps.

- Requirements: backup and restore configuration, on page 473
- Guidelines and limitations for backup and restore, on page 474
- Best practices for backup and restore, on page 476

Procedure

Step 1 Log into the Firewall Management Center you want to restore.

Step 2 Select **System** (⚙) > **Tools** > **Backup/Restore**.

The Backup Management page lists all locally and remotely stored backup files. You can click a backup file to view its contents.

If the backup file is not in the list and you have it saved on your local computer, click **Upload Backup**; see Manage backups and remote storage, on page 499.

Step 3 Select the backup file you want to restore and click **Restore**.

Step 4 Select from the available components to restore, then click **Restore** again to begin.

Step 5 Monitor progress in the Message Center.

If you are restoring configurations, you can log back in after the Firewall Management Center reboots.

The Firewall Management Center is restored from the selected backup file with the chosen components.

What to do next

- If necessary, reconfigure any licensing settings that you reverted before the restore. If you notice licensing conflicts or orphan entitlements, contact Cisco TAC.
- If necessary, reconfigure remote storage and audit log server certificate settings. These settings are not included in backups.
- Update the SRU and VDB. If the restored SRU or the VDB version is older than the one available on the Cisco Support & Download site, ensure to update the VDB to the latest version before deploying any changes to the device.
- Deploy configurations:
  - If you have performed restoration with the latest Firewall Management Center backup and no configuration changes are deployed to the managed Threat Defense devices after the backup, you are not required to perform deployment after restoration.
  - For all other conditions, you must perform force deployment.



Note For information about how to force deploy, see refer the "Redeploy Existing Configurations to a Device" section in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Restore Firewall Threat Defense from Backup

Device backup and restore is intended for RMA. Restoring configurations overwrites *all* configurations on the device, including the management IP address. It also reboots the device.

In case of hardware failure, this procedure outlines how to replace the, standalone or high availability. It assumes you have access to a successful backup of the device or devices you are replacing; see Back up a Device from the Firewall Management Center, on page 480.

This procedure applies to all platforms except for the Firepower 4100/9300 and Firewall Threat Defense Virtual. For zero-touch restore on the ISA 3000 using an SD card, see Zero-Touch Restore Firewall Threat Defense from Backup: ISA 3000, on page 489.

For high availability devices, you can use this procedure to replace all peers. To replace all, perform all steps on all devices simultaneously, except the **restore** CLI command itself.

Note Do *not* unregister from the Firewall Management Center, even when disconnecting a device from the network. For Firewall Threat Defense high availability devices, do *not* suspend or break high availability. Maintaining these links ensures replacement devices can automatically reconnect after restore.

Before you begin

You must read and understand the requirements, guidelines, limitations, and best practices. You do not want to skip any steps or ignore security concerns. Careful planning and preparation can help you avoid missteps.

- Requirements: backup and restore configuration, on page 473
- Guidelines and limitations for backup and restore, on page 474
- Best practices for backup and restore, on page 476

Procedure

Step 1 Contact Cisco TAC for replacement hardware.

Obtain an identical model, with the same number of network modules and same type and number of physical interfaces. You can begin the RMA process from the Cisco Returns Portal.

Step 2 Locate a successful backup of the faulty device.

Depending on your backup configuration, device backups may be stored:

- On the faulty device itself in `/var/sf/backup`.
- On the Firewall Management Center in `/var/sf/remote-backup`.
- In a remote storage location.

For Firewall Threat Defense high availability devices, you back up the group as a unit. For high availability devices, the backup process produces unique backup files, with each device's role indicated in the backup file name.

If the only copy of the backup is on the faulty device, copy it somewhere else now. If you reimage the device, the backup will be erased. If something else goes wrong, you may not be able to recover the backup. For more information, see Manage backups and remote storage, on page 499.

The replacement device will need the backup, but can retrieve it with SCP during the restore process. We recommend you put the backup somewhere SCP-accessible to the replacement device. Or, you can copy the backup to the replacement device itself.

Step 3 Remove (unrack) the faulty device.

Disconnect all interfaces. In Firewall Threat Defense high availability deployments, this includes the failover link.

See the hardware installation and getting started guides for your model: http://www.cisco.com/go/ftd-quick.

Note
Do *not* unregister from the Firewall Management Center, even when disconnecting a device from the network. For Firewall Threat Defense high availability devices, do *not* suspend or break high availability. Maintaining these links ensures replacement devices can automatically reconnect after restore.

Step 4 Install the replacement device and connect it to the management network.

Connect the device to power and the management interface to the management network. In Firewall Threat Defense high availability deployments, connect the failover link. However, do *not* connect the data interfaces.

See the hardware installation guide for your model: http://www.cisco.com/go/ftd-quick.

Step 5 (Optional) Reimage the replacement device.

In an RMA scenario, the replacement device will arrive configured with factory defaults. If the replacement device is not running the same major version as the faulty device, we recommend you reimage.

See the *Cisco Secure Firewall ASA and Secure Firewall Threat Defense Reimage Guide*.

Step 6 Perform initial configuration on the replacement device.

Access the Firewall Threat Defense CLI as the `admin` user. A setup wizard prompts you to configure the management IP address, gateway, and other basic network settings.

Do not set the same management IP address as the faulty device. This can cause problems if you need to register the device in order to patch it. The restore process will correctly reset the management IP address.

See the initial configuration topics in the getting started guide for your model: http://www.cisco.com/go/ftd-quick.

Note
If you need to patch the replacement device, start the Firewall Management Center registration process as described in the getting started guide. If you do not need to patch, do *not* register.

Step 7 Make sure the replacement device is running the same software version, *including patches*, as the faulty device.

Ensure that the existing device should not be deleted from the Firewall Management Center. The replacement device should be unmanaged from the physical network and the new hardware as well as the replacing Firewall Threat Defense patch should have the same version. The Firewall Threat Defense CLI does not have an upgrade command. To patch:

a) From the Firewall Management Center web interface, complete the device registration process.

   Create a new AC policy and use the default action "Network Discovery". Leave this policy as is; do not add any features or modifications. This is being used to register the device and deploy a policy with no features so that you do not require licenses, and you will then be able to patch the device. Once backup is restored, it should restore the licensing and policy into the expected state.

b) Patch the device: https://www.cisco.com/go/ftd-upgrade.

c) Unregister the freshly patched device from the Firewall Management Center.

   If you do not unregister, you will have a ghost device registered to the Firewall Management Center after the restore process brings your "old" device back up.

Step 8 Make sure the replacement device has access to the backup file.

The restore process can retrieve the backup with SCP, so we recommend you put the backup somewhere accessible. Or, you can manually copy the backup to the replacement device itself, to `/var/sf/backup`.

Step 9 From the Firewall Threat Defense CLI, restore the backup.

Access the Firewall Threat Defense CLI as the `admin` user. You can use the console or you can SSH to the newly configured management interface (IP address or hostname). Keep in mind that the restore process will change this IP address.

To restore:

- With SCP: **restore remote-manager-backup location** *scp-hostname username filepath backup tar-file*
- From the local device: **restore remote-manager-backup** *backup tar-file*

In Firewall Threat Defense high availability deployments, make sure you choose the appropriate backup file: primary vs secondary. The role is noted in the backup file name. If you are restoring all devices, do this sequentially. Do not run the **restore** command on the next device until the restore process completes for the first device, including the reboot.

Step 10 Log into the Firewall Management Center and wait for the replacement device to connect.

When the restore is done, the device logs you out of the CLI, reboots, and automatically connects to the Firewall Management Center. At this time, the device should appear out of date.

Step 11 Before you deploy, perform any post-restore tasks and resolve any post-restore issues:

- Resolve licensing conflicts or orphan entitlements. Contact Cisco TAC.
- Resume high availability synchronization. From the Firewall Threat Defense CLI , enter `configure high-availability resume`. See *Suspend and Resume High Availability* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

  Note
  Skip this step in Version 7.2.10–7.2.x, where resume is automatic.

- Re-add/re-enroll all VPN certificates. The restore process removes VPN certificates from Firewall Threat Defense devices, including certificates added after the backup was taken. See *Managing VPN Certificates* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Step 12 Deploy configurations.

You must deploy. After you restore a device, you must force deploy from the Device Management page. See *Redeploy Existing Configurations to a Device* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Note
When you restore a Firewall Threat Defense backup to one of the Firewall Threat Defense high availability (HA) nodes, any mismatch of encryption configuration between two nodes put the Firewall Threat Defense HA into active-active mode, which prevents the IPsec tunnel from being established.

Both the primary and secondary nodes must use the same encryption method to establish an IPsec tunnel. Disable HA on the restored Firewall Threat Defense. Then, use Firewall Management Center to deploy the Firewall Threat Defense HA pair, ensuring that both nodes use the same encryption method as the restored Firewall Threat Defense and then enable the HA in restored Firewall Threat Defense.

Step 13 Connect the device's data interfaces.

See the hardware installation guide for your model: http://www.cisco.com/go/ftd-quick.

The device has been successfully restored from backup and is operational with the restored configuration.

What to do next

Verify that the restore succeeded and the replacement device is passing traffic as expected.

Zero-Touch Restore Firewall Threat Defense from Backup: ISA 3000

This procedure allows you to replace a failed ISA 3000 device, either standalone or in an HA pair, using a backup stored on an SD card.

Device backup and restore is intended for RMA. Restoring configurations overwrites *all* configurations on the device, including the management IP address. It also reboots the device.

In case of hardware failure, this procedure outlines how to replace the ISA 3000, either standalone or in an HA pair. It assumes you have a backup of the failed unit on an SD card; see Back up a Device from the Firewall Management Center, on page 480.

For high availability devices, you can use this procedure to replace all peers. To replace all, perform all steps on all devices simultaneously, except the **restore** CLI command itself.

Note Do *not* unregister from the Firewall Management Center, even when disconnecting a device from the network. For Firewall Threat Defense high availability devices, do *not* suspend or break high availability. Maintaining these links ensures replacement devices can automatically reconnect after restore.

Before you begin

You must read and understand the requirements, guidelines, limitations, and best practices. You do not want to skip any steps or ignore security concerns. Careful planning and preparation can help you avoid missteps.

- Requirements: backup and restore configuration, on page 473
- Guidelines and limitations for backup and restore, on page 474
- Best practices for backup and restore, on page 476

Procedure

Step 1 Contact Cisco TAC for replacement hardware.

Obtain an identical model, with the same number of network modules and same type and number of physical interfaces. You can begin the RMA process from the Cisco Returns Portal.

Step 2 Remove the SD card from the faulty device, and unrack the device.

Disconnect all interfaces. In Firewall Threat Defense HA deployments, this includes the failover link.

Note

Do *not* unregister from the Firewall Management Center, even when disconnecting a device from the network. For Firewall Threat Defense high availability devices, do *not* suspend or break high availability. Maintaining these links ensures replacement devices can automatically reconnect after restore.

Step 3 Rerack the replacement device, and connect it to the management network. In Firewall Threat Defense HA deployments, connect the failover link. However, do *not* connect the data interfaces.

If you need to reimage the device or apply a software patch, connect the power connector.

Step 4 (May be required) Reimage the replacement device.

In an RMA scenario, the replacement device will arrive configured with factory defaults. If the replacement device is not running the same major version as the faulty device, you need to reimage. Obtain the installer from https://www.cisco.com/go/isa3000-software.

See the *Cisco Secure Firewall ASA and Secure Firewall Threat Defense Reimage Guide* to reimage.

Step 5 (May be required) Make sure the replacement device is running the same Secure Firewall software version, *including the same patch version*, as the faulty device. If you need to patch the device, you can connect to Secure Firewall Device Manager (Firewall Device Manager) to install the patch.

This procedure assumes you have a factory default configuration. If you already configured the device, you can log into Firewall Device Manager and go directly to the **Device** > **Upgrades** page to install the patch.

In either case, obtain the patch package from https://www.cisco.com/go/isa3000-software.

a) Connect your computer directly to the inside (Ethernet 1/2) interface, and access Firewall Device Manager on the default IP address: **https://192.168.95.1**.
b) Enter the **admin** username and the default password **Admin123**, then click **Login**.
c) Complete the setup wizard. Keep in mind that you are not going to retain anything you configure in Firewall Device Manager; you only want to get past any initial configuration so you can apply the patches, so it doesn't matter what you enter in the setup wizard.
d) Go to the **Device** > **Upgrades** page.

   The **System Upgrade** section shows the currently running software version.

e) Upload the patch file by clicking **Browse**.
f) Click **Install** to start the installation process.

   Information next to the icon indicates whether the device will reboot during installation. You are automatically logged out of the system. Installation might take 30 minutes or more.

   Wait before logging into the system again. The Device Summary, or System monitoring dashboard, should show the new version.

   Note
   Do not simply refresh the browser window. Instead, delete any path from the URL, and reconnect to the home page. This ensures that cached information gets refreshed with the latest code.

Step 6 Insert the SD card in the replacement device and then do the following:

a) Power on or reboot the device and shortly after it starts the bootup, depress and hold the Reset button for no fewer than 3 seconds and no longer than 15 seconds.

   If you used Firewall Device Manager to install a patch, you can reboot from the **Device** > **System Settings** > **Reboot/Shutdown** page. From the Firewall Threat Defense CLI, use the **reboot** command. If you have not yet attached power, attach it now.

Use a standard size #1 paper clip with wire gauge 0.033 inch or smaller to depress the Reset button. The restoration process is triggered during bootup. The device restores the configuration, and then reboots. The device will then register with the Firewall Management Center automatically.

If you are restoring both devices in an HA pair, do this sequentially. Do not restore the second device until the restore process completes for the first device, including the reboot.

b) Log into the Firewall Management Center and wait for the replacement device to connect.

At this time, the device should appear out of date.

Step 7 Before you deploy, perform any post-restore tasks and resolve any post-restore issues:

- Resolve licensing conflicts or orphan entitlements. Contact Cisco TAC.
- Resume high availability synchronization. From the Firewall Threat Defense CLI , enter `configure high-availability resume`. See *Suspend and Resume High Availability* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

  Note
  Skip this step in Version 7.2.10–7.2.x, where resume is automatic.

- Re-add/re-enroll all VPN certificates. The restore process removes VPN certificates from Firewall Threat Defense devices, including certificates added after the backup was taken. See *Managing VPN Certificates* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Step 8 Deploy configurations.

You must deploy. After you restore a device, you must force deploy from the Device Management page. See *Redeploy Existing Configurations to a Device* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Note
When you restore a Firewall Threat Defense backup to one of the Firewall Threat Defense high availability (HA) nodes, any mismatch of encryption configuration between two nodes put the Firewall Threat Defense HA into active-active mode, which prevents the IPsec tunnel from being established.

Both the primary and secondary nodes must use the same encryption method to establish an IPsec tunnel. Disable HA on the restored Firewall Threat Defense. Then, use Firewall Management Center to deploy the Firewall Threat Defense HA pair, ensuring that both nodes use the same encryption method as the restored Firewall Threat Defense and then enable the HA in restored Firewall Threat Defense.

Step 9 Connect the device's data interfaces.

See the hardware installation guide for your model: http://www.cisco.com/go/ftd-quick.

The replacement device is restored with the backed-up configuration and reconnects to the management center. The device is ready for deployment with the current configuration.

What to do next

Verify that the restore succeeded and the replacement device is passing traffic as expected.

Restore Firewall Threat Defense from Backup: Firepower 4100/9300 Chassis

Device backup and restore is intended for RMA. Restoring configurations overwrites *all* configurations on the device, including the management IP address. It also reboots the device.

In case of hardware failure, this procedure outlines how to replace a Firepower 4100/9300, standalone or in a High Availability pair. It assumes you have access to successful backups of:

- The logical device or devices you are replacing; see Back up a Device from the Firewall Management Center, on page 480.
- FXOS configurations; see Export an FXOS configuration file, on page 481.

For high availability devices, you can use this procedure to replace all peers. To replace all, perform all steps on all devices simultaneously, except the **restore** CLI command itself.

Note Do *not* unregister from the Firewall Management Center, even when disconnecting a device from the network. For Firewall Threat Defense high availability devices, do *not* suspend or break high availability. Maintaining these links ensures replacement devices can automatically reconnect after restore.

Before you begin

You must read and understand the requirements, guidelines, limitations, and best practices. You do not want to skip any steps or ignore security concerns. Careful planning and preparation can help you avoid missteps.

- Requirements: backup and restore configuration, on page 473
- Guidelines and limitations for backup and restore, on page 474
- Best practices for backup and restore, on page 476

Procedure

Step 1 Contact Cisco TAC for replacement hardware.

Obtain an identical model, with the same number of network modules and same type and number of physical interfaces. You can begin the RMA process from the Cisco Returns Portal.

Step 2 Locate a successful backup of the faulty device.

Depending on your backup configuration, device backups may be stored:

- On the faulty device itself in `/var/sf/backup`.
- On the Firewall Management Center in `/var/sf/remote-backup`.
- In a remote storage location.

For Firewall Threat Defense high availability devices, you back up the group as a unit. For high availability devices, the backup process produces unique backup files, with each device's role indicated in the backup file name.

If the only copy of the backup is on the faulty device, copy it somewhere else now. If you reimage the device, the backup will be erased. If something else goes wrong, you may not be able to recover the backup. For more information, see Manage backups and remote storage, on page 499.

The replacement device will need the backup, but can retrieve it with SCP during the restore process. We recommend you put the backup somewhere SCP-accessible to the replacement device. Or, you can copy the backup to the replacement device itself.

Step 3 Locate a successful backup of your FXOS configurations.

Step 4 Remove (unrack) the faulty device.

Disconnect all interfaces. In Firewall Threat Defense high availability deployments, this includes the failover link.

See the hardware installation and getting started guides for your model: http://www.cisco.com/go/ftd-quick.

Note
Do *not* unregister from the Firewall Management Center, even when disconnecting a device from the network. For Firewall Threat Defense high availability devices, do *not* suspend or break high availability. Maintaining these links ensures replacement devices can automatically reconnect after restore.

Step 5 Install the replacement device and connect it to the management network.

Connect the device to power and the management interface to the management network. In Firewall Threat Defense high availability deployments, connect the failover link. However, do *not* connect the data interfaces.

See the hardware installation guide for your model: http://www.cisco.com/go/ftd-quick.

Step 6 (Optional) Reimage the replacement device.

In an RMA scenario, the replacement device will arrive configured with factory defaults. If the replacement device is not running the same major version as the faulty device, we recommend you reimage.

See the instructions on restoring the factory default configuration in the appropriate Cisco Firepower 4100/9300 FXOS Firepower Chassis Manager Configuration Guide.

Step 7 Make sure FXOS is running a compatible version.

You must be running a compatible FXOS version before you re-add logical devices. You can use Firewall Chassis Manager to import your backed-up FXOS configurations: Importing a Configuration File, on page 495.

Step 8 Use Firewall Chassis Manager to add logical devices and perform initial configurations.

Do not set the same management IP addresses as the logical device or devices on the faulty chassis. This can cause problems if you need to register a logical device in order to patch it. The restore process will correctly reset the management IP address.

See the Firewall Management Center deployment chapter in the getting started guide for your model: http://www.cisco.com/go/ftd-quick.

Note
If you need to patch a logical device, register to the Firewall Management Center as described in the getting started guide. If you do not need to patch, do *not* register.

Step 9 Make sure the replacement device is running the same software version, *including patches*, as the faulty device.

Ensure that the existing device should not be deleted from the Firewall Management Center. The replacement device should be unmanaged from the physical network and the new hardware as well as the replacing Firewall Threat Defense patch should have the same version. The Firewall Threat Defense CLI does not have an upgrade command. To patch:

a) From the Firewall Management Center web interface, complete the device registration process.

   Create a new AC policy and use the default action "Network Discovery". Leave this policy as is; do not add any features or modifications. This is being used to register the device and deploy a policy with no features so that you do not require licenses, and you will then be able to patch the device. Once backup is restored, it should restore the licensing and policy into the expected state.

b) Patch the device: https://www.cisco.com/go/ftd-upgrade.

c) Unregister the freshly patched device from the Firewall Management Center.

   If you do not unregister, you will have a ghost device registered to the Firewall Management Center after the restore process brings your "old" device back up.

Step 10 Make sure the replacement device has access to the backup file.

The restore process can retrieve the backup with SCP, so we recommend you put the backup somewhere accessible. Or, you can manually copy the backup to the replacement device itself, to `/var/sf/backup`.

Step 11 From the Firewall Threat Defense CLI, restore the backup.

Access the Firewall Threat Defense CLI as the `admin` user. You can use the console or you can SSH to the newly configured management interface (IP address or hostname). Keep in mind that the restore process will change this IP address.

To restore:

- With SCP: **restore remote-manager-backup location** *scp-hostname username filepath backup tar-file*
- From the local device: **restore remote-manager-backup** *backup tar-file*

In Firewall Threat Defense high availability deployments, make sure you choose the appropriate backup file: primary vs secondary. The role is noted in the backup file name. If you are restoring all devices, do this sequentially. Do not run the **restore** command on the next device until the restore process completes for the first device, including the reboot.

Step 12 Log into the Firewall Management Center and wait for the replacement device to connect.

When the restore is done, the device logs you out of the CLI, reboots, and automatically connects to the Firewall Management Center. At this time, the device should appear out of date.

Step 13 Before you deploy, perform any post-restore tasks and resolve any post-restore issues:

- Resolve licensing conflicts or orphan entitlements. For additional assistance and support, contact Cisco TAC.
- Re-add/re-enroll all VPN certificates. The restore process removes VPN certificates from Firewall Threat Defense devices, including certificates added after the backup was taken. See *Managing VPN Certificates* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Step 14 Deploy configurations.

You must deploy. After you restore a device, you must force deploy from the Device Management page. See *Redeploy Existing Configurations to a Device* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Note
When you restore a Firewall Threat Defense backup to one of the Firewall Threat Defense high availability (HA) nodes, any mismatch of encryption configuration between two nodes put the Firewall Threat Defense HA into active-active mode, which prevents the IPsec tunnel from being established.

Both the primary and secondary nodes must use the same encryption method to establish an IPsec tunnel. Disable HA on the restored Firewall Threat Defense. Then, use Firewall Management Center to deploy the Firewall Threat Defense HA pair, ensuring that both nodes use the same encryption method as the restored Firewall Threat Defense and then enable the HA in restored Firewall Threat Defense.

Step 15 Connect the device's data interfaces.

See the hardware installation guide for your model: http://www.cisco.com/go/ftd-quick.

Verify that the restore succeeded and the replacement device is passing traffic as expected.

Importing a Configuration File

You can use the configuration import feature to apply configuration settings that were previously exported from your Firepower 4100/9300 chassis . This feature allows you to return to a known good configuration or to recover from a system failure.

Note This procedure explains how to use Firewall Chassis Manager to import FXOS configurations before you restore the software. For the CLI procedure, see the appropriate version of the Cisco Firepower 4100/9300 FXOS CLI Configuration Guide .

Before you begin

Review the Guidelines and Restrictions .

Procedure

Step 1 Choose **System** > **Tools** > **Import/Export** on the Firewall Chassis Manager .

Step 2 To import from a local configuration file:

a) Click **Local** .
b) Click **Choose File** to navigate to and select the configuration file that you want to import.
c) Click **Import** .
A confirmation dialog box opens asking you to confirm that you want to proceed and warning you that the chassis might need to restart.
d) Click **Yes** to confirm that you want to import the specified configuration file.
The existing configuration is deleted and the configuration specified in the import file is applied to the Firepower 4100/9300 chassis . If there is a breakout port configuration change during the import, the Firepower 4100/9300 chassis will need to restart.

Step 3 To import from a configuration file on a remote server:

a) Click **Remote**.

b) Choose the protocol to use when communicating with the remote server. This can be one of the following: FTP, TFTP, SCP, or SFTP.

c) If you are using a non-default port, enter the port number in the **Port** field. Enter the hostname or IP address of the location where the backup file is stored. This can be a server, storage array, local drive, or any read/write media that the Firepower 4100/9300 chassis can access through the network.

If you use a hostname rather than an IP address, you must configure a DNS server.

d) Enter the username the system should use to log in to the remote server. This field does not apply if the protocol is TFTP.

e) Enter the password for the remote server username. This field does not apply if the protocol is TFTP.

Note
The password must not exceed 64 characters. If you enter a password more than 64 character, Firewall Chassis Manager will display an error stating that `property pwd of org-root/cfg-exp-policy-default is out of range .`

f) In the **File Path** field, enter the full path to the configuration file including the file name.

g) Click **Yes** to confirm that you want to import the specified configuration file.
The existing configuration is deleted and the configuration specified in the import file is applied to the Firepower 4100/9300 chassis . If there is a breakout port configuration change during the import, the Firepower 4100/9300 chassis will need to restart.

Restore Firewall Threat Defense Virtual from Backup

This task restores a faulty or failed Firewall Threat Defense Virtual device from a backup, ensuring business continuity and minimal downtime.

Use this procedure to replace a faulty or failed Firewall Threat Defense Virtual device.

For high availability devices, you can use this procedure to replace all peers. To replace all, perform all steps on all devices simultaneously, except the **restore** CLI command itself.



Note Do *not* unregister from the Firewall Management Center, even when disconnecting a device from the network. For Firewall Threat Defense high availability devices, do *not* suspend or break high availability. Maintaining these links ensures replacement devices can automatically reconnect after restore.

Before you begin

You must read and understand the requirements, guidelines, limitations, and best practices. You do not want to skip any steps or ignore security concerns. Careful planning and preparation can help you avoid missteps.

- Requirements: backup and restore configuration, on page 473
- Guidelines and limitations for backup and restore, on page 474
- Best practices for backup and restore, on page 476

Procedure

Step 1 Locate a successful backup of the faulty device.

Depending on your backup configuration, device backups may be stored:

- On the faulty device itself in `/var/sf/backup`.
- On the Firewall Management Center in `/var/sf/remote-backup`.
- In a remote storage location.

For Firewall Threat Defense high availability devices, you back up the group as a unit. For high availability devices, the backup process produces unique backup files, with each device's role indicated in the backup file name.

If the only copy of the backup is on the faulty device, copy it somewhere else now. If you reimage the device, the backup will be erased. If something else goes wrong, you may not be able to recover the backup. For more information, see Manage backups and remote storage, on page 499.

The replacement device will need the backup, but can retrieve it with SCP during the restore process. We recommend you put the backup somewhere SCP-accessible to the replacement device. Or, you can copy the backup to the replacement device itself.

Step 2 Remove the faulty device.

Shut down, power off, and delete the virtual machine. For procedures, see the documentation for your virtual environment.

Step 3 Deploy a replacement device.

See https://www.cisco.com/go/ftdv-quick.

Step 4 Perform initial configuration on the replacement device.

Use the console to access the Firewall Threat Defense CLI as the `admin` user. A setup wizard prompts you to configure the management IP address, gateway, and other basic network settings.

Do not set the same management IP address as the faulty device. This can cause problems if you need to register the device in order to patch it. The restore process will correctly reset the management IP address.

See the CLI setup topics in the getting started guide: https://www.cisco.com/go/ftdv-quick.

Note
If you need to patch the replacement device, start the Firewall Management Center registration process as described in the getting started guide. If you do not need to patch, do *not* register.

Step 5 Make sure the replacement device is running the same software version, *including patches*, as the faulty device.

Ensure that the existing device should not be deleted from the Firewall Management Center. The replacement device should be unmanaged from the physical network and the new hardware as well as the replacing Firewall Threat Defense patch should have the same version. The Firewall Threat Defense CLI does not have an upgrade command. To patch:

a) From the Firewall Management Center web interface, complete the device registration process.

Create a new AC policy and use the default action "Network Discovery". Leave this policy as is; do not add any features or modifications. This is being used to register the device and deploy a policy with no features so that you do not require licenses, and you will then be able to patch the device. Once backup is restored, it should restore the licensing and policy into the expected state.

b) Patch the device: https://www.cisco.com/go/ftd-upgrade.

c) Unregister the freshly patched device from the Firewall Management Center.

If you do not unregister, you will have a ghost device registered to the Firewall Management Center after the restore process brings your "old" device back up.

Step 6 Make sure the replacement device has access to the backup file.

The restore process can retrieve the backup with SCP, so we recommend you put the backup somewhere accessible. Or, you can manually copy the backup to the replacement device itself, to `/var/sf/backup`.

Step 7 From the Firewall Threat Defense CLI, restore the backup.

Access the Firewall Threat Defense CLI as the `admin` user. You can use the console or you can SSH to the newly configured management interface (IP address or hostname). Keep in mind that the restore process will change this IP address.

To restore:

- With SCP: **restore remote-manager-backup location** *scp-hostname username filepath backup tar-file*
- From the local device: **restore remote-manager-backup** *backup tar-file*

In Firewall Threat Defense high availability deployments, make sure you choose the appropriate backup file: primary vs secondary. The role is noted in the backup file name. If you are restoring all devices, do this sequentially. Do not run the **restore** command on the next device until the restore process completes for the first device, including the reboot.

Step 8 Log into the Firewall Management Center and wait for the replacement device to connect.

When the restore is done, the device logs you out of the CLI, reboots, and automatically connects to the Firewall Management Center. At this time, the device should appear out of date.

Step 9 Before you deploy, perform any post-restore tasks and resolve any post-restore issues:

- Resolve licensing conflicts or orphan entitlements. For additional assistance and support, contact Cisco TAC.
- Re-add/re-enroll all VPN certificates. The restore process removes VPN certificates from Firewall Threat Defense devices, including certificates added after the backup was taken. See *Managing VPN Certificates* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Step 10 Deploy configurations.

You must deploy. After you restore a device, you must force deploy from the Device Management page. See *Redeploy Existing Configurations to a Device* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Note
When you restore a Firewall Threat Defense backup to one of the Firewall Threat Defense high availability (HA) nodes, any mismatch of encryption configuration between two nodes put the Firewall Threat Defense HA into active-active mode, which prevents the IPsec tunnel from being established.

Both the primary and secondary nodes must use the same encryption method to establish an IPsec tunnel. Disable HA on the restored Firewall Threat Defense. Then, use Firewall Management Center to deploy the Firewall Threat Defense HA pair, ensuring that both nodes use the same encryption method as the restored Firewall Threat Defense and then enable the HA in restored Firewall Threat Defense.

Step 11 Add and configure data interfaces.

See the getting started guide: https://www.cisco.com/go/ftdv-quick.

The Firewall Threat Defense Virtual device is successfully restored from backup and returns to operational status with all previous configurations and policies intact.

What to do next

Verify that the restore succeeded and the replacement device is passing traffic as expected.

Manage backups and remote storage

This task allows you to manage backup files stored on appliances, including transferring files to and from remote storage locations, enabling or disabling remote storage, and performing file operations such as upload, download, and deletion.

Backups are stored as unencrypted archive (.tar) files. The file name includes identifying information that can include:

- The name of the backup profile or scheduled task associated with the backup.
- The display name or IP address of the backed-up appliance.
- The appliance's role, such as a member of an HA pair.

We recommend you back up appliances to a secure remote location and verify transfer success. Backups left on an appliance may be deleted, either manually or by the upgrade process; upgrades purge locally stored backups. For more information on your options, see Backup storage locations, on page 501.



Caution Especially because backup files are unencrypted, do *not* allow unauthorized access. If backup files are modified, the restore process will fail. Keep in mind that anyone with the Admin/Maint role can access the Backup Management page, where they can move and delete files from remote storage.

Procedure

Step 1 Select **System** (⚙) > **Tools** > **Backup/Restore**.

The Backup Management page lists available backups. It also lists how much disk space you have available to store backups. Backups can fail if there is not enough space.

Step 2 Do one of the following:

Table 46: Remote Storage and Backup File Management

To	Do This
Enable or disable remote storage for backups without having to edit the Firewall Management Center system configuration.	Click **Enable Remote Storage for Backups**. This option appears only after you configure remote storage. Toggling it here also toggles it in the system configuration (**System** (⚙) > **Configuration** > **Remote Storage Device**). **Tip** To quickly access your remote storage configuration, click **Remote Storage** at the upper right of the Backup Management page. **Note** To store backup on the remote storage location, you must also enable the **Retrieve to Management Center** option (see Back up a Device from the Firewall Management Center, on page 480).
Move a file between the Firewall Management Center and the remote storage location.	Click **Move**. You can move a file back and forth as many times as you want. This will delete—not copy—the file from the current location. When you move a backup file from remote storage to the Firewall Management Center, where it is stored on the Firewall Management Center depends on the kind of backup: • Firewall Management Center backups: `/var/sf/backup` • Device backups: `/var/sf/remote-backup`
View the contents of the backup.	Click the backup file.
Delete a backup file.	Choose a backup file and click **Delete**. You can delete both locally and remotely stored backup files.
Upload a backup file from your com puter.	Click **Upload Backup**, choose a backup file, and click **Upload Backup** again.
Download a backup to your computer.	Choose a backup file and click **Download**. Unlike moving a backup file, this does not delete the backup from the Firewall Management Center. Store your downloaded backup in a secure location.

You have successfully managed backup files on your system. The selected operations have been completed, including enabling remote storage, transferring files, or performing file management tasks as specified.

Backup storage locations

This reference provides information about backup storage options for Firewall Management Centers and managed devices.

Table 47: Backup storage locations

Location	Details
Remote, by mounting a network volume (NFS, SMB, SSHFS).	**Note** Backup is stored on a remote storage location only when you have configured remote storage and enabled the **Retrieve to Management Center** option (see Back up a Device from the Firewall Management Center, on page 480). In the Firewall Management Center's system configuration, you can mount an NFS, SMB, or SSHFS network volume as remote storage for Firewall Management Center and device backups; see Remote storage devices, on page 114.) After you do this, all subsequent Firewall Management Center backups *and Firewall Management Center-initiated device backups* are copied to that volume, but you can still use the Firewall Management Center to manage them (restore, download, upload, delete, move). Note that only the Firewall Management Center mounts the network volume. Managed device backup files are routed through the Firewall Management Center. Make sure you have the bandwidth to perform a large data transfer between the Firewall Management Center and its devices. For more information, see Guidelines for Downloading Data from the Firepower Management Center to Managed Devices (Troubleshooting TechNote).
Remote, by copying (SCP).	**Note** Backup is stored on a remote storage location only when you have configured remote storage and enabled the **Retrieve to Management Center** option (see Back up a Device from the Firewall Management Center, on page 480). For the Firewall Management Center, you can use a **Copy when complete** option to securely copy (SCP) completed backups to a remote server. Compared with remote storage by mounting a network volume, **Copy when complete** cannot copy to NFS or SMB volumes. You cannot provide CLI options or set a disk space threshold, and it does not affect remote storage of reports. You also cannot manage backup files after they are copied out. This option is useful if you want to store backups locally *and* SCP them to a remote location. **Note** If you configure SSHFS remote storage in the Firewall Management Center system configuration, do *not* copy backup files to the same directory using **Copy when complete**.

Location	Details
Local, on the Firewall Management Center.	If you do not configure remote storage by mounting a network volume, you can save backup files on the Firewall Management Center: • Firewall Management Center backups are saved to `/var/sf/backup`. • Device backups are saved to `/var/sf/remote-backup` on the Firewall Management Center if you enable the **Retrieve to Management Center** option when you perform the backup.
Local, on the device internal flash memory.	Device backup files are saved to `/var/sf/backup` on the device if you: • Do not configure remote storage by mounting a network volume. • Do not enable **Retrieve to Management Center**.
Local, on the device SD card.	For the ISA 3000, when you back up the device to the local `/var/sf/backup` internal flash memory location, if you have an SD card installed, the backup is automatically copied to the SD card at `/mnt/disk3/backup/` for use with zero-touch restore.

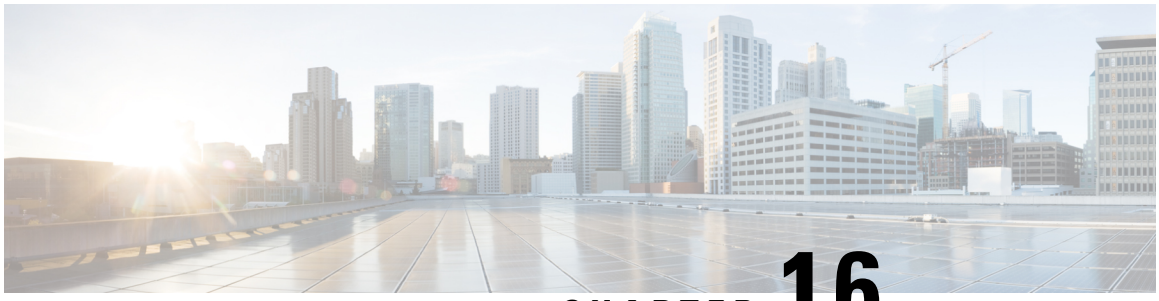
History for backup and restore

This reference provides a chronological history of backup and restore features, including deprecated functionality and version-specific changes across different releases.

Table 48: History for backup and restore

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Deprecated: Cisco AMP Cloud connection backups.	7.2.10 7.0.7	Any	Public and private AMP cloud connections are no longer backed up. You must reconfigure them after restore.
Deprecated: Backup of locally stored reports on the management center.	7.2.11 7.0.8	Any	Locally stored reports are no longer backed up. You should store reports in a secure remote location.
Single backup file for high availability Firewall Management Centers.	7.4.1 7.2.6	Any	When performing a configuration-only backup of the active Firewall Management Center in a high availability pair, the system now creates a single backup file which you can use to restore either unit.
Zero-touch restore for the ISA 3000 using the SD card.	7.0.0	7.0.0	When you perform a local backup, the backup file is copied to the SD card if present. To restore the configuration on a replacement device, simply install the SD card in the new device, and depress the Reset button for 3 to 15 seconds during the device bootup.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Back up and restore FTD container instances.	6.7.0	6.7.0	You can now use the FMC to perform on-demand remote backups of FTD container instances on the Firepower 4100/9300.
No longer need to match VDBs to restore.	6.6.0	Any	Restoring the FMC from backup now replaces the existing VDB with the VDB in the backup file. You no longer need to match VDB versions before you restore.
Automatically scheduled backups.	6.5.0	Any	For new or reimaged FMCs, the setup process creates a weekly scheduled task to back up FMC configurations and store them locally.
On-demand remote backups of managed devices.	6.3.0	6.3.0	You can now use the FMC to perform on-demand remote backups of certain managed devices. For supported platforms, see Requirements: backup and restore configuration, on page 473. New/modified screens: **System** > **Tools** > **Backup/Restore** > **Managed Device Backup** New/modified FTD CLI commands: **restore**



CHAPTER 16

Scheduling

- Task scheduling, on page 505
- Guidelines for scheduled tasks, on page 505
- Scheduled task types, on page 506
- Schedule a task, on page 511
- History for scheduling, on page 511

Task scheduling

Task scheduling is a system management feature that allows you to automate routine jobs to run once or on a recurring basis.

Time zones and seasonal time changes

Tasks are scheduled in Coordinated Universal Time (UTC), which remains constant year-round. Scheduled tasks do not adjust automatically for local time changes, such as daylight saving time or summer time. For example, a task scheduled for 2:00 a.m. during standard time runs at 3:00 a.m. after the local clock moves forward.

Automatically scheduled tasks

The system automatically schedules these tasks:

- A one-time task to download and install the latest VDB, at initial setup.
- A weekly task to download the latest available patches, maintenance releases, and, starting at initial setup.
- A weekly task to perform a locally stored configuration-only backup of the Firewall Management Center, starting at initial setup..
- A daily task to updated certificate revocation lists (CRL), when you configure user or audit log certificates.

Guidelines for scheduled tasks

This section provides general scheduling guidelines. Refer to Scheduled task types, on page 506 for task-specific guidelines and prerequisites.

User role requirements

Admin or Maintenance

When to run tasks

Follow these guidelines when scheduling tasks:

- Automatically scheduled tasks: Review automatically scheduled tasks to make sure they run at the right time for your environment.
- Bandwidth: Schedule tasks that require large amounts of bandwidth during periods of low network use.
- Traffic inspection and flow: Schedule tasks that might interrupt traffic during maintenance windows.

Scheduling task sequences

This table lists outcomes that require multiple tasks in sequence. Leave enough time so that each task can finish before the next begins.

Table 49: Outcomes requiring multiple scheduled tasks

Outcome	Tasks
VDB update	**1.** Download Latest Update **2.** Install Latest Update **3.** Deploy Policies
Software upgrade	**1.** Download Latest Update **2.** Push Latest Update (devices only) **3.** Install Latest Update **4.** Deploy Policies **Note** For more information about software upgrades, including additional restrictions and recommended actions if an upgrade is unresponsive or fails, refer to the upgrade guide for your Firewall Management Center version: Secure Firewall Threat Defense upgrade guides for Firewall Management Center.
Recommended rules update	**1.** Cisco Recommended Rules **2.** Deploy Policies

Scheduled task types

This section describes the available scheduled task types (**Job Type** in the scheduler), as well as options, guidelines, and requirements for each task.

Backup

Purpose: Back up the Firewall Management Center and managed devices. Initial setup schedules a weekly configuration backup of the Firewall Management Center, stored locally.

Options:

- **Backup Type**: Device or Firewall Management Center
- **Backup Profile**: Firewall Management Center backups require a backup profile. Refer to Create a backup profile, on page 482.
- **Retrieve to Management Center**: Devices only. This option controls device backups irrespective of remote storage is enabled or not.
  - Enabled (default): If remote storage is configured, the backup files are saved to remote storage, else the device backup files are saved to the Firewall Management Center in `/var/sf/remote-backup/`.
  - Disabled: Saves device backups to the device in `/var/sf/backup/`.

  For more information about managing backup, refer to Manage backups and remote storage, on page 499.

Guidelines:

- Supported devices: Some devices, such as devices in the public cloud, cannot be backed up. Refer to Requirements: backup and restore configuration, on page 473.
- Prerequisite configurations: Firewall Management Center backups require a backup profile. Refer to Create a backup profile, on page 482.
- Simultaneous backups: Back up no more than 20 devices per task. Do not schedule multiple backup tasks for the same time; start with 30 minutes between backups.
- Bandwidth: If you are transferring backup files, consider scheduling backups during periods of low network use.

Cisco recommended rules

Purpose: Generate rule state recommendations and modify intrusion policies based on network discovery data.

Options:

- **Policies**: The intrusion policies where you want to generate recommendations.

Guidelines:

- Licenses: IPS or Protection
- Prerequisite configurations: You must have at least one custom intrusion policy with recommendations enabled to schedule this task.
- Save changes before the task runs: If your intrusion policies have unsaved changes, recommendations are not applied. Discard your changes and commit the policy to apply recommendations.

- Deploy during a maintenance window to implement changes: Snort may need to restart. Restarting the Snort process briefly interrupts traffic flow and inspection on all devices, including devices configured for high availability or clustering.

Deploy policies

Purpose: Deploy configuration changes from the Firewall Management Center to managed devices. You may want to schedule a deploy after recommended rules, software, or VDB updates.

Options:

- **Device**: The devices where you want to deploy the policies.
- **Skip deployment for up-to-date devices**: Enabled by default to improve performance during the deployment process.

Guidelines:

- Traffic inspection and flow: When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. Refer to Snort restart traffic behavior and Configurations that restart the snort process when deployed or activated.
- Concurrent deployments: Scheduled policy deployments do not run if a manual policy deployment is already in progress. You cannot deploy from the web interface while a scheduled deployment is running.

Download CRL

Purpose: Download certificate revocation list (CRL) updates. The system automatically schedules daily CRL updates when you configure user or audit log certificates in the system configuration. Use the scheduler to change the update interval or run a one-time update. Disabling the configurations removes the task.

Guidelines:

- Internet access: The Firewall Management Center must have internet access to download updates. Refer to Communication Ports and Internet Access, on page 1063.

Download latest update

Purpose: Download the latest maintenance releases, patches, and VDB update. Initial setup schedules a weekly download of the latest applicable updates.

Options:

- **Update Items**: Choose **Vulnerability Database**, **Software**, or both.

Guidelines:

- Supported software upgrade types: You can download maintenance updates and patches. Major upgrades are not supported.
- Required task order: To update the VDB or to upgrade the Firewall Management Center, download then install. To upgrade a managed device, download, push, then install. Each task must finish before the next begins.

- Internet access: The Firewall Management Center must have internet access to download updates. Refer to Communication Ports and Internet Access, on page 1063.
- Bandwidth: Consider scheduling downloads during periods of low network use.

Push latest update

Purpose: Push a downloaded software upgrade package to a managed device so that you can upgrade the device.

Options:

- **Device**: Choose a device or device group.

Guidelines:

- Supported upgrade types: You can push maintenance updates and patches. Major upgrades are not supported.
- Required task order: To upgrade the Firewall Management Center, download then install. To upgrade a managed device, download, push, then install. Each task must finish before the next begins.
- Bandwidth: Consider scheduling pushes during periods of low network use.

Install latest update

Purpose: Install a maintenance release, patch, or VDB update.

Options:

- **Update Items**: Choose **Vulnerability Database** or **Software**.
- **Device**: For software updates, you can choose the Firewall Management Center, a device, or a device group. For VDB updates, you must choose the Firewall Management Center.

Guidelines for VDB updates:

- Required task order: To update the VDB, download then install. Download must finish before install begins.
- Do not perform tasks related to mapped vulnerabilities while the VDB is updating. Even if the Message Center shows no progress for several minutes or indicates that the update has failed, do not restart the update. Instead, contact Cisco TAC.
- Deploy during a maintenance window to implement changes: Snort typically restarts during the first deployment after VDB update. Restarting the Snort process briefly interrupts traffic flow and inspection on all devices, including devices configured for high availability or clustering.

Guidelines for software upgrades:

- Supported upgrade types: You can install maintenance updates and patches. Major upgrades are not supported.
- Required task order: To upgrade the Firewall Management Center, download then install. To upgrade a managed device, download, push, then install. Each task must finish before the next begins.

- Grouped devices: Use simple device groups to upgrade multiple devices with one scheduled task. The system upgrades the targets one at a time. In high availability or clustered deployments, devices use the normal upgrade behavior.
- Upgrade devices during a maintenance window due to effects on traffic inspection and flow: For high availability or clustered devices, traffic inspection and flow are not interrupted during upgrade. For standalone devices, interface configurations determine whether traffic is dropped or passed without inspection.
- Deploy during a maintenance window to implement changes: You must deploy after software upgrades. Snort typically restarts during this first deployment. Restarting the Snort process briefly interrupts traffic flow and inspection on all devices, including devices configured for high availability or clustering.

Nmap scan

Purpose: Nmap scans refresh static operating system, application, and server data previously supplied by Nmap. You can also use these scans to test for unidentified applications and servers.

Options:

- **Nmap Remediation**: The specific Nmap remediation to run.
- **Nmap Target**: The scan target.
- **Domain**: In a multidomain deployment, the domain whose network map you want to augment.

Guidelines:

- Prerequisite configurations: You must configure Nmap scanning to schedule this task. This includes creating an Nmap instance, scan target, and remediation. Refer to Nmap scanning guidelines for additional guidelines.
- Scan regularly: The system does not automatically update network map data replaced by Nmap unless you delete the scan targets from the network map and allow the system to rediscover them. Schedule regular scans to keep your network map current.
- Bandwidth: Consider scheduling expensive scans, such as portscans, during periods of low network use.

Report

Purpose: Schedule report generation.

Options:

- **Report Template**: Use a system-provided template or create your own to generate the report. To get reports by email, edit the report template. The scheduler email option sends only task status and does not send the report by email.
- **If report is empty, still attach to email**: Receive reports as email attachments even when reports have no data; for example, when no events of a certain type occurred during the report period.

Update URL filtering database

Purpose: Obtain the latest URL filtering data from Cisco. By default, when you enable URL filtering, automatic updates are enabled. However, if you need to control exactly when these updates occur, use the scheduler.

Guidelines:

- Licenses: URL Filtering
- Prerequisite configurations: You must enable URL filtering to schedule this task. Disable automatic updates on **Integration** > **Other Integrations** > **Cloud Services**.
- Update size, duration, and bandwidth: Daily updates are typically small. With longer intervals, expect larger downloads and additional time for the changes to propagate, and consider scheduling during periods of low network use.
- Internet access: The Firewall Management Center must have internet access to download updates. Refer to Communication Ports and Internet Access, on page 1063.

Schedule a task

Before you begin

Before you schedule a task:

- Review the guidelines and prerequisites for scheduled tasks in general: Guidelines for scheduled tasks, on page 505.
- Review the guidelines and prerequisites that apply to your task: Scheduled task types, on page 506.
- (Optional) Set up an email relay server to get task status messages: Configure a mail relay host and notification address, on page 71.

Follow these steps to schedule a task.

Procedure

Step 1 Choose **System** (⚙) > **Tools** > **Scheduling** and click **Add Task**.

Step 2 Choose a **Job Type** and configure the task-specific parameters.

Refer to Scheduled task types, on page 506 for the specific options and restrictions for each task.

Step 3 Specify the schedule details: once or recurring, start time, and frequency.

Step 4 Enter a name for the task, and optionally, a comment and email addresses for status notifications.

You can see comments when you view task details in the calendar.

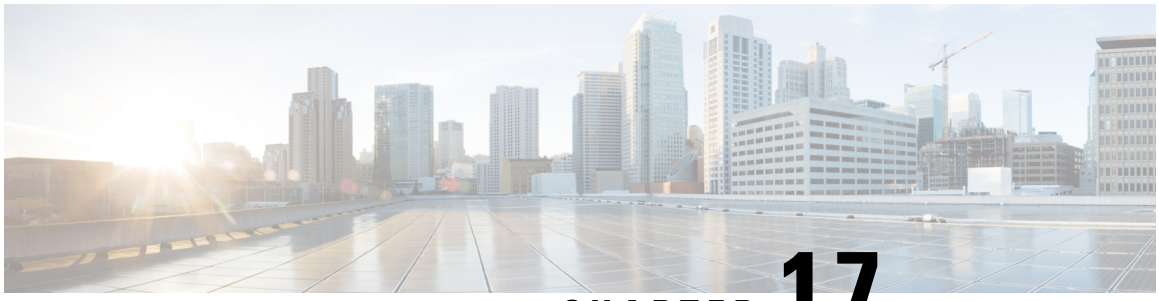
Step 5 Click **Save**.

Your task is scheduled and will run at the configured time.

History for scheduling

This table provides the feature history for the task scheduler.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Automatic software downloads and configuration backups.	6.5	Any	Initial setup schedules weekly tasks to: • Download the latest available software updates for the Firewall Management Center and its managed devices. • Perform a locally stored configuration-only backup. We recommend you review these tasks and adjust as necessary.
Schedule remote backups of many managed devices.	6.4	Any	Schedule device backups. New/modified screens: When configuring a recurring backup, you can now choose a **Backup Type**: Firewall Management Center vs device. Platform restrictions: Device must support on-demand backup. Refer to Requirements: backup and restore configuration, on page 473.



CHAPTER 17

Import/Export

The following topics explain how to use the Import/Export feature:

- Import/Export feature overview, on page 513
- Requirements and prerequisites for import/export, on page 516
- Guidelines for Import/Export, on page 516
- Export configurations, on page 517
- Import configurations, on page 518

Import/Export feature overview

You can use the Import/Export feature to copy configurations between Firewall Management Centers. Import/Export is not a backup tool but can simplify the process of adding a new Firewall Management Center.

Configuration export and import capabilities

You can export a single configuration, or you can export a set of configurations (of the same type or of different types) with a single action. When you later import the package onto another Firewall Management Center, you can choose which configurations in the package to import.

An exported package contains revision information for that configuration, which determines whether you can import that configuration onto another Firewall Management Center. When the Firewall Management Center are compatible, but the package includes a duplicate configuration, the Firewall Management Center offers resolution options.

Exceptions to export behavior

When you export a configuration, the Firewall Management Center also exports other required configurations. For example, exporting an access control policy also exports any subpolicies it invokes, objects and object groups it uses, ancestor policies, and so on. As another example, if you export a platform settings policy with external authentication enabled, the authentication object is exported as well.

There are some exceptions, however:

- System-provided databases and feeds—The Firewall Management Center does not export URL filtering category and reputation data, Cisco Intelligence Feed data, or the geolocation database (GeoDB). Each Firewall Management Center needs to obtain up-to-date information from Cisco.

- Global Security Intelligence lists—The Firewall Management Center exports Global Security Intelligence Block and Do Not Block lists associated with exported configurations. The import process converts these lists to user-created lists, then uses those new lists in the imported configurations. This ensures that imported lists do not conflict with existing Global Block and Do Not Block lists. To use Global lists on the importing Firewall Management Center, manually add the lists to your imported configurations.
- Intrusion policy shared layers—The export process breaks intrusion policy shared layers. The previously shared layer is included in the package, and imported intrusion policies do not contain shared layers.
- Intrusion policy default variable set—The export package includes a default variable set with custom variables and system-provided variables with user-defined values. The import process updates the default variable set on the importing Firewall Management Center with the imported values. However, the import process does **not** delete custom variables not present in the export package. The import process also does not revert user-defined values on the importing Firewall Management Center, for values not set in the export package. Therefore, an imported intrusion policy may behave differently than expected if the importing Firewall Management Center has differently configured default variables.
- Custom user objects—If you have created custom user groups or objects in your Firewall Management Center, and if such a custom user object is a part of any rule in your access control policy, note that the export file (.sfo) does not carry the user object information (including user realms) , and therefore while importing such a policy, any reference to such custom user objects will be removed and will not be imported to the destination Firewall Management Center. To avoid detection issues due to the missing user group, add the customized user objects manually to the new Firewall Management Center, and re-configure the access control policy after import.

Import objects and object groups

When you import objects and object groups:

- Generally, the import process imports objects and groups as new, and you cannot replace existing objects and groups. However, if network and port objects or groups in an imported configuration match existing objects or groups, the imported configuration reuses the existing objects/groups, rather than creating new objects/groups. The Firewall Management Center determines a match by comparing the name (minus any autogenerated number) and content of each network and port object/group.
- If the names of imported objects match existing objects on the importing Firewall Management Center, the Firewall Management Center appends autogenerated numbers to the imported object and group names to make them unique.
- You must map any security zones and interface groups used in the imported configurations to matching-type zones and groups managed by the importing Firewall Management Center.
- If you export a configuration that uses PKI objects containing private keys, the Firewall Management Center decrypts the private keys before export. On import, the Firewall Management Center encrypts the keys with a randomly generated key.

Conflict resolution for duplicate configurations

When you attempt to import a configuration, the Firewall Management Center determines whether a configuration of the same name and type already exists. When an import includes a duplicate configuration, the Firewall Management Center offers resolution options:

Default resolution behavior

- **Keep existing**

  The Firewall Management Center does not import that configuration.

- **Replace existing**

  The Firewall Management Center overwrites the current configuration with the configuration selected for import. This option is not available if the duplicate is in ancestor or descendant domain.

- **Keep newest**

  The Firewall Management Center imports the selected configuration only if its timestamp is more recent than the timestamp on the current configuration. This option is not available if the duplicate is in ancestor or descendant domain.

Note If you import a configuration that contains Microsoft Active Directory users and groups we strongly recommend you download all users and groups after the import to avoid issues in SSL policies, access control policies, and possibly other policies. (**Integration** > **Other Integrations** > **Realms** > **Realms**, then click **Download Now**( ) (**Download Now**).

- **Import as new**

  The Firewall Management Center imports the selected duplicate configuration, appending a system-generated number to the name to make it unique. (You can change this name before completing the import process.) The original configuration remains unchanged.

Note If you modify an imported configuration on the Firewall Management Center and later re-import that configuration to the same Firewall Management Center, you must choose which version of the configuration to keep.

File list resolution behavior

When you import an access control policy with a file policy that uses clean or custom detection file lists, and a file list presents a duplicate name conflict, the Firewall Management Center offers conflict resolution options as described, but the action taken varies as described in the table:

Resolution Option	Resolution Behavior for the Access Control Policy
Keep existing	Unchanged
Replace existing	Imported as new; file lists are merged
Import as new	Imported as new; file lists are merged
Keep newest and access control policy being imported is the newest	Imported as new; file lists are merged

Resolution Option	Resolution Behavior for the Access Control Policy
Keep newest and existing access control policy is the newest	Unchanged

Requirements and prerequisites for import/export

Version requirements

Both Firewall Management Centers must meet version compatibility requirements before you configure import/export.

- Both Firewall Management Centers need to be the same software version.
- For access control and subpolicies (including intrusion policies), ensure the intrusion rule update version matches.

  You cannot use the Import/Export feature to update intrusion rules; you must download and apply the latest rule update version.

Domain requirements

Both Firewall Management Centers must have the same domain hierarchy.



Note If you import a configuration without the same domains, to recover you must first rename the domains and re-export the configuration. Then you can add those new domains to the target Firewall Management Center and import the new configuration.

User roles

- Admin

Guidelines for Import/Export

Supported configurations

You can import and export the following configurations:

- Access control policies and the policies they invoke: prefilter, network analysis, intrusion, SSL, file, Service Policy.
- When exporting policy configuration with dynamic ISE-SGT (Security Group Tag) attributes, ensure that the target Firewall Management Center has the same ISE server configuration. Otherwise, export may lead to object corruption or policy misconfiguration on the target.
- Intrusion policies, independently of access control
- NAT policies

- FlexConfig policies. However, the contents of any secret key variables are cleared when you export the policy. You must manually edit the values of all secret keys after importing a FlexConfig policy that uses secret keys.
- Platform settings
- Health policies
- Alert responses
- Application detectors (both user-defined and those provided by Cisco Professional Services)
- Dashboards
- Custom tables
- Custom workflows
- Saved searches
- Custom user roles
- Report templates
- Third-party product and vulnerability mappings
- Users and groups for user control

Export configurations

Depending on the number of configurations being exported and the number of objects those configurations reference, the export process may take several minutes.

Tip Many list pages include a **YouTube EDU** ( ) icon next to list items. Where this icon is present, you can use it as a quick alternative to the export procedure that follows.

Procedure

Step 1 Choose **System** ( ) > **Tools** > **Import/Export**.

Step 2 Click **Collapse** ( ) and **Expand** ( ) to collapse and expand the list of available configurations.

Step 3 Check the configurations you want to export and click **Export**.

Step 4 Follow your web browser's prompts to save the exported package to your computer.

Import configurations

Depending on the number of configurations being imported and the number of objects those configurations reference, the import process may take several minutes.

Note If you log out of the Firewall Management Center, if you change to a different domain, or if your user session expires after you click **Import**, the import process continues in the background until it is complete. We recommend that you wait for the import process to complete before creating any new objects or policies. Attempting to create them while the import process is still running might result in failures.

Procedure

Step 1 On the importing Firewall Management Center, choose **System** (⚙) > **Tools** > **Import/Export**.

Step 2 Click **Upload Package**.

Step 3 Enter the path to the exported package or browse to its location, then click **Upload**.

Step 4 If there are no version mismatches or other issues, choose the configurations you want to import, then click **Import**.
If you do not need to perform any conflict resolution or interface object mapping, the import completes and a success message appears. Skip the rest of this procedure.

Step 5 If prompted, on the Import Conflict Resolution page, map interface objects used in the imported configurations to zones and groups with matching interface types managed by the importing Firewall Management Center.

Interface object type (security zone or interface group) and interface type (passive, inline, routed, and so on) of source and destination objects must match. For information, see Interfaces.

If the configurations you are importing reference security zones or interface groups that do not already exist, you can map them to existing interface objects, or create new ones.

Note
For individual access control policies, you have the option of replacing an existing policy with the imported ones. However, for nested access control policies, you can only import them as new policies.

Step 6 Click **Import**.

Step 7 If prompted, on the Import Resolution page, expand each configuration and choose the appropriate option as described in Conflict resolution for duplicate configurations, on page 514.

Step 8 Click **Import**.

Step 9 Update all feeds.

For example, go to **Objects** > **Object Management** > **Security Intelligence** and click the **Update Feed** button on the URL, Network, and DNS Lists and Feeds pages.

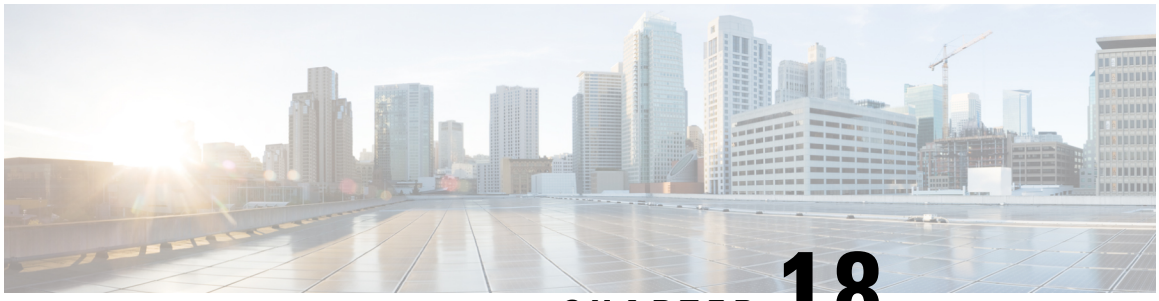
Imported policies do not include feed contents.

Step 10 Wait for all feed updates to complete before deploying the policies to devices.

What to do next

Note If you import a configuration that contains Microsoft Active Directory users and groups we strongly recommend you download all users and groups after the import to avoid issues in SSL policies, access control policies, and possibly other policies. (**Integration** > **Other Integrations** > **Realms** > **Realms**, then click **Download Now**( ) (**Download Now**).

- Optionally, view a report summarizing the imported configurations; see View task messages, on page 450.



CHAPTER 18

Data Purge and Storage

Data purge and storage is a system management process that controls the retention, archiving, and removal of historical configuration data and operational logs within the management environment.

- Categorization of remote data storage options.
- Management of historical configuration updates.
- Tracking of feature deprecations and enhancements.

- Data stored on the Firewall Management Center, on page 521
- External data storage, on page 523
- History for data storage, on page 525

Data stored on the Firewall Management Center

This reference provides information about the various types of data stored on the Firewall Management Center and directs you to relevant documentation for managing each data type.

For	See
General information about data storage on the Firewall Management Center	Disk usage widget categories and features, on page 367
Purging old data	Purge data from the Firewall Management Center Database , on page 522
Allowing external access to the data on the Firewall Management Center (this is an advanced feature)	External database access, on page 72
Backups	Manage backups and remote storage, on page 499 and subtopics
Reports	Configure local storage, on page 114
Events	Connection Logging, on page 737 Database, on page 67 and subtopics

For	See
Network discovery data	*Network Discovery Data Storage Settings* and subsequent topics in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Files	Information about storing files in the *Network Malware Protection and File Policies* chapter of the *Cisco Secure Firewall Management Center Device Configuration Guide*, including best practices. *Tuning File and Malware Inspection Performance and Storage* *Cisco Secure Firewall Management Center Device Configuration Guide*
Packet data	*Edit General Settings* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Users and user activity	*The Users Database* in the *Cisco Secure Firewall Management Center Device Configuration Guide* *The User Activity Database* in the *Cisco Secure Firewall Management Center Device Configuration Guide*

Purge data from the Firewall Management Center Database

Use the database purge page to purge discovery, identity, connection, and security intelligence data files from the Firewall Management Center databases. Note that the respective processes restart with purging.

Caution Purging a database removes the data you specify from the Firewall Management Center. After the data is deleted, it *cannot* be recovered.

Before you begin

You must have Admin or Security Analyst privileges to purge data. Perform this action in the global domain.

Follow these steps to purge data from the Firewall Management Center database:

Procedure

Step 1 Choose **System** (⚙) > **Tools** > **Data Purge**.

Step 2 Under **Discovery and Identity**, perform any or all of these actions:

- Check the **Network Discovery Events** check box to remove all network discovery events from the database.
- Check the **Hosts** check box to remove all hosts and Host Indications of Compromise flags from the database.

- Check the **User Activity** check box to remove all user activity events from the database.
- Check the **User Identities** check box to remove all user login and user history data from the database, as well as User Indications of Compromise flags.

Step 3 Under **Connections**, perform any or all of these actions:

- Check the **Connection Events** check box to remove all connection data from the database.
- Check the **Connection Summary Events** check box to remove all connection summary data from the database.
- Check the **Security Intelligence Events** check box to remove all security intelligence data from the database.

Note
Checking the **Connection Events** check box does not remove Security Intelligence events. Connections with Security Intelligence data will still appear in the Security Intelligence event page (available under the **Analysis** > **Connections** > **Security Intelligence Events** menu). Correspondingly, checking the **Security Intelligence Events** check box does not remove connection events with associated security intelligence data.

Step 4 Click **Purge Selected Events**.
Once the items are purged, the appropriate processes restart.

The selected data types are permanently removed from the Firewall Management Center database to optimize storage and performance.

External data storage

Remote data storage supports storing data under different categories. The classification is as follows:

For	See
Backups	Manage backups and remote storage, on page 499 and subtopics Remote storage devices, on page 114 and subtopics
Reports	Remote storage devices, on page 114 and subtopics Moving reports to remote storage, on page 558
Events	Information about syslog and other resources in Event Analysis Using External Tools, on page 635 Remote data storage in Cisco Secure Cloud Analytics, on page 524 Remote data storage on a Secure Network Analytics appliance, on page 525 If you store connection events remotely, consider disabling storage of connection events on your Firewall Management Center. For information, see Database, on page 67 and subtopics.

Important If you will use syslog or store events externally, avoid special characters in object names such as policy and rule names. Object names should not contain special characters, such as commas, that the receiving application may use as separators.

Comparison of Security Analytics and Logging remote event storage options

Different options for storing event data externally to your Firewall Management Center:

On Premises	SaaS
You purchase, license, and set up the storage system behind your firewall.	You purchase licenses and a data storage plan and send your data to the Cisco cloud.
Supported event types: • **Connection** • **Security Intelligence** • **Intrusion** • **File and Malware** • **LINA**	Supported event types: • Connection • Security Intelligence • Intrusion • File and Malware
Supports both syslog and direct integration.	Supports both syslog and direct integration.
• **View all events on the Secure Network Analytics Manager.** • **Cross-launch from FMC event viewer to view events on the Secure Network Analytics Manager.** • **View remotely stored connection and Security Intelligence events in FMC**	View events in Security Cloud Control or Secure Network Analytics, depending on your license. Cross-launch from FMC event viewer.
For more information, see the links in Remote data storage on a Secure Network Analytics appliance, on page 525**.**	For more information, see the links in Remote data storage in Cisco Secure Cloud Analytics, on page 524.

Remote data storage in Cisco Secure Cloud Analytics

Remote data storage in Cisco Secure Cloud Analytics is a security data management capability that

- sends select Secure Firewall event data to Secure Cloud Analytics using Security Analytics and Logging (SaaS),
- supports connection, Security Intelligence, intrusion, file, and malware events, and
- allows event transmission either directly or via syslog.

Additional information

For more information, refer to *Secure Firewall Management Center and Cisco Security Analytics and Logging (SaaS) Integration Guide*.



Important If you will use syslog or store events externally, avoid special characters in object names such as policy and rule names. Object names should not contain special characters, such as commas, that the receiving application may use as separators.

Remote data storage on a Secure Network Analytics appliance

A remote data storage on a Secure Network Analytics appliance is a storage solution that

- provides additional data storage capacity when your Secure Firewall appliance storage is insufficient
- uses Security Analytics and Logging (On Premises) to store Secure Firewall data externally, and
- allows you to view connection events in Firewall Management Center even when data is stored remotely.

Remote data storage capabilities

You can view connection events in Firewall Management Center even if they are stored on a Secure Network Analytics appliance. See Work in Firewall Management Center with connection events stored on a Secure Network Analytics appliance, on page 685.

For complete information, see the documentation available from Cisco Security Analytics and Logging.



Note If you will use syslog or store events externally, avoid special characters in object names such as policy and rule names. Object names should not contain special characters, such as commas, that the receiving application may use as separators.

History for data storage

This reference details the historical updates in data storage functionalities and configurations. It provides insights into deprecated and improved features across various product versions.

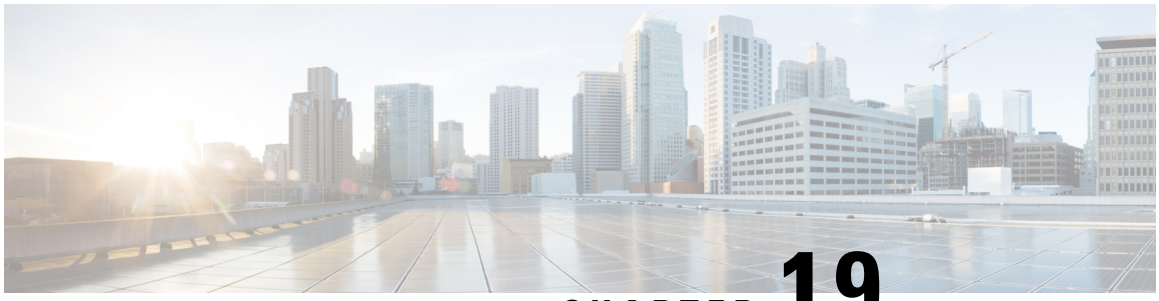
Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Exempt low priority connection events from event rate limits	7.0	Any	If you choose not to store connection events on the Firewall Management Center because you are storing them on a remote volume, those events do not count towards the flow rate limits for your Firewall Management Center hardware device. If you send events to Security Analytics and Logging (On Premises) using the new 7.0 configurations, you configure this setting as part of that integration. Otherwise, see information about the Connection Database in Database event limits, on page 68. New/Modified pages: None. Behavior change only.
Improved process for sending events to a Secure Network Analytics appliance	7.0	Any	A new wizard streamlines sending events directly to a Secure Network Analytics appliance using Security Analytics and Logging (On Premises). The wizard also allows you to see remotely stored connection events while viewing event pages on your Firewall Management Center, and to cross-launch from Firewall Management Center to view events on your Secure Network Analytics appliance. If you have already configured your system to send events using syslog, events will continue to be sent using syslog unless you disable those configurations. For details, see the documentation referenced in Remote data storage on a Secure Network Analytics appliance, on page 525. New/Modified pages: The **System** > **Logging** > **Security Analytics & Logging** page now displays the wizard instead of the configuration for creating cross-launch options.
Remote data storage on a Secure Network Analytics appliance	6.7	Any	You can now store large volumes of Firepower event data remotely, using Security Analytics and Logging (On Premises). When viewing events in Firewall Management Center, you can quickly cross-launch to view events in your remote data storage location. Supported events: Connection, Security Intelligence, intrusion, file, and malware. Events are sent using syslog. This solution depends on availability of Stealthwatch Management Console (SMC) Virtual Edition running Stealthwatch Enterprise (SWE) version 7.3. See Remote data storage on a Secure Network Analytics appliance, on page 525.
Remote data storage in Cisco Secure Cloud Analytics	6.4	Any	Use syslog to send select Firepower data using Security Analytics and Logging (SaaS). Supported events: Connection, Security Intelligence, intrusion, file, and malware. For details, see the *Firepower Management Center and Cisco Security Analytics and Logging (SaaS) Integration Guide* at https://cisco.com/go/firepower-sal-saas-integration-docs.



PART V

Reporting and Alerting

- Reports, on page 529
- External Alerting with Alert Responses, on page 561
- External Alerting for Intrusion Events, on page 573



CHAPTER 19

Reports

The following topics describe how to work with reports:

- Requirements and prerequisites for reports, on page 529
- Reports, on page 529
- Risk reports, on page 530
- Standard reports, on page 531
- Generated reports, on page 556
- Report template feature history, on page 559

Requirements and prerequisites for reports

Lists the requirements and user roles needed to generate and access reports in the system, including special considerations for risk-related reporting.

Model support

Any.

Supported domains

Any

User roles

- Admin
- Maintenance User (risk reports only)
- Security Analyst

Reports

The system offers two types of reports:

- Risk reports, on page 530: Offers high-level summaries of risks found on your network.
- Standard reports: Provides detailed, customizable reports about all aspects of your system.

Risk reports

A risk report is a portable, high-level summary that

- presents high-level, easy-to-interpret overviews of risks identified within your organization,
- includes recommendations for addressing identified security concerns, and
- enables sharing of risk information with individuals who may not have access to your security system or specialist expertise.

Risk report usage

You can use risk reports to share information about areas of risk and recommendations for addressing these risks with stakeholders who may not be network security experts.

Risk report templates

This topic provides a brief description of various risk report templates.

- **Advanced Malware Risk Report**
- **Attacks Risk Report**. These are the fields in this report:
  - **Total Attacks**: The total number of IPS events.
  - **Relevant Attacks**: The number of IPS events with impact flag equal to 1.
  - **Hosts Targeted**: The number of unique destination IP addresses from IPS events with impact flag equal to 1.
  - **Irrelevant Attacks**: The percentage of IPS events with impact flag not equal to 1.
  - **Events Requiring Attention**: The percentage of IPS events with impact flag equal to 1.
  - **Hosts Connected to CnC Servers**: The total number of unique hosts with the IOC category: "`CnC Connected`."
- **Network Risk Report**

Generate, view, and print risk reports

Templates for standard reports do not apply to risk reports.

Reports pertain to the current domain.

Each risk report generates as an HTML file.

To schedule risk report generation, refer to Scheduling, on page 505.

Before you begin

- Make sure your system is configured to detect the risks that you want to summarize.

- If you want to email the report and you have not yet configured a relay host, you can do so now. For information, refer to Configure a mail relay host and notification address, on page 71.

Follow these steps to generate, view, and print risk reports:

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Report Templates**.

Step 3 Click **Generate Report** for the desired report.

Step 4 Enter information.

- Information that you enter in the Input Parameters section will appear on the title page of the report. You can leave these fields blank.

Step 5 Click **Generate**.

Step 6 Click **OK**.

What to do next

- To view, download, move, or delete a risk report, refer to Generated reports, on page 556.
- You can print to PDF any risk report from most supported browsers. For best results, enable background colors, images, and optionally headers and footers, in the print or print preview settings of your browser. Supported page sizes are A4 and US letter.

Standard reports

A report is a document file formatted in PDF, HTML, or CSV with the content you want to communicate. A report template specifies the data searches and formats for the report and its sections. The system includes a powerful report designer that automates the design of report templates. You can replicate the content of any event view table or dashboard graphic displayed in the web interface.

Report templates and customization

The system provides a flexible reporting system that allows you to quickly and easily generate multi-section reports with the event views or dashboards that appear on your Firewall Management Center. You can also design your own custom reports from scratch.

You can build as many report templates as you need. Each report template defines the individual sections in the report and specifies the database search that creates the report's content, as well as the presentation format (table, chart, detail view, and so on) and the time frame.

Your template also specifies document attributes, such as the cover page and table of contents and whether the document pages have headers and footers (available only for reports in PDF format). You can export a report template in a single configuration package file and import it for reuse on another Firewall Management Center.

You can include input parameters in a template to expand its usefulness. Input parameters allow you to produce tailored variations of the same report. When you generate a report with input parameters, the generation process prompts you to enter a value for each input parameter.

The values you type constrain the report contents on a one-time basis. For example, you can place an input parameter in the destination IP field of the search that produces an intrusion event report; at report generation time, you can specify a department's network segment when prompted for the destination IP address. The generated report then contains only information concerning that particular department.

About designing reports

A report is a structured information resource that

- presents network data, performance metrics, and system information,
- enables analysis and documentation by organizing information in a readable format, and
- can be customized to highlight specific data or trends as required.

Report templates

A report template is a report configuration tool that

- enables you to define the content and format for each section of a report,
- allows you to configure document attributes such as cover page, table of contents, and page headers and footers, and
- remains available for future reuse until you delete it.

Report template design and format options

Your reports contain one or more information sections. You choose the format (text, table, or chart) for each section individually. The format you select for a section may constrain the data that can be included. For example, you cannot show time-based information in certain tables using a pie chart format. You can change the data criteria or format of a section at any time to obtain optimum presentation.

You can base a report's initial design on a predefined event view, or you can start your design by importing content from any defined dashboard, workflow, or summary. You can also start with an empty template, adding sections and defining their attributes one by one.



Note In a multidomain deployment, you can view but not edit report templates belonging to ancestor domains. To generate reports from these templates, you must copy them to your current domain.

Report template section fields

The table describes the fields you can use to build sections in a report template. Some fields are available only for specific section types. After you select a section format, the system displays the appropriate fields.

Field Name	Section Types	Definition
Format	n/a	Choose the format of the section data: **Bar chart** (): Compares quantities of the selected variables. **Line chart** (): Shows trends/changes over time of a selected variable. Available only for time-based tables. **Pie chart** (): Shows each selected variable as a percentage of the whole. Variables with quantities of zero are dropped from the chart. Very small quantities are clustered into a category labeled **Other**. **Table view** (): Shows values of attributes for each record. Not available for summary or statistical data. **Detail view** (): Shows complex object data associated with certain events, such as packets (for intrusion events) and host profiles (for host events). This format is available only for certain event types that involve such objects. Output may degrade performance if large numbers are requested.
Table	All	Choose the table from which the section data is extracted.
Preset	All	Predefined searches. Select an appropriate preset to initialize the search criteria when you define a new search.
Search or Filter	All	For most tables, you can constrain a report using a predefined or saved **Search**. You can also create a new search by clicking **Edit** (). For the Application Statistics table, you use a user-defined application **Filter** to constrain a report.
X-Axis	Bar chart Line chart Pie chart	Available data for the X-axis of the selected chart. For line charts, the X-axis value is always **Time**. For bar and pie charts, you cannot select **Time** as the X-axis value.
Y-Axis	Bar chart Line chart Pie chart	Available data for the Y-axis of the selected chart.
Section Description	All	Descriptive text that precedes the search data in the section. Enter a combination of text and input parameters. The default for a new section is `$<Time Window>` and `$<Constraints>`.
Time Window	All	The time window for the data that appears in the section. If the section searches time-based tables, you can select the check box to inherit the report's global time window. Alternatively, you can set a specific time window for the section.

Field Name	Section Types	Definition
Data Source	All	If you used the wizard to configure remote (external) data storage using Security Analytics and Logging (On Premises), you can choose the data source to use for connection and Security Intelligence events. Options are: • **Auto**: Show data stored on the Firewall Management Center if available. If data on the Firewall Management Center is not available for the entire selected time window, show only remotely stored data. • **Local**: Show only data that is stored on the Firewall Management Center, regardless of the time window selected. Choose this option to include data that is not available on the remote volume, such as events generated from devices that are not configured to send events to the remote volume. • **Extended**: Show only data that is stored on the remote volume.
Maximum Results	Table view Detail view	The maximum number of matching records to include. You can include fewer records in a PDF report than in a CSV or HTML report. The web interface uses warning and error icons to indicate when the number is too large. Hover your pointer over the icon to see the limits.
Results	Bar chart Pie chart	Choose either **Top** or **Bottom** and enter the number of matching records you want to use to build the chart.
Color	Bar chart Line chart	Colors for graphed data in the section.

Report templates

A report template is a framework for organizing report sections that

- allows each section to be independently built from its own database query,
- supports creation through several methods, such as starting from an existing template, importing dashboards or workflows, or basing on an event view, and
- provides features for customizing the format and content of each section.

Report template creation methods

You can build a new report template by creating a new template, using an existing template, basing a template off an event view, or importing a dashboard or workflow.

Report template creation methods include:

- **New template creation**: Create an entirely new template by generating the framework and then designing individual template sections
- **Existing template copying**: Copy an existing template and edit its attributes to create a new report template

- **Event view-based creation**: Create a report template from an event view and modify it by adding, modifying, or deleting sections
- **Import-based creation**: Import dashboards, workflows, and statistics summaries to quickly create reports with sections for each widget graphic or event view

If you do not want to copy an existing report template, you can create an entirely new template. The first step in creating a template is to generate the framework that allows you to add and format the sections. Then, in the order you prefer, you design the individual template sections and set attributes for the report document.

Each template section consists of a dataset generated by a search or filter, and has a format specification (table, pie chart, and so on) that determines the mode of presentation. You further determine section content by selecting the fields in the data records you want to include in the output, as well as the time frame and number of records to show.



Note Use the section preview utility to check the column selection and output characteristics such as pie chart colors. It is not a reliable indicator of the correctness of your configured search.

The report you generate from the template has several document attributes that span all sections and control features, such as the cover page, headers and footers, page numbering, and so on.

Note that if you selected CSV as your document format, you have no document attributes to set.

If you identify a good model among your existing templates, you can copy the template and edit its attributes to create a new report template. Cisco also provides a set of predefined report templates, visible on the **Reports Tab** in the list of templates.

From an event view, you can create a report template and modify it to meet your needs. You can add additional sections, modify automatically included sections, and delete sections.

You can quickly create a new report by importing dashboards, workflows, and statistics summaries. The import creates a section for each widget graphic in your dashboard and each event view in your workflow. You can delete any unnecessary sections to focus on the most important information.

Create a custom report template

Custom report templates allow you to define the structure, content, and format of reports. You can configure sections, add input parameters, and set attributes for different output formats.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Report Templates**.

Step 3 Click **Create Report Template**.

Step 4 Enter a name for your new template in the **Report Title** field.

Step 5 To add an input parameter to the report title, place your cursor in the title where the parameter value should appear, then click insert **Input Parameter** (+).

Step 6 Use the set of add under the Report Sections title bar to insert sections as necessary.

Step 7 Configure section content as described in Report templates, on page 538.

Tip
You can click **Preview** at the bottom of the section window to view the column layout or graphic format you chose.

Step 8 Click **Advanced** to set attributes for PDF and HTML reports as described in Document attributes in report templates, on page 548.

Step 9 Click **Save**.

If you see an error, look for a yellow triangle beside the results value in each section. If you see any such triangles, do one of these:

- For each field that displays a yellow triangle, mouse over the triangle and reduce the number of results to the number indicated.
- Click **Generate** and include an output format other than PDF.

Create a report template from an existing template

Use this procedure when you need a new report template that is similar to an existing one, saving time by starting with a pre-configured template rather than creating one from scratch.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Report Templates**.

Step 3 Click **Copy** () next to the report template you want to copy.

Step 4 In the **Report Title** field, enter a name.

Step 5 Make changes to the template as needed.

Step 6 Click **Save**.

Create a report template from an event view

This task guides you through creating a report template based on selected events in an event view. You can customize report sections and layout, enhancing reporting flexibility.

Procedure

Step 1 Populate an event view with the events you want in the report:

- Use an event search to define the events you want to view.
- Drill down through a workflow until you have the appropriate events in your event view.

Step 2 From the event view page, click **Report Designer**.

The Report Sections page displays a section for each view in the captured workflow.

Step 3 Optionally, enter a new name in the **Report Title** field and click **Save**.

Step 4 You can:

- Add a cover page, table of contents, starting page number, or header and footer text: Click **Advanced** settings.
- Add page breaks: Click **Add Page Break** (), and drag the new page break object from the template bottom to the front of the section that should start the new page.
- Add text sections: Click **Add Text Section** ( ), and drag the new text section from the template bottom to the place where you want it to appear in the report template.
- Change the title of a section: Click the section title in the title bar, enter the section title, and click **OK**.
- Configure the report sections: Adjust the field settings in each section.

  Tip
  To view the current column layout or chart formatting for a section, click the section's **Preview** link.

- Exclude template sections from the report: Click **Delete** () in the section's title bar, and confirm the deletion.

  Note
  The last report section in some workflows contains detail views that show packets, host profiles, or vulnerabilities, depending on the workflow. Retrieving large numbers of events with these detail views when generating your report may affect performance of the Firewall Management Center.

Step 5 Click **Save**.

Create a report template by importing a dashboard or workflow

Use this task when you want to create standardized, reusable report templates by pulling in content from dashboards, workflows, or summaries. This feature lets you import widgets or event views as report sections and customize them.

Procedure

Step 1 Identify the dashboard, workflow, or summary you want to replicate in your report.

Step 2 Choose **Overview** > **Reporting**, and click **Report Templates**.

Step 3 Click **Create Report Template**.

Step 4 Enter a name for your new report template in the **Report Title** field, and click **Save**.

Step 5 Click **Import Section** ().

You can choose any of the data sources described in Data source options in the import report sections window, on page 538.

Step 6 Choose a dashboard, workflow, or summary from the drop-down menus.

Step 7 For the data sources you want to add, click **Import**.

For dashboards, each widget graphic will have its own section; for workflows, each event view will have its own section.

Step 8 Make changes to the content of your sections as needed.

Note
The last report section in some workflows contains detail views that show packets, host profiles, or vulnerabilities, depending on the workflow. Retrieving large numbers of events with these detail views when generating your report may affect performance of the Firewall Management Center.

Step 9 Click **Save**.

Data source options in the import report sections window

The import report sections window provides several data source options that enable you to import specific content into your reports.

The table summarizes each option and the type of data it imports.

Table 50: Data source options on import report sections window

Select this option...	To import...
Import Dashboard	any custom analysis widget on the selected dashboard.
Import Workflow	any predefined or custom workflow. Selections have the format: `Table - Workflow name` For example, **Connection Events - Traffic by Port** imports the views in the Traffic by Port workflow generated from the Connection Events table.
Import Summary Sections	any of these generic summaries: • Intrusion Detailed Summary • Intrusion Short Summary • Discovery Detailed Summary • Discovery Short Summary

Report templates

Report template configuration is a customization capability that

- enables modification of report section attributes to tailor content and presentation,
- queries database tables to generate content for each section, and
- allows users to define time windows, searches, filters, and data fields displayed.

Report template configuration capabilities

Each section in a report template queries a database table to generate content for that section. Changing the section's data format uses the same data query, but modifies the fields that appear in the section according to the analytical purpose of the format type. For example, the table view of intrusion events populates the section with a large number of data fields per event record, while a pie chart section shows the portion of all matching records that each selected attribute represents, with no details about individual events.

Bar chart sections compare the total counts of matching records that have specific attributes. Line charts summarize changes in the matching records over time with respect to a single attribute. Line charts are available only for data that is time-based, not for information about hosts, users, third-party vulnerabilities, and so on.

The search or filter in a report section specifies the database query on which the section content is based. For most tables, you can constrain a report using a predefined or saved search, or you can create a new search on the fly:

- Predefined searches serve as examples for searching certain event tables and can provide quick access to important information about your network that you may want to include in reports.
- Saved event searches include all public event searches that you or others have created, plus all your saved private event searches.
- Saved searches for the current report template are accessible only in the report template itself. The search names of saved report template searches end with the string "Custom Search." Users create these searches while designing reports.

For the Application Statistics table, you use a user-defined application filter to constrain a report.

If you include table data in a section, you can choose which fields in the data record to show. All fields in the table are available for inclusion or exclusion. You select fields that accomplish the purpose of the report, then order and sort them accordingly.

You can add text sections to your templates to provide custom text, such as an introduction, for the whole report or for individual sections.

You can add page breaks before or after any section in the template. This feature is particularly helpful for multi-section reports with text pages that introduce the various sections.

A report template's time window defines the template's reporting period.

Note Security Analysts can edit only report templates they created. In multidomain deployments, you cannot edit report templates from ancestor domains, but you can copy them to create descendant versions.

Configure table and data format in a report template section

Report template sections require configuration of both the data source table and the output format to properly display reporting information. The available output formats depend on the selected table type.

Procedure

Step 1 Choose **Overview** > **Reporting**, click the **Report Templates** tab, and then click **Create Report Template**.

Step 2 In the report template section, use the **Table** drop-down menu to choose the table to query.

The **Format** field represents each of the output formats available for the table you chose.

Step 3 Choose the applicable output format for the section.

Step 4 To change the search constraints, click **Edit** ( ) next to the **Section description field** or **Filter** field.

Step 5 For graphic output formats (pie chart, bar chart, and so on), adjust the **X-Axis** and **Y-Axis** parameters using the drop-down menus.

When you choose a value for the X-axis, only compatible values appear in the Y-axis drop-down menu, and vice versa.

Step 6 For table output, choose the columns, order of appearance, and sort order in your output.

Step 7 Click **Save**.

Related Topics

Report template section fields, on page 532

Specify search or filter criteria for a report template

Use this task when creating or customizing a report template to ensure the resulting report displays only the data you need.

Procedure

Step 1 In the report template section, choose the database table to query from the **Table** drop-down menu:

- For most tables, the **Search** drop-down list appears.
- For the Application Statistics table, the **Filter** drop-down list appears.

Step 2 Choose the search or filter you want to use to constrain the report.

You can view the search criteria or create a new search by clicking **Edit** ( ).

Modify fields in a report table

Use this task to tailor report tables to show the information that matters most for your analysis or audience.

Procedure

Step 1 For table-format report sections, click the **Edit** ( ) icon next to the **Fields** parameter.

Step 2 If you want to modify the section, you must add and delete columns, and drag a column into the order you want.

Step 3 If you want to change the sort order of any column, you must use the drop-down lists next to each column name to set the sort order and priority.

Step 4 Click **OK**.

Add a text section to a report template

Text sections allow you to include introductions or explanatory content in your report templates. These sections support varied font styles, input parameters that update dynamically, and imported images for enhanced presentation.

Tip Text sections are useful for introductions to your report or your report sections.

Procedure

Step 1 In the report template editor, click **Add Text Section** ( ).

Step 2 Drag the new text section to its intended position in the report template.

Step 3 If you want to position the text section first or last on a page, add page breaks before or after the text section.

Step 4 If you want to change the text section's generic name, click section's name in the title bar, and enter a new name.

Step 5 Add formatted text and images to the body of the text section.

You can include input parameters that dynamically update when you generate the report.

Step 6 Click **Save**.

Related Topics

Input parameters, on page 544

Add a page break to a report template

Use page breaks in report templates to visually separate sections, making reports more readable and easier to review.

Procedure

Step 1 In the report template editor, click **Add Page Break** ( ).

A page break appears at the bottom of the template.

Step 2 Drag the page break to its intended location, before or after a section.

Step 3 Click **Save**.

Global time windows in report template sections

A global time window is a report template feature that

- applies to report templates with time-based data such as intrusion or discovery events,
- provides default time settings that time-based sections inherit when created, and
- can be modified to change local time windows for sections configured to inherit the global setting.

Time window inheritance behavior

Changing the global time window changes the local time window for the sections that are configured to inherit the global time window. You can disable time window inheritance for an individual section by clearing its **Inherit Time Window** check box. You can then edit the local time window.

Note Global time window inheritance applies only to report sections with data from time-based tables, such as intrusion events and discovery events. For sections that report on network assets (hosts and devices) and related information (such as vulnerabilities), you must set each time window individually.

Set the global time window for a report template and its sections

A report template can contain multiple sections, each with its own time range. For instance, you might set the first section as a monthly summary and configure other sections to provide weekly details. This process enables you to modify the global time window for the template while also allowing section-level adjustments.

Procedure

Step 1 In the report template editor, click **Generate**.

Step 2 To modify the global time window, click **Time Window** ( ).

Step 3 Modify time settings in **Events Time Window**.

Step 4 Click **Apply**.

Step 5 Click **Generate** to generate the report and **Yes** to confirm.

Set the local time window for report template sections

Adjusting the time window allows you to focus each section on a relevant date range, especially for statistical data that supports only sliding windows.

Procedure

Step 1 On the Report Sections page of a template, clear the **Inherit Time Window** check box for the section if it is present.

Step 2 To change the section's local time window, click **Time Window** ( ).

Note
Sections with data from statistics tables can have only sliding time windows.

Step 3 Click **Apply** on the Events Time Window.

Step 4 Click **Save**.

Rename a report template section

Use this task when reorganizing report templates or when a section name no longer fits its content.

Procedure

Step 1 In the report template editor, click the current section name in the section header.

Step 2 Enter a new name for the section.

Step 3 Click **OK**.

Preview a report template section

The preview function shows the field layout and sort order for table views and important legibility characteristics of graphics, such as pie chart colors.

Procedure

Step 1 At any time while editing a report template section, click **Preview** for the section.

Step 2 Close the preview by clicking **OK**.

Report template searches

A report template search is a configuration feature that

- defines the criteria for retrieving data to populate sections of a report,
- enables the inclusion of both predefined and custom searches within a report template, and
- utilizes a search editor to manage, view, and customize available searches.

Search editor capabilities

The system provides a search editor, allowing you to view existing searches in report templates and to define new custom searches tailored to your reporting needs.

Search in report template sections

Use this task when you need to filter or locate specific data in a report template section, either by customizing search criteria or basing your search on predefined queries.

Procedure

Step 1 From the relevant section in the report template, click **Edit** ( ) next to the **Search** field.

Step 2 If you want to base a custom search on a predefined search, you must choose a predefined search from the **Saved Searches** drop-down list.

This list includes all available predefined searches for this table, including system-wide and report-specific predefined searches.

Step 3 Edit the search criteria in the appropriate fields.

For certain fields, your constraints can include the same operators (<, <>, and so on) as event searches. If you enter multiple criteria, the search returns only the records that match all the criteria.

Step 4 If you want to insert an input parameter from the drop-down menu instead of entering a constraint value, you must click **Input Parameter** ( ).

Note
When you edit the constraints of a reporting search, the system saves your edited search under the following name: *section* `custom search`, where *section* is the name in the section title bar followed by the string `custom search`. To have meaningful names for your saved custom searches, be sure you change the section name before you save the edited search. You cannot rename a saved reporting search.

Step 5 Click **OK**.

Input parameters

An input parameter is a report template attribute that

- allows the report to collect dynamic values from the user during generation,
- supports both predefined (system-resolved) and user-defined (customizable) constraints, and
- enables customization of report sections and fields such as search filters, email text, and file names.

Input parameter types

There are two types of input parameters:

- *Predefined input parameters*: These are resolved by internal system functions or configuration information. For example, at report generation time, the system replaces the `$<Time>` parameter with the current date and time.
- *User-defined input parameters*: These supply constraints in report section searches. Constraining a search with an input parameter instructs the system to collect a value at generation time from the person who requests the report. In this way, you can dynamically tailor a report at generation time to show a particular

subset of data without changing the template. For example, you can provide an input parameter for the **Destination IP** field of a report section's search. Then, when you generate the report, you can enter the IP network segment for a particular department to get data for that department only.

You can also define string-type input parameters to add dynamic text in certain fields of your report, such as in emails (subject or body), report file names, and text sections. You can personalize reports for different departments, with customized report file names, email addresses, and email messages, using the same template for all.

Report template predefined input parameters

This reference provides information about predefined input parameters that can be inserted into report templates to include dynamic content such as logos, dates, system information, and other contextual data.

Table 51: Predefined input parameters

Insert this parameter...	...to include this information in your template:
`$<Logo>`	The selected uploaded logo
`$<Report Title>`	The report title
`$<Time>`	The date and time of day the report ran, with one-second granularity
`$<Month>`	The current month
`$<Year>`	The current year
`$<System Name>`	The name of the Firewall Management Center
`$<Model Number>`	The model number of the Firewall Management Center
`$<Time Window>`	The time window currently applied to the report section
`$<Constraints>`	The search constraints currently applied to the report section

Table 52: Predefined input parameter usage

Parameter	Report Template Cover Page	Report Template Report Title	Report Template Section Description	Report Template Text Section	Generate Report File Name	Generate Report Email Subject, Body
`$<Logo>`	yes	no	no	no	no	no
`$<Report Title>`	yes	no	yes	yes	yes	yes
`$<Time>`	yes	yes	yes	yes	yes	yes
`$<Month>`	yes	yes	yes	yes	yes	yes
`$<Year>`	yes	yes	yes	yes	yes	yes
`$<System Name>`	yes	yes	yes	yes	yes	yes

Parameter	Report Template Cover Page	Report Template Report Title	Report Template Section Description	Report Template Text Section	Generate Report File Name	Generate Report Email Subject, Body
`$<Model Number>`	yes	yes	yes	yes	yes	yes
`$<Time Window>`	no	no	yes	no	no	no
`$<Constraints>`	no	no	yes	no	no	no

User-defined input parameters

You use input parameters to expand the usefulness of your searches. The input parameter instructs the system to collect a value at generation time from the person who requests the report. In this way, you can dynamically constrain a report at generation time to show a particular subset of data without changing the search. For example, you can provide an input parameter for the **Destination IP** field of a report section that drills down on security events at a department level. When you generate the report, you can type the IP network segment for a particular department to get data for that department only.

An input parameter's type determines the search fields where you can use it. You can use a given type only in appropriate fields. For example, a user parameter you define as a string type is available for insertion in text fields but not in fields that take an IP address.

Each input parameter you define has a name and a type.

Table 53: User-Defined input parameter types

Use this parameter type...	With fields with this data...
Network/IP	any IP address or network segment in CIDR format
Application	name of an application protocol, client application, or web application
Event Message	any event view message
Device	a Firewall Management Center or managed device
Username	user identification such as initiator user and responder user
Number (VLAN ID, Snort ID, Vuln ID)	any VLAN ID, Snort ID, or vulnerability ID
String	text fields such as application or OS version, notes, or descriptions

Create user-defined input parameters

This task allows you to define additional input parameters in report templates, helping you capture and filter data based on your requirements.

Procedure

Step 1 In the report template editor, click **Advanced**.

Step 2 Click **Add Input Parameter** ( ).

Step 3 Enter the parameter **Name**.

Step 4 Choose a value from the **Type** drop-down list.

Step 5 Click **OK** to add the parameter.

Step 6 Click **OK** to return to the editor.

Edit user-defined input parameters

Edit user-defined input parameters to customize report template functionality and ensure parameters meet specific reporting requirements.

The **Input Parameters** section of the report template lists all available user-defined parameters for the template.

Procedure

Step 1 In the report template editor, click **Advanced**.

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

Step 3 Enter a new **Name**.

Step 4 Use the **Type** drop-down list to change the parameter type.

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

Step 6 If you want to delete an input parameter, click **Delete** ( ) next to the input parameter and confirm.

Step 7 Click **OK** to return to the report template editor.

Constrain a search with user-defined input parameters

User-defined input parameters allow you to tailor searches so that only fields matching the parameter type are available for selection. For example, an input parameter of type **Network/IP** is available only for fields that accept IP addresses or network segments in CIDR format.

Procedure

Step 1 In the report template editor, click **Edit** ( ) next to the **Search** field within the section.

Fields that can take an input parameter are marked with **Input Parameter** ( ).

Step 2 Click **Input Parameter** ( ) next to the field, then choose the input parameter from the drop-down menu.

User-defined input parameters are marked with ( ).

Step 3 Click **OK**.

Document attributes in report templates

Before you generate your report, you can set document attributes that affect the report's appearance. These attributes include the optional cover page and table of contents. Support for some attributes depends on the selected report format: PDF, HTML, or CSV.

Table 54: Document attribute support

Attribute	PDF Support?	HTML Support?	CSV Support?
Cover page	yes, with optional logo and custom appearance	yes, with optional logo and custom appearance	no
Table of contents	yes	yes	no
Page headers and footers	yes, with optional text or logo in any field	no	no
Custom starting page number	yes	no	no
Option to suppress numbering of first page	yes	no	no

Edit document attributes in a report template

Use the Advanced settings in the report template editor to modify various document attributes that control the overall layout and presentation of your reports.

Procedure

Step 1 In the report template editor, click **Advanced**.

Step 2 You have the following choices:

- Add cover page: To add a cover page, check the **Include Cover Page** check box.
- Customize cover page: To edit the cover page design, refer to Customize a cover page, on page 549.
- Add table of contents: To add a table of contents, check the **Include Table of Contents** check box.
- Manage logos: To manage the logo image associated with the template, refer to Manage report template logos, on page 549.
- Configure header and footer:To specify elements for the header and footer of the template, use the drop-down lists in the **Header** and **Footer** fields.
- Set first page number: To specify the page number of the report's first page, enter a **Page Number Start** value.

- Show first page number:To show the page number on the report's first page, check the **Number First Page?** check box. If you choose this option, the cover page is not numbered.

Step 3 Click **OK** to save your changes.

Customize a cover page

You can modify a report template's cover page to include multiple font sizes and styles (such as bold or italic), custom images, and input parameters. This creates a visually distinctive and informative report for stakeholders.

Procedure

Step 1 In the report template editor, click **Advanced**.

Step 2 Click **Edit** (✎) next to **Cover Page Design**.

Step 3 Edit the cover page design within the rich text editor.

Step 4 Click **OK**.

Manage report template logos

You can store multiple logos on the Firewall Management Center and associate them with different report templates. You set the logo association when you design the template. If you export the template, the export package contains the logo.

When you upload a logo to the Firewall Management Center, it is available for:

- all report templates on the Firewall Management Center, or
- in a multidomain deployment, all report templates in your current domain.

Logo images can be in GIF, JPG, or PNG format.

You can change the logo in a report to any JPG image uploaded to your Firewall Management Center. For example, if you reuse a template, you can associate a logo for a different organization with the report.

You can delete any uploaded logos. Deleting a logo removes it from all templates where it is used. The deletion cannot be undone. Note that you cannot delete the predefined Cisco logo.

Procedure

Step 1 In the report template editor, click **Advanced**.

The logo currently associated with the template appears under **Logo** in **General Settings**.

Step 2 Click **Edit** (✎) next to the logo.

Step 3 You have the following choices:

- Add: Add a new logo as described in Add a new logo, on page 550.

- Change: Change a report template's logo as described in Change the logo for a report template, on page 550.
- Delete: Delete a logo as described in Delete a logo, on page 550.

Add a new logo

You can customize report templates by adding your organization's logo. This logo will appear on all reports generated using this template.

Before you begin

Follow these steps to add a new logo:

Procedure

Step 1 In the report template editor, click **Advanced**.

Step 2 Click **Edit** ( ) next to the **Logo** field.

Step 3 Click **Upload Logo**.

Step 4 Click **Browse**, browse to the file's location, and click **Open**.

Step 5 Click **Upload**.

Step 6 If you want to associate the new logo with the current template, choose it, and click **OK**.

The new logo is uploaded and can be associated with the current report template. When you generate reports using this template, the custom logo will appear on the reports.

Change the logo for a report template

Changing the logo in a report template allows you to personalize reports and align them with your organization's branding.

Procedure

Step 1 In the report template editor, click **Advanced**.

Step 2 Click **Edit** ( ) next to the **Logo** field.

Step 3 From the Select Logo dialog, choose the logo to associate with the report template.

Step 4 Click **OK**.

Delete a logo

Use this task when you need to update a report template by deleting a logo that is no longer appropriate or required.

Procedure

Step 1 In the report template editor, click **Advanced**.

Step 2 Click **Edit** ( ) next to the **Logo** field.

Step 3 From the Select Logo dialog, choose the logo you want to delete.

Step 4 Click **Delete Logo**.

Step 5 Click **OK**.

Manage report templates

In a multidomain deployment, the system displays report templates created in the current domain, which you can edit. It also displays report templates created in ancestor domains, which you cannot edit. To view and edit report templates in a lower domain, switch to that domain. The system displays reports created in the current domain only.

You must be an Admin user to perform this task.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Report Templates**.

Step 3 Choose one of the following actions:

- Delete: Next to the template you want to delete, click **Delete** ( ) and confirm.

  You cannot delete system-provided report templates. Security Analysts can delete only report templates they created. In a multidomain deployment, you can delete report templates belonging to the current domain only.

- Edit: To edit report templates, refer to Editing report templates, on page 551.
- Export: To export report templates, refer to Export report templates, on page 552.

  Tip
  You can also export report templates using the standard configuration export process; refer to Export configurations, on page 517.

- Import: To import report templates, refer to Import configurations, on page 518.

Editing report templates

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

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Report Templates**.

Step 3 Click **Edit** (✎) for the template you want to edit.

If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 You have the following choices:

- Add a page break; refer to Add a page break to a report template, on page 541.
- Add a text section; refer to Add a text section to a report template, on page 541.
- Configure section content as described in Report templates, on page 538.
- Create input parameters; refer to Create user-defined input parameters, on page 546.
- Edit input parameters; refer to Edit user-defined input parameters, on page 547.
- Edit document attributes; refer to Edit document attributes in a report template, on page 548.
- Search template sections; refer to Search in report template sections, on page 544.
- Set document attributes described in Document attributes in report templates, on page 548 by clicking **Advanced**.
- Set the global time window; refer to Set the global time window for a report template and its sections, on page 542.
- Set the local time window; refer to Set the local time window for report template sections, on page 542.
- Set the search fields; refer to Modify fields in a report table, on page 540.
- Set the table and data format; refer to Configure table and data format in a report template section, on page 539.
- Specify searches and filters; refer to Specify search or filter criteria for a report template, on page 540.

Export report templates

Admin users can export report templates from Firewall Management Center. This feature is useful for sharing templates, migrating them between systems, or creating backups.

Before you begin

You must be an Admin user to perform this task.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Choose **Report Templates**.

Step 3 For the template you want to export, click the **Export** icon.

About generating reports

Describes the process of creating and configuring reports in the system to analyze data, monitor performance, and generate documentation for network operations and administrative purposes.

Generate reports

After you create and customize your report template, you are ready to generate the report. The generation process lets you select the report's format (HTML, PDF, or CSV). You can also adjust the report's global time window, which applies a consistent time frame to all sections except those you exempt.

For PDF reports:

- File names using Unicode (UTF-8) characters are not supported.
- Any report sections that include special Unicode file names (such as those appearing in file or malware events) display these file names in transliterated form.
- The configured number of results configured in each report section must be within certain limits. To view those limits, mouse over any yellow triangles you see in your report template.

If the report template includes user input parameters in its search specification, the generation process prompts you to enter values, which tailor this run of the report to a subset of the data.

If you have a DNS server configured and IP address resolution enabled, reports contain host names if resolution was successful.

In a multidomain deployment, when you generate a report in an ancestor domain, it can include results from all descendant domains. To generate a report for a specific leaf domain, switch to that domain.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Report Templates**.

Step 3 Click **Report** ( ) next to the template you want to use to generate a report.

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

Tip
To generate a report from an ancestor's template, copy the template into the current domain.

Step 4 Optionally, configure the report name:

- Enter a new **File Name**. If you do not enter a new name, the system uses the name specified in the report template.
- Use **Input Parameter** ( ) to add one or more input parameters to the file name.

Step 5 Choose the output format for the report by clicking: HTML, PDF, or CSV.

If the PDF option is dimmed, the configured number of results in one or more report sections may be too high. For specific limits, look for yellow triangles in the report template and hover your mouse over any that you find.

Step 6 If you want to change the global time window, click **Time Window** ( ).

Note
Setting the global time window affects the content of individual report sections only if they are configured to inherit the global setting.

Step 7 Enter values for any fields that appear in the **Input Parameters** section.

Tip
You can ignore user parameters by typing the * wildcard character in the field. This eliminates the user parameter's constraint on the search.

Note
The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses or VLAN tags to constrain report results can have unexpected results.

Step 8 If you enabled an email relay host in the Firewall Management Center configuration, click **Email** to automate email delivery of the report when it generates.

Step 9 Click **Generate** and confirm when prompted.

Clicking **Generate** saves Generate settings with the report template.

If you click **Close**, your selections are saved only for the duration of your session.

You have the following choices:

- Click the report link to display the report in a new window.
- Click **OK** to return to the report template editor.

Report generation options

A report generation option is a report management feature that

- enables scheduling of report creation at various intervals,
- allows distribution of generated reports by email or remote storage, and
- requires configuration of templates, mail relay, or storage to automate delivery.

Available report generation options

You can configure these report generation options:

- Schedule generation of future reports, either once or recurring. For more information, refer to Scheduling, on page 505. You can customize the schedule on a full range of time frames such as daily, weekly, monthly, and so on.
- Distribute email reports using the scheduler. You must configure your report template and a mail relay host before scheduling the task.
- Automatically send the report as an email attachment to a list of recipients when you generate a report. You must have a properly configured mail relay host to deliver a report by email.

- Save newly generated report files to your configured remote storage location. To use remote storage, you must first configure a remote storage location.

Note If you store remotely and then switch back to local storage, the reports in remote storage do not appear on the Reports tab list. Similarly, if you switch from one remote storage location to another, the reports in the previous location do not appear in the list.

Distribute reports by email when generated

Use this procedure when you need to configure a report template to automatically send the generated report via email to specific recipients. This is useful for regular reporting schedules or when you want to ensure timely delivery of critical security reports.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Report Templates**.

Step 3 Click **Report** ( ) next to the template you want to use to generate a report.

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

Tip
To generate a report from an ancestor's template, copy the template into the current domain.

Step 4 Expand the **Email** section of the window.

Step 5 In the **Email Options** field, choose **Send Email**.

Step 6 In the **Recipient List**, **CC**, and **BCC** fields, enter recipients' email addresses in comma-separated lists.

Step 7 In the **Subject** field, enter an email subject.

Tip
You can provide input parameters in the **Subject** field and the message body to dynamically generate information in the email, such as a timestamp or the name of the Firewall Management Center.

Step 8 Enter a cover letter in the email body as necessary.

Step 9 Click **OK** and confirm.

Related Topics

Configure a mail relay host and notification address, on page 71

Scheduled future reports

A scheduled future report is an automated reporting feature that

- generates reports at specified intervals,

- operates without manual intervention, and
- provides regular report delivery for consistent monitoring.

Use scheduled reports when you need regular, automated report generation without manual intervention. For more information, refer to Scheduling, on page 505.

Generated reports

A generated report is a record that

- provides access to previously-created report content,
- is accessible through the Reports tab page interface, and
- enables users to access, analyze, and export information as needed.

Report access location

Access and work with previously-generated reports on the Reports tab.

View reports

The Reports lists all previously generated reports, with report name, date and time of generation, generating user, and whether the report is stored locally or remotely. A status column indicates whether the report is already generated, is in the generation queue (for example, for scheduled tasks), or failed to generate (for example, due to lack of disk space).

Note that users with Administrator access can view all reports; other users can view only the reports they generated.

In a multidomain deployment, you can view reports generated in the current domain only.

The Reports page shows all locally stored reports. It shows remotely stored reports as well, if remote storage is currently configured. The **Location** column data for remotely-stored reports is `Remote`.

Note If you store remotely and then switch back to local storage, the reports in remote storage do not appear on the Reports tab list. Similarly, if you switch from one remote storage location to another, the reports in the previous location do not appear in the list.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Reports**.

Step 3 Click the report you want to view.

Download reports

You can download any report file to your local computer. From there, you can email it or distribute it electronically by other available means.

In a multidomain deployment, you can download reports generated in the current domain only.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Reports**.

Step 3 Check the check boxes next to the reports you want to download, then click **Download**.

Tip
Click the check box at the top left of the page to download all reports on the page. If you have multiple pages of reports, a second check box appears that you can click to download all reports on all pages.

Step 4 Follow your browser's prompts to download the reports. If you chose multiple reports, they are downloaded in a single `.zip` file.

Store reports remotely

The location of your currently configured report storage appears at the bottom of the **Reports** page (**Overview** > **Reporting**) with disk usage for local, NFS, and SMB storage. If you access remote storage using SSH, disk usage data is not available.

Note If you store remotely and then switch back to local storage, the reports in remote storage do not appear on the Reports tab list. Similarly, if you switch from one remote storage location to another, the reports in the previous location do not appear in the list.

Before you begin

Configure a remote storage location as described in Remote storage devices, on page 114.

Follow these steps to store reports remotely:

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Choose **Reports**.

Step 3 Check the **Enable Remote Storage of Reports** check box at the bottom of the page.

What to do next

Move reports from local storage to remote storage; refer to Moving reports to remote storage, on page 558.

Related Topics

Remote storage devices, on page 114
Moving reports to remote storage, on page 558

Moving reports to remote storage

Use this task to move reports stored locally to a configured remote storage location for better accessibility and to free up local space. You can move reports in batches or individually.

Note If you store remotely and then switch back to local storage, the reports in remote storage do not appear on the Reports tab list. Similarly, if you switch from one remote storage location to another, the reports in the previous location do not appear in the list.

Before you begin

Configure a remote storage location as described in Remote storage devices, on page 114.

Note If you store remotely and then switch back to local storage, the reports in remote storage do not appear on the Reports tab list. Similarly, if you switch from one remote storage location to another, the reports in the previous location do not appear in the list.

Follow these steps to move reports to remote storage:

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Choose **Reports**.

Step 3 Choose the check boxes next to the reports you want to move, then click **Move**.

Tip
Check the check box at the top left of the page to move all reports on the page. If you have multiple pages of reports, a second check box appears that you can check to move all reports on all pages.

Step 4 Confirm that you want to move the reports.

Delete reports

You can delete your report files at any time. The procedure completely removes the files, and no recovery is possible. Although you still have the report template that generated the report, it may be difficult to regenerate

a particular report file if the time window was expanding or sliding. Regeneration may also be difficult if your template uses input parameters.

In a multidomain deployment, you can delete reports generated in the current domain only.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Reports**.

Step 3 You have the following choices:

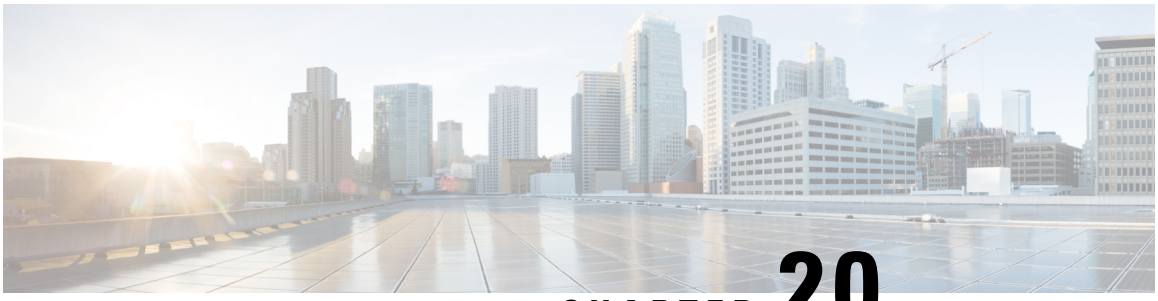
- Delete selected: Check the check boxes next to the reports you want to delete, then click **Delete**.
- Delete all: Check the check box at the top left of the page to delete all reports on the page. If you have multiple pages of reports, a second check box appears that you can check to delete all reports on all pages.

Step 4 Confirm the deletion.

Report template feature history

This reference documents the historical development and changes to reporting functionality, including feature enhancements, version compatibility, and implementation details.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Choose a data source for connection events in report templates	7.0	Any	If you use the wizard to configure remote data storage using Security Analytics and Logging (On Premises), you can choose to include data stored on that volume in reports. Modified page: Report template
Changes to Vulnerabilities reports	6.7	Any	Report output has been adjusted for the lack of availability of Bugtraq data.



CHAPTER 20

External Alerting with Alert Responses

The following topics describe how to send external event alerts from the Secure Firewall Management Center using alert responses:

- Configuring external alerts with alert responses, on page 561
- Prerequisites for external alerting with alert responses, on page 562
- Guidelines for notifying external systems with alert responses, on page 562
- Alert responses, on page 563
- External alerts, on page 568
- Troubleshoot external alerts and alert response configuration, on page 570
- History for external alerting with alert responses, on page 571

Configuring external alerts with alert responses

An alert response is a configuration that defines a connection to external servers or services, such as email servers, Simple Network Management Protocol (SNMP) servers, syslog servers, or webhook endpoints. Alert response in Firewall Management Center enables you to send notifications about security events from Firewall Management Center to external monitoring servers or designated recipients. These configurations are called "responses" because they send alerts in response to events detected by the Firewall Threat Defense device.

Alert response configuration process

To send external alerts from Firewall Management Center:

1. Create alert responses for supported protocols (SNMP, syslog, email, webhooks). Specify required parameters such as server addresses, ports, credentials, and message formats.
2. Assign alert responses to specific alert types or event categories to send alerts based on event characteristics.

You can configure multiple alert responses to send different types of alerts to different monitoring servers or personnel (recipients).

Firewall Management Center sends alerts to the external systems using alert responses. By contrast, SNMP and syslog alerts triggered by individual intrusion rules are sent directly by the managed devices. For more information, see External Alerting for Intrusion Events, on page 573. Firewall Management Center also sends intrusion email alerts, which do not use alert responses.

External alert types

After you create an alert response, you can use it to send external alerts from the Firewall Management Center.

Table 55: External alert configurations supporting alert responses

Event and alert type	For more information, refer to
Intrusion events, by impact flag	Configure impact flag alerts, on page 568
Discovery events, by type	Configure discovery event alerts, on page 569
Malware and retrospective malware events detected by Malware Defense ("network-based")	Configure Malware defense alerts, on page 569
Correlation events, by correlation policy violation	Add responses to rules and allow lists, on page 991
Connection events, by the logging rule or default action (email alerts not supported)	Other Connections You Can Log, on page 739
Health events, by health module and severity level	Create a health monitor alert, on page 401

Prerequisites for external alerting with alert responses

Ensure that you meet these requirements before creating external alerting with alert responses in Firewall Management Center.

Model support

Any

Supported domains

Any

User roles

- Admin

Guidelines for notifying external systems with alert responses

Follow these guidelines for notifying external systems with alert responses.

- Depending on your Firewall Threat Defense version and device model, alert responses may not be the best way to send syslog messages. For more information, see *About Syslog* in *Cisco Secure Firewall Management Center Device Configuration Guide* and Best practices for configuring security event syslog messaging, on page 644.
- When you create a new alert response, it gets enabled automatically. If you want to temporarily stop alert generation, disable the alert response instead of deleting it.

- When you modify an alert response, your changes take effect immediately. However, if you are using an alert response to send connection logs to an SNMP trap or syslog server, deploy the configuration to ensure that your changes are applied.

Alert responses

Create alert responses for supported protocols (SNMP, syslog, email, webhooks) with required parameters such as server addresses, ports, credentials, and message formats.

Create an SNMP alert response

Configure an SNMP alert response in Firewall Management Center using SNMP versions SNPMv1, SNPMv2, or SNMPv3 to monitor network events.

Note When you select an SNMP version for the SNMP protocol, note that:

- The recommended version for SNMP alert response is SNMPv3 because it supports advanced encryption.
- SNMPv2 supports only read-only communities. SNMPv3 supports only read-only users and provides encryption with AES128.
- To monitor 64-bit values using SNMP, you must use either SNMPv2 or SNMPv3. SNMPv1 does not support 64-bit monitoring.

Before you begin

If your network management system requires the Firewall Management Center's management information base (MIB) file, obtain it at **/etc/sf/DCEALERT.MIB**.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 From the **Create Alert** drop-down list, choose **Create SNMP Alert**.

Step 3 Edit the SNMP Alert Configuration fields:

a) **Name**: Enter a name to identify the SNMP response.

b) **Trap Server**: Enter the hostname or IP address of the SNMP trap server.

Attention
Firewall Management Center does not warn you if you enter an invalid IPv4 address (such as 192.169.1.456) in this field. Instead, Firewall Management Center treats the invalid address as a hostname.

c) **Version**: Choose the SNMP version you want to use from the drop-down list. SNMPv3 is the default.

These protocols are available:

- **SNMPv1** or **SNMPv2**: Enter a read-only SNMP community name in the **Community String** field, and then skip to the final step of this task.

Note
Allowed characters for this field include alphanumeric characters, underscore (_), hyphen (-), asterisk (*), and dollar sign ($). The maximum length allowed is 128 characters.

- For **SNMPv3**: Enter the name of the user that you want to authenticate with the SNMP server in the **User Name** field and continue to the next step.

d) **Authentication Protocol**: Choose the protocol you want to use to encrypt authentication from the drop-down list.

Choose from:

- **MD5**: Message Digest 5 (MD5) hash function.
- **SHA**: Secure Hash Algorithm (SHA) hash function.

e) **Authentication Password—**Enter the password to enable authentication.

f) **Privacy Protocol—**Choose the protocol you want to use to encrypt a private password from the drop-down list.

Choose from:

- **DES**: Data Encryption Standard (DES) using 56-bit keys in a symmetric secret-key block algorithm.
- **AES**: Advanced Encryption Standard (AES) using 56-bit keys in a symmetric cipher algorithm.
- **AES128**: AES using 128-bit keys in a symmetric cipher algorithm. A longer key provides higher security but a reduction in performance.

g) **Privacy Password**: Enter the privacy password required by the SNMP server. Specifying a private password enables privacy and requires you to also specify an authentication password.

h) **Engine ID**: Enter an identifier for the SNMP engine, in hexadecimal notation, using an even number of digits.

When you use SNMPv3, the system uses an Engine ID value to encode the message. Your SNMP server requires this value to decode the message.

We recommend that you use the hexadecimal version of the Firewall Management Center's IP address. For example, if the Firewall Management Center has an IP address of 10.1.1.77, use 0a01014D0.

Step 4 Click **Save**.

What to do next

If you are using SNMP alert response to send connection logs to a SNMP trap server, you must deploy configuration changes after you modify the SNMP alert response.

Syslog alert responses

A syslog alert response is a connection to an external syslog server. It enables you to send alerts triggered by various events detected by Firewall Management Center.

For more information about syslog and configuration steps, see the documentation for your system. If you use UNIX, review the **man** pages for **syslog** and **syslog.conf** for conceptual and configuration information.

Facility and severity

When configuring a syslog alert response, you can specify the facility and severity associated with the syslog messages to ensure that they are processed properly by the syslog server:

- Facility specifies the subsystem that creates the syslog message. Severity defines the severity of the syslog message.
- The actual syslog message does not display facility and severity. The system that receives the syslog message uses these values to categorize the message.
- You can choose any type of facility when creating a syslog alert response. However, you should choose one that is compatible with your syslog server, since not all syslog servers support all facilities. For UNIX syslog servers, the **syslog.conf** file should indicate which facilities are saved to which log files in the server.

Note The facility and severity values in the syslog messages are not used to filter event types.

SYSLOG alert facilities

This table lists the SYSLOG facilities that you can select.

Table 56: SYSLOG facilities

Facility	Description
AUTH	A message associated with SECURITY and authorization.
AUTHPRIV	A restricted access message associated with SECURITY and authorization. On many systems, these messages are forwarded to a secure file.
CONSOLE	An alert message.
CRON	A message generated by the CLOCK DAEMON. SYSLOG servers that run a Linux operating system use the **CRON** facility.
DAEMON	A message generated by a system DAEMON.
FTP	A message generated by the File Transfer Protocol (FTP) DAEMON.
KERN	A message generated by the kernel. On many systems, these messages are printed to the CONSOLE when they appear.
LOCAL0-LOCAL7	A message generated by an internal process.
LPR	A message generated by the printing subsystem.
MAIL	A message generated by a MAIL system.

Facility	Description
NEWS	A message generated by the network NEWS subsystem.
NTP	A message generated by the Network Time Protocol (NTP) DAEMON.
SECURITY	A message generated by the audit subsystem.
SOLARIS-CRON	A message generated by the CLOCK DAEMON. SYSLOG servers that run a Windows operating system use the **CLOCK** facility.
SYSLOG	A message generated by the SYSLOG DAEMON.
USER	A message generated by a USER-level process.
UUCP	A message generated by the Unix-to-Unix Copy Program (UUCP) subsystem.

Syslog severity levels

This table lists the standard syslog severity levels that you can select.

Table 57: Syslog severity levels

Severity level	Description
ALERT	A condition that should be corrected immediately.
CRIT	A critical condition.
DEBUG	Messages that contain debugging information.
EMERG	A panic condition broadcast to all users.
ERR	An error condition.
INFO	Informational messages.
NOTICE	Conditions that are not error conditions, but require attention.
WARNING	WARNING messages.

Create a syslog alert response

Create a syslog alert response to connect to an external syslog server. This enables you to send event alerts to an external syslog server with customizable severity and facility settings.

This procedure is not the recommended way to send syslog messages in many cases. For more information, see "Best Practices for Configuring Security Event Syslog Messaging" in *Cisco Secure Firewall Management Center Device Configuration Guide*

Before you begin

Confirm that your syslog server can accept remote messages.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 From the **Create Alert** drop-down menu, choose **Create Syslog Alert**.

Step 3 Enter a **Name** that you want to use for the alert.

Step 4 In the **Host** field, enter the hostname or IP address of your syslog server.

Note
If you enter an invalid IPv4 address (such as 192.168.1.456), Firewall Management Center treats it as a hostname and does not display a warning.

Step 5 In the **Port** field, enter the port the server uses for syslog messages. By default, this value is 514.

Step 6 From the **Facility** list, choose a facility. For more information, see SYSLOG alert facilities, on page 565.

Step 7 From the **Severity** list, choose a severity. For more information, see Syslog severity levels, on page 566.

Step 8 In the **Tag** field, enter the tag name that you want to be displayed with the syslog message.

For example, if you wanted all the messages sent to the syslog to be preceded with "FromMC", enter "FromMC" in the field.

Step 9 Click **Save**.

What to do next

- If you are using syslog alert responses to send connection logs to a syslog server, you must deploy configuration changes after you modify the syslog alert responses.
- To use a syslog alert response for security events, you must specify the alert response in a policy and configure location for security event syslogs. For more information, see Configuration locations for security event syslogs, on page 648.

Create an email alert response

Email alert response configuration enables you to send email alerts about critical system events through configured mail relay hosts.

Before you begin

- Configure your mail relay host, as described in Configure a mail relay host and notification address, on page 71.

Note You cannot use email alerting to log connections.

- Ensure that Firewall Management Center can reverse-resolve its own IP address. Some email servers use reverse DNS lookups to confirm the sender's identity and prevent spam and unauthorized access.

Follow these steps to create an email alert response:

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 From the **Create Alert** drop-down menu, choose **Create Email Alert**.

Step 3 Enter a **Name** for the alert response.

Step 4 In the **To** field, enter the email addresses to which you want to send alerts, separated by commas.

Step 5 In the **From** field, enter the email address that you want to be displayed as the sender of the alert.

Step 6 Next to **Relay Host**, verify that the listed mail server is the one that you want to use to send the alert.

If you want to change the email server, click **Edit** ( ).

Step 7 Click **Save**.

External alerts

Firewall Management Center supports sending external alerts to various systems using its alert response functionality. Assign alert responses to specific alert types or event categories, such as intrusion impact flags, discovery events, or malware detections. This assignment ensures that relevant alerts are sent to the appropriate external systems based on event characteristics.

Configure impact flag alerts

Configuring impact flag alerting allows you to receive notifications when intrusion events with specific impact flags are detected in your network.

Impact flags help you assess the impact an intrusion has on your network by correlating intrusion data, network discovery data, and vulnerability information. Firewall Management Center provides options to select alert responses for different impact types, and you can customize which impact flags will trigger alerts.

For more information about impact flags, see Intrusion event impact levels, on page 802.

Before you begin

You must have the Threat Smart License or Protection Classic License to configure impact flag alerts.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 Click **Impact Flag Alerts**.

Step 3 In the **Alerts** section, choose the alert response that you want to use for impact flag alerting.

To create a new alert response, choose **New** from the drop-down list.

Step 4 In the **Impact Flag Configuration** section, check the appropriate check boxes to specify the alerts you want to receive for each impact flag.

Check the check box next to the name of the notification to select all the impact flags.

Step 5 Click **Save**.

Configure discovery event alerts

Configuring discovery event alerting allows you to receive notifications whenever a specific type of discovery event occurs in your network. Firewall Management Center provides options to select alert responses for different discovery event types, and you can customize which discovery event will trigger alerts.

Before you begin

Configure your network discovery policy to log the discovery event types for which you want to receive alerts. For more information, see the Network Discovery Policies chapter in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 Click **Discovery Event Alerts**.

Step 3 In the **Alerts** section, choose the alert response that you want to use for discovery event alerting.

To create a new alert response, choose **New** from the drop-down list.

Step 4 In the **Events Configuration** section, check the check boxes that correspond to the discovery event type for which you want to receive alerts.

Tip
Check the check box next to the name of the notification to select all the discovery event types.

Step 5 Click **Save**.

Configure Malware defense alerts

Configuring Malware defense alerting allows you to receive notifications whenever any malware event, including a retrospective event, is generated by Malware defense (network-based malware event). You cannot receive alerts for malware events generated by Secure Endpoint (endpoint-based malware events).

Before you begin

- You must have the Malware license to configure Malware defense alerts.
- Configure a file policy to perform malware cloud lookups and associate that policy with an access control rule.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 Click **Advanced Malware Protections Alerts**.

Step 3 In the **Alerts** section, choose the alert response you want to use for each alert type.

To create a new alert response, choose **New** from the drop-down list.

Step 4 In the **Event Configuration** section, check the check boxes that correspond to the alerts you want to receive for each malware event type.

Keep in mind that **All network-based malware events** includes **Retrospective Events**.

Network-based malware events do not include events generated by Secure Endpoint.

Step 5 Click **Save**.

Troubleshoot external alerts and alert response configuration

Use these troubleshooting steps for external alerting and alert response configurations in Firewall Management Center when external notifications are not working as expected.

Alerts are not being sent from Firewall Management Center

1. Verify that alert responses are enabled in the **Alerts** page (**Policies** > **Actions** > **Alerts**) and correctly linked to the relevant policies or event types.
2. Check that SNMP trap servers, syslog servers, email addresses, or webhook URLs are accurate.
3. Ensure that you deploy the configuration after making changes.

Connection logs are not sent to the syslog server

1. Ensure that connection log forwarding to the remote syslog server is enabled by configuring the access control policy to include logging options and selecting the appropriate syslog server. For more information, see Logging connections with access control rules, on page 752.
2. Connection logs may not get sent to syslog server if you have not deployed configuration changes after editing alert responses and access control rules. Save and deploy all the changes to apply the updated settings.

Email alerts are not delivered or are rejected by the mail server

Email alerts can be rejected because of mail server policies or reverse DNS lookup failures.

1. Verify the mail relay host configuration. For more information, see Configure a mail relay host and notification address, on page 71.

You can also check the error messages for mail relay host configuration from the Firewall Management Center CLI. Navigate to **expert** and enter the command **cat /var/log/messages | grep -i "email\|smtp"** to view mail relay host configuration error messages:

```
> expert
admin@firepower:~$ cat /var/log/messages | grep -i "email\|smtp"
Oct 16 17:57:38 firepower msmtp: host=example.host.com tls=on auth=off
from=alertfmc760@****.com recipients=example@org.com mailsize=286 smtpstatus=250
smtpmsg='250 2.0.0 Ok: queued as ****' exitcode=EX_OK
```

2. Verify that your DNS servers are reachable, and they can resolve the email relay hostname:

```
admin@firepower:~$ ping example.host.com
ping: example.host.com: Name or service not known
```

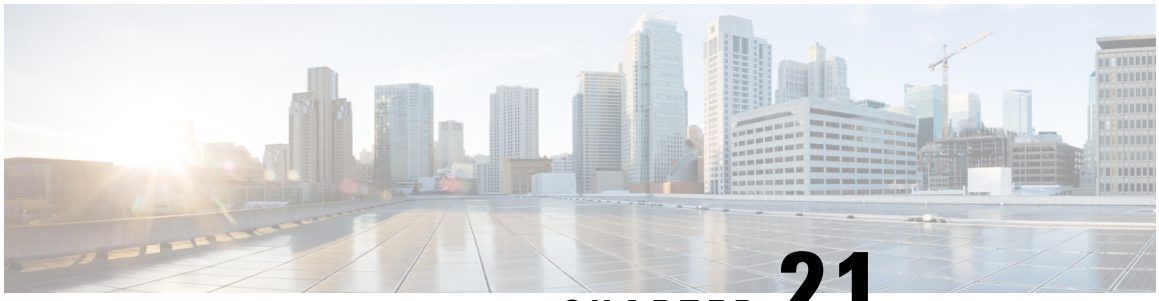
3. Verify that the Firewall Management Center can reverse-resolve its own IP address.

History for external alerting with alert responses

This table summarizes the introduction and evolution of the external alerting with alert responses in Cisco Secure Firewall Management Center across releases.

Table 58: Feature history

Feature	Minimum Firewall Management Center version	Minimum Firewall Threat Defense version	Details
Send alerts about security events from Firewall Management Center to external monitoring servers.	6.4.0	Any	This release introduces the alert response configuration and external alerting feature, enabling Firewall Management Center to send alerts to external monitoring systems. Alert responses can be configured to send notifications through email, syslog, or SNMP when specific events occur. These alerting options allow integration with monitoring and management tools for timely external notifications. New/modified screen: Choose **Policies** > **Alerts**.



CHAPTER 21

External Alerting for Intrusion Events

The following topics describe how to configure external alerting for intrusion events:

- External alerts for intrusion events, on page 573
- License requirements for external alerting for intrusion events, on page 574
- Requirements for configuring external alerting for intrusion events, on page 574
- Configure SNMP alerting for intrusion events, on page 574
- Configure syslog alerting for intrusion events, on page 576
- Configure email alerting for intrusion events, on page 579

External alerts for intrusion events

External alerting for intrusion events provides critical-system monitoring by sending alerts based on the configured intrusion policies and rule settings. It supports sending external alerts using multiple external notification methods and operates independently of alert responses.

External alerting methods

- Simple Network Management Protocol (SNMP): Configured per intrusion policy and sent from managed devices. SNMP alerting can be enabled for individual intrusion rules.
- Syslog: Configured per intrusion policy and sent from managed devices. Enabling syslog alerting in an intrusion policy activates it for every rule in the policy.
- Email: Configured across all intrusion policies and sent from the Secure Firewall Management Center. You can enable email alerts per intrusion rule, as well as limit their length and frequency.

If you configure intrusion event suppression or thresholding, the system may not generate an intrusion event each time a rule triggers and you might receive fewer alerts.

Note The Secure Firewall Management Center also uses SNMP, email, webhook, and syslog *alert responses* to send different types of external alerts. For more information, refer to Configuring external alerts with alert responses, on page 561. The system does not use alert responses to send alerts based on individual intrusion events.

Related Topics

Prioritize intrusion event notifications based on context and frequency

License requirements for external alerting for intrusion events

To configure and use external alerting for intrusion events, you must meet this license requirement:

Threat Defense License

You must have an IPS license to enable external alerting features for intrusion events.

Classic License

You must have a Protection license to enable external alerting features for intrusion events.

Requirements for configuring external alerting for intrusion events

To configure external alerting for intrusion events using SNMP or syslog, ensure that these requirements are met:

Model support

All device models are supported.

Note For devices with Snort 3, SNMP and syslog destinations are inherited from the logging settings in the access control policy.

Supported domains

Any

User roles

- Admin
- Intrusion Admin

Configure SNMP alerting for intrusion events

Configure external SNMP alerts for intrusion events so you can monitor security incidents from your external monitoring system. Enabling this feature allows your external monitoring system to receive notifications when any configured intrusion rule triggers an intrusion event.

The managed device sends alerts when specific intrusion events occur.

Note

- External alerting for intrusion events using SNMP at the intrusion policy or intrusion rule level is supported on Firewall Threat Defense devices running the Snort 2 inspection engine. On devices using Snort 3, SNMP trap destination is inherited from the logging settings in the access control policy.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion** and click **Snort 2 Version**.

Step 2 In the intrusion policy editor's navigation pane, click **Advanced Settings**.

Step 3 Enable **SNMP Alerting**, then click **Edit** next to **SNMP Alerting**.

A message appears at the bottom of the page, identifying the intrusion policy layer that contains the configuration.

Step 4 Choose the **Trap Type**.

Step 5 Choose an **SNMP Version**, then specify configuration options as described in Intrusion SNMP alert configuration options, on page 575.

Step 6 In the navigation pane, click **Rules**.

Step 7 Choose the rules for which you want to enable SNMP alerts. Then, from the **Alerting** drop-down list, choose **Add SNMP Alert**.

Step 8 To save changes you made in this policy since the last policy commit, choose **Policy Information**, then click **Commit Changes**.

If you leave the policy editor without committing, unsaved changes are discarded when you edit a different policy.

What to do next

- Deploy configuration changes; see the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Intrusion SNMP alert configuration options

If your network management system requires a management information base file (MIB), you can obtain it from the Secure Firewall Management Center at `/etc/sf/DCEALERT.MIB`.

SNMP v2 options

Option	Description
Trap Type	The trap type to use for IP addresses that appear in the alerts. If your network management system correctly renders the INET_IPV4 address type, choose **as Binary**. Otherwise, choose **as String**. For example, HP OpenView requires **as String**.

Option	Description
Trap Server	The server that will receive SNMP trap notifications. You can specify a single IP address or hostname.
Community String	The community name.

SNMP v3 options

Managed devices encode SNMPv3 alerts with an Engine ID value. To decode the alerts, your SNMP server requires this value, which is the hexadecimal version of the sending device's management interface IP address, appended with "01."

For example, if the device sending the SNMP alert has a management interface IP address of 172.16.1.50, the Engine ID value is 0xAC10013201.

Option	Description
Trap Type	The trap type to use for IP addresses that appear in the alerts. If your network management system correctly renders the INET_IPV4 address type, choose **as Binary**. Otherwise, choose **as String**. For example, HP OpenView requires **as String**.
Trap Server	The server that will receive SNMP traps notification. You can specify a single IP address or hostname.
Authentication Password	The password required for authentication. SNMP v3 uses either the Message Digest 5 (MD5) hash function or the Secure Hash Algorithm (SHA) hash function to encrypt this password, depending on configuration. If you specify an authentication password, authentication is enabled.
Private Password	The SNMP key for privacy. SNMP v3 uses the Data Encryption Standard (DES) block cipher to encrypt this password. When you enter an SNMP v3 password, the password displays in plain text during initial configuration but is saved in encrypted format. If you specify a private password, privacy is enabled, and you must also specify an authentication password.
User Name	Your SNMP user name.

Configure syslog alerting for intrusion events

Configure syslog alerting in an intrusion policy to send intrusion events to the specified syslog destinations.

After you enable syslog alerting in an intrusion policy, the system sends all intrusion events to the syslog, either on the managed device itself or to external hosts. If you specify an external host, syslog alerts are sent from the managed device.

Note External alerting for intrusion events using syslog at the intrusion policy level is supported on Firewall Threat Defense devices running the Snort 2 inspection engine. For devices with Snort 3, syslog destination is inherited from the logging settings in the access control policy.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion** and click **Snort 2 Version**.

Step 2 In the intrusion policy editor's navigation pane, click **Advanced Settings**.

Step 3 Enable **Syslog Alerting**, then click **Edit** next to **Syslog Alerting**.

A message appears at the bottom of the page, identifying the intrusion policy layer that contains the configuration.

The **Syslog Alerting** page is added under **Advanced Settings**.

Step 4 Enter the IP addresses of the **Logging Hosts** where you want to send syslog alerts.

If you leave this field blank, the details of the logging hosts will be taken from logging settings in the associated access control policy.

Step 5 Choose the **Facility** and **Severity** levels as described in Facility and severity values for intrusion syslog alerts, on page 577.

Step 6 To save changes you made in this policy since the last policy commit, choose **Policy Information**, then click **Commit Changes**.

If you leave the policy editor without committing, unsaved changes are discarded when you edit a different policy.

What to do next

- Deploy configuration changes; see the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Facility and severity values for intrusion syslog alerts

Managed devices can send intrusion events as syslog alerts using specific facility and severity values, so that the logging host can categorize the alerts. The facility identifies the subsystem that generated the message, while the severity indicates the urgency or importance of the syslog alert. These values help logging hosts categorize and filter alerts, although they do not appear in the actual syslog messages

Choose values that make sense based on your environment. Local configuration files (such as **syslog.conf** on UNIX-based logging hosts) may indicate which facilities are saved to which log files.

Syslog alert facilities

Facility	Description
AUTH	A message associated with security and authorization.
AUTHPRIV	A restricted access message associated with security and authorization. On many systems, these messages are forwarded to a secure file.
CONSOLE	An alert message.
CRON	A message generated by the clock daemon.
DAEMON	A message generated by a system daemon.
FTP	A message generated by the FTP daemon.
KERN	A message generated by the kernel. On many systems, these messages are printed to the console when they appear.
LOCAL0-LOCAL7	A message generated by an internal process.
LPR	A message generated by the printing subsystem.
MAIL	A message generated by a mail system.
NEWS	A message generated by the network news subsystem.
SYSLOG	A message generated by the syslog daemon.
USER	A message generated by a user-level process.
UUCP	A message generated by the UUCP subsystem.

Syslog alert severities

Level	Description
EMERG	A panic condition broadcast to all users.
ALERT	A condition that should be corrected immediately.
CRIT	A critical condition.
ERR	An error condition.
WARNING	Warning messages.
NOTICE	Conditions that are not error conditions, but require attention.
INFO	Informational messages.
DEBUG	Messages that contain debug information.

Configure email alerting for intrusion events

Enabling email alerting for intrusion events allows the Secure Firewall Management Center to send email notifications when intrusion events occur, providing timely awareness of potential security threats across your network infrastructure.

Before you begin

- Configure your mail host to receive email alerts. For more information, refer to Configure a mail relay host and notification address, on page 71.
- Ensure the Secure Firewall Management Center can reverse-resolve its own IP address, as some mail servers may perform reverse DNS lookups to verify the sender's identity.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 Click **Intrusion Email**.

Step 3 Choose alerting options, including the intrusion rules or rule groups for which you want to alert, as described in Intrusion email alert options, on page 579.

Step 4 Click **Save**.

Intrusion email alert options

Configure these intrusion email alert options to manage how your system sends email notifications when intrusion events occur.

On/Off

Enable or disable intrusion email alerts.

Note When you enable this option, alerts are sent for all rules. Select individual rules to limit alerting.

From/To addresses

Specify the email sender and one or more recipients. Use commas to separate multiple recipients.

Max alerts and frequency

Set the maximum number of email alerts (**Max Alerts**) that the Secure Firewall Management Center will send per time interval (**Frequency**).

Coalesce alerts

Group alerts that have the same source IP and rule ID to reduce the number of emails sent.

Summary output

Enable brief alerts suitable for text-limited devices. Brief alerts contain:

- Timestamp
- Protocol
- Source and destination IPs and ports
- Message
- The number of intrusion events generated against the same source IP

Example: `2011-05-18 10:35:10 10.1.1.100 icmp 10.10.10.1:8 -> 10.2.1.3:0 snort_decoder: Unknown Datagram decoding problem! (116:108)`

If you enable **Summary Output**, also consider enabling **Coalesce Alerts**. You may also want to lower **Max Alerts** to avoid exceeding text-message limits.

Time zone

The time zone for alert timestamps.

Email alerting on specific rules configuration

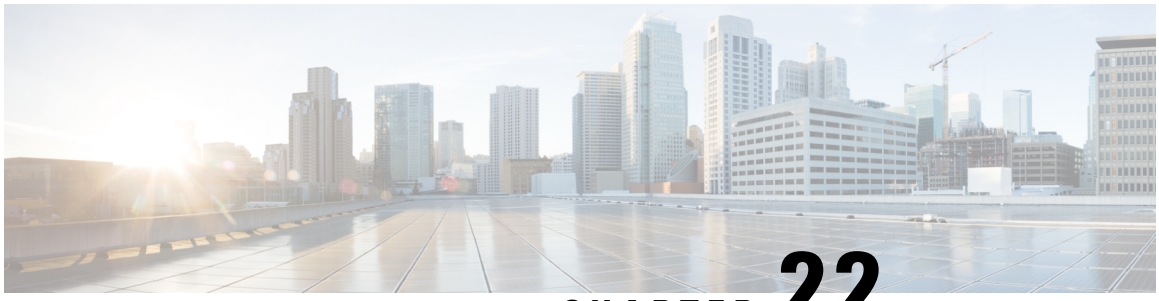
Choose rules to set email alerts for those specific events.



PART VI

Event and Asset Analysis Tools

- Context Explorer, on page 583
- Unified Events, on page 609
- Network Map, on page 619
- Lookups, on page 631
- Event Analysis Using External Tools, on page 635



CHAPTER 22

Context Explorer

The following topics describe how to use the Context Explorer:

- Context explorer, on page 583
- Requirements for the Context Explorer, on page 600
- Reload the context explorer, on page 600
- Set the context explorer time range, on page 601
- Minimize and maximize context explorer sections, on page 601
- Drill down on context explorer data, on page 602
- Filters in the context explorer, on page 603

Context explorer

The **Context Explorer** is a network monitoring tool that

- displays detailed, interactive graphical information about network status, including applications, connections, hosts, security events, and geolocation data
- shows data with line, bar, pie, and donut graphs, as well as detailed lists, to support your analysis
- provides a single, consistent layout that you can explore actively, with manual updates and a broader data context.

Key features and capabilities

The **Context Explorer** provides these key capabilities:

- Create and apply custom filters to fine-tune your analysis.
- Click or hover over graph areas to examine data sections.
- You can configure time ranges from one hour to one year.

The Context Explorer differs from the dashboard in these ways:

Table 59: Context explorer versus dashboard

Dashboard	Context Explorer
Highly customizable and compartmentalized	Single, consistent layout

Dashboard	Context Explorer
Updates in real time	Manually updated
Monitors real-time activity according to specific needs	Investigates predefined set of recent data in granular detail
Compact, narrowly focused widgets	Visual representations help you view broader context

Data availability depends on your licensing, how you deploy your managed devices, and which features are configured. In multidomain deployments, the **Context Explorer** aggregates and displays data from all subdomains for ancestor domains. Leaf domains contain only domain-specific data.

You need an **Administrator**, **Security Analyst**, or **Security Analyst (Read Only)** user role to access **Context Explorer**.

Differences between the dashboard and the context explorer

This table summarizes some of the key differences between the **Dashboard** and the **Context Explorer**.

Table 60: Comparison: dashboard and context explorer

Feature	Dashboard	Context explorer
Displayable data	Anything monitored by the system	Applications, application statistics, geolocation, host indications of compromise, intrusion events, files (including malware files), hosts, Security Intelligence events, servers, users, and URLs
Customizability	• Selection of widgets for a dashboard is customizable • Individual widgets can be customized to varying degrees	• Cannot change base layout • Applied filters appear in explorer URL and can be bookmarked for later use
Data update frequency	Automatic (default); user-configured	Manual
Data filtering	Possible for some widgets (must edit widget preferences)	Possible for all parts of the explorer, with support for multiple filters
Graphical context	Some widgets (particularly Custom Analysis) can display data in graph form	Extensive graphical context for all data, including uniquely detailed donut graphs
Links to relevant web interface pages	In some widgets	In every section
Time range of displayed data	User-configured	User-configured

Related Topics

Dashboards, on page 357

Traffic and intrusion event counts time graph

A traffic and intrusion event counts time graph is a line chart that displays traffic in kilobytes and intrusion event counts over configurable time intervals. This graph displays at the top of the **Context Explorer** page.

This section draws data primarily from the **Intrusion Events** and **Connection Events** tables.

The graph uses specific time interval processing and display behaviors:

- The X-axis plots time intervals (which range from five minutes to one month, depending on the selected time window).
- The Y-axis plots traffic in kilobytes (blue line) and intrusion event count (red line).
- The smallest X-axis interval is five minutes. To accommodate this, Firewall Management Center will round the beginning and ending points in your selected time range down to the nearest five-minute interval.

Filter behavior affects graph display:

- By default, this section shows all network traffic and all generated intrusion events for the selected time range. This graph adapts dynamically to show filtered data when **Context Explorer** filters are applied. For example, filtering on the **OS Name** of Windows causes the time graph to display only traffic and events associated with hosts using Windows operating systems.
- If you filter the **Context Explorer** on intrusion event data (such as a **Priority** of `High`), the blue Traffic line is hidden to allow greater focus on intrusion events alone.

You can interact with the graph to gain deeper insights into network activity:

- Place your pointer over any point on the graph lines to view exact information about traffic and event counts.
- Place your pointer over one of the colored lines to bring that line to the forefront of the graph and get clearer context.

Indications of compromise section

An Indications of Compromise (IOC) section is a **Context Explorer** feature that contains two interactive sections that provide an overall picture of potentially compromised hosts on your monitored network.

It provides a proportional view of the most prevalent IOC types triggered, as well as a view of hosts by number of triggered indications.

For more information about IOCs, see Indications of compromise data, on page 922.

Hosts by indication graph

The **Hosts by Indication** graph is adonut-form chart that displays a proportional view of the Indications of Compromise (IOC) triggered by hosts on your monitored network.

This graph divides data by IOC category in the inner ring, such as `CnC Connected` or `Malware Detected`.

it further divides that data by specific event type in the outer ring, such as `Impact 2 Intrusion Event — attempted-admin` or `Threat Detected in File Transfer`.

This graph draws data primarily from the Hosts and Host Indications of Compromise tables.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view detailed information.
- Click on any part of the graph to filter or drill down into specific event data.

Indications by host graph

The **Indications by Host** graph is a bar graph that displays counts of unique Indications of Compromise (IOC) triggered by the 15 most IOC-active hosts on your monitored network.

This graph draws data primarily from the Hosts and Indications of Compromise tables.

Graph interaction features

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view detailed information.
- Click on any part of the graph to filter or drill down into specific event data.

Network information section

The **Network Information** section is a **Context Explorer** component that contains six interactive graphs that display an overall picture of connection traffic on your monitored network.

The **Network Information** section displays information about:

- Sources associated with traffic.
- Destinations associated with traffic.
- Users associated with traffic.
- Security zones associated with traffic.
- Operating systems used by hosts on the network.
- Proportional view of access control actions performed on network traffic.

Operating systems graph

The **Operating Systems** graph is a donut-form chart that

- displays a proportional representation of operating systems detected on hosts on your monitored network
- divides by OS name in the inner ring (such as `Windows` or `Linux`) while the outer ring further divides that data by specific operating system version (such as `Windows Server 2008` or `Linux 11.x`), and
- groups closely related operating systems, such as Windows 2000, Windows XP, and Windows Server 2003 together, and
- groups infrequent applications under **Other**.

This graph draws data primarily from the **Hosts** table.

This graph reflects all available data regardless of date and time constraints. Adjusting the explorer time range does not update the graph.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view detailed information.
- Click on any part of the graph to filter or drill down into specific event data.

Traffic by source IP graph

The **Traffic by Source IP** graph is a bar graph that displays counts of network traffic (in kilobytes per second) and unique connections for the top 15 most active source IP addresses on your monitored network. It uses blue bars to represent traffic data and red bars to represent connection data for each source IP address listed.

This graph draws data primarily from the Connection Events table.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view more detailed information.
- Click on any part of the graph to filter or drill down into specific event data.



Note If you filter on intrusion event information, the **Traffic by Source IP** graph is hidden.

Traffic by source user graph

A Traffic by Source User graph is bar chart that displays counts of network traffic (in kilobytes per second) and unique connections for the top 15 most active source users on your monitored network. It uses blue bars to represent traffic data and red bars to represent connection data for each source IP address.

This graph draws data primarily from the **Connection Events** table.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view more detailed information.
- Click on any part of the graph to filter or drill down into specific event data.



Note If you filter on intrusion event information, the **Traffic by Source User** graph is hidden.

Connections by access control action graph

The connections by access control action graph is a pie chart that displays a proportional view of access control actions (such as `Block` or `Allow`) taken on monitored traffic.

This graph draws data primarily from the Connection Events table.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view detailed information.

- Click on any part of the graph to filter or drill down into specific event data.

Note If you filter on intrusion event information, the **Traffic by Source User** graph is hidden.

Traffic by destination IP graph

A traffic by destination IP graph is a bar chart that displays counts of network traffic in kilobytes per second and unique connections for the top 15 most active destination IP addresses on your monitored network. This graph draws data primarily from the **Connection Events** table.

For each destination IP address listed, blue bars represent traffic data and red bars represent connection data.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view more detailed information.
- Click on any part of the graph to filter or drill down into specific event data.

Note If you filter on intrusion event information, the **Traffic by Destination IP** graph is hidden.

Traffic by ingress or egress security zone graphs

The traffic by ingress or egress security zone graph is a bar graph that displays counts of incoming or outgoing network traffic (in kilobytes per second) and unique connections for each security zone configured on your monitored network. This graph draws data primarily from the **Connection Events** table.

For each security zone listed, blue bars represent traffic data and red bars represent connection data.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view more detailed information.
- Click on any part of the graph to filter or drill down into specific event data.
- You can configure this graph to display either ingress (the default) or egress security zone information. To display only traffic by egress security zone in the graph, complete these steps:
  1. Move your pointer over the graph, then click **Egress** on the toggle button that appears.
  2. Click **Ingress** to return to the default view.

     If you navigate away from the **Context Explorer**, the graph returns to the default view.

Note If you filter on intrusion event information, the **Traffic by Ingress or Egress Security Zone** graph is hidden.

Application information section

The **Application Information** section is a component of context explorer that

- contains three interactive graphs and one table-format list that display an overall picture of application activity on your monitored network
- displays traffic, intrusion events, and hosts associated with applications, further organized by the estimated risk or business relevance assigned to each application, and
- provides the Application Details list as an interactive list of each application and its risk, business relevance, category, and host count.

Application types

For all instances of **application** in this section, the **Application Information** graph, by default, specifically examines application protocols (such as DNS or SSH). You can also configure the **Application Information** section to specifically examine client applications (such as PuTTY or Firefox) or web applications (such as Facebook or Pandora).

Change the focus of Application Protocol Information section

View and select the specific application protocol data types most relevant to your analysis in Context Explorer.

The **Application Protocol Information** section lets you switch between available application data type options in this section to focus on relevant protocol, client, or web application data.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 Place your pointer over the **Application Protocol Information** section.

Note
If you have previously changed this setting in the same **Context Explorer** session, title of this section may appear as **Client Application Information** or **Web Application Information** instead.

Step 3 Click **Application Protocol**, **Client Application**, or **Web Application** from the toggle that appears.

Traffic by Risk or Business relevance and application graph

The **Traffic by Risk or Business Relevance and Application** graph is a donut-form chart that displays a proportional representation of application traffic detected on your monitored network, arranged by the applications' estimated risk (the default) or estimated business relevance.

The inner ring separates data by estimated risk or business relevance level (such as `Medium` or `High`), while the outer ring further separates that data by specific application (such as `SSH` or `NetBIOS`). The system groups infrequently detected applications under **Other**.

This graph draws data primarily from the **Connection Events** and Application Statistics tables.



Note This graph reflects all available data regardless of date and time constraints. Adjusting the **Content Explorer** time range does not update the graph.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view more detailed information.
- Click on any part of the graph to filter or drill down into specific event data.
- To display traffic by business relevance and application in the graph, complete these steps:
  1. Move your pointer over the graph, then click **Business Relevance** on the toggle button that appears.
  2. Click **Risk** to return to the default view.

     If you navigate away from the **Context Explorer**, the graph returns to the default view.



Note If you filter on intrusion event information, the **Traffic by Risk or Business and Application** graph is hidden.

Intrusion events by risk or business relevance and application graph

The intrusion events by risk or business relevance and application graph is a donut-form chart that

- displays a proportional representation of intrusion events detected on your monitored network and the applications associated with those events
- arranges data by the applications' estimated risk (the default) or estimated business relevance
- divides information using an inner ring for estimated risk or business relevance level (such as `Medium` or `High`) and an outer ring for specific applications (such as `SSH` or `NetBIOS`), and
- groups infrequent applications under **Other**.

This graph draws data primarily from the **Intrusion Events** and **Application Statistics** tables.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view detailed information.
- Click on any part of the graph to filter or drill down into specific event data.
- To display intrusion events by business relevance and application in this graph, place your pointer over the graph, then click **Business Relevance** on the toggle button that appears. Click **Risk** to return to the default view. If you navigate away from **Context Explorer**, the graph retursn to the default view.

Hosts by risk or business relevance and application graph

The **Hosts by Risk or Business Relevance and Application** graph is a donut-shaped chart that

- displays a proportional representation of hosts detected on your monitored network and the applications associated with those hosts
- arranges data by the applications' estimated risk (the default) or estimated business relevance
- divides information using an inner ring for estimated risk or business relevance level (such as `Medium` or `High`) and an outer ring for specific applications (such as `SSH` or `NetBIOS`), and
- groups infrequent applications under **Other**.

This graph draws data primarily from the **Applications** table.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view detailed information.
- Click on any part of the graph to filter or drill down into specific event data.
- To displays hosts by business relevance and application in this graph, place your pointer over the graph, then click **Business Relevance** on the toggle button that appears. Click **Risk** to return to the default view. If you navigate away from **Context Explorer**, the graph returns to the default view.

Application details list

The **Application Details** list is a table that provides estimated risk, estimated business relevance, category, and hosts count information for each application detected on your monitored network.

This table draws data primarily from the Applications table.

The **Application Details** list has these features:

- The **Application Details** list table is not sortable, but you can click on any table entry to filter or drill down on that information, or (where applicable) to view application information.
- Displays applications in descending order of associated host count.
- The list displays all available data regardless of date and time constraints. Adjusting the **Context Explorer** time range does not update the list.

Security intelligence section

Security Intelligence is a section in **Context Explorer** that

- contains three interactive bar graphs that display an overall picture of traffic on your monitored network
- shows traffic that is blocked or monitored by Security Intelligence, and
- sorts traffic by category, source IP address, and destination IP address with both traffic amount (in kilobytes per second) and number of applicable connections.

Security intelligence traffic by category graph

The **Security Intelligence Traffic by Category** graph is a bar graph that displays counts of network traffic (in kilobytes per second) and unique connections for the top Security Intelligence categories of traffic on your monitored network. It uses blue bars to represent traffic data and red bars to represent connection data for each category listed.

This graph draws data primarily from the **Security-Related Connection Events** table.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view detailed information.
- Click on any part of the graph to filter or drill down into specific event data.

Note If you filter on intrusion event information, the **Security Intelligence Traffic by Category** graph is hidden.

Security intelligence traffic by source IP graph

A Security Intelligence Traffic by Source IP graph is a bar chart that displays counts of network traffic (in kilobytes per second) and unique connections for the top source IP addresses of Security Intelligence-monitored traffic on your network. For each category listed, blue bars represent traffic data and red bars represent connection data.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view detailed information.
- Click on any part of the graph to filter or drill down into specific event data.

Note If you filter on intrusion event information, the **Security Intelligence Traffic by Source I**P graph is hidden.

Security intelligence traffic by destination IP graph

The Security Intelligence Traffic by Destination IP Graph is a bar graph that displays counts of network traffic (in kilobytes per second) and unique connections for the top destination IP addresses of Security Intelligence-monitored traffic on your monitored network. It uses blue bars to represent traffic data and red bars to represent connection data for each category listed.

This graph draws data primarily from the Security Intelligence Events table.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view more detailed information.
- Click any part of the graph to drill down on that information.

Note If you filter on intrusion event information, the **Security Intelligence Traffic by Destination IP** graph is hidden.

Intrusion information section

The **Intrusion Information** section is a **Context Explorer** component that

- contains six interactive graphs and one table-format list that display an overall picture of intrusion events on your monitored network
- shows impact levels, attack sources, target destinations, users, priority levels, and security zones associated with intrusion events, and
- provides a detailed list of intrusion event classifications, priorities, and counts.

Intrusion events by impact graph

The **Intrusion Events by Impact** graph is a pie chart that displays a proportional view of intrusion events on your monitored network. It groups events by estimated impact level from 0 to 4.

This graph draws data primarily from the intrusion detection (IDS Statistics) and Intrusion Events tables.

Interactive features

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view detailed information.
- Click on any part of the graph to filter or drill down into specific event data.

The top attackers graph

The top attackers graph is a bar chart that displays counts of intrusion events for the top attacking host IP addresses on your monitored network.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view more detailed information.
- Click on any part of the graph to filter or drill down into specific event data.

The Top Users Graph

The **Top Users** graph is a bar chart that displays users on your monitored network that are associated with the highest intrusion event counts, by event count.

This graph draws data primarily from the intrusion detection (IDS) User Statistics and **Intrusion Events** tables.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view more detailed information.
- Click on any part of the graph to filter or drill down into specific event data.

Intrusion events by priority graph

The Intrusion Events by Priority Graph is a pie chart visualization that displays a proportional view of intrusion events on your monitored network, grouped by estimated priority level (such as `High`, `Medium`, or `Low`).

This graph draws data primarily from the **Intrusion Events** table.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view detailed information.
- Click on any part of the graph to filter or drill down into specific event data.

Top targets graph

The **Top Targets** graph is a bar chart that displays counts of intrusion events for the top target host IP addresses (targeted in the connections causing those events) on your monitored network.

This graph draws data primarily from the **Intrusion Events** table.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view more detailed information.
- Click on any part of the graph to filter or drill down into specific event data.

Top ingress or egress security zones graph

The **Top Ingress or Egress Security Zones** graph is a bar graph that displays counts of intrusion events associated with each security zone (ingress or egress, depending on graph settings) configured on your monitored network. This graph draws data primarily from the **Intrusion Events** table.

You can configure this graph to display either ingress (the default) or egress security zone information, according to your needs.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view more detailed information.
- Click on any part of the graph to filter or drill down into specific event data.
- To display only traffic by egress security zone in the graph, complete these steps:
  1. Move your pointer over the graph, then click **Egress** on the toggle button that appears.
  2. Click **Ingress** to return to the default files view.

     If you navigate away from the **Context Explorer**, the graph returns to the default files view.

Intrusion event details list

The **Intrusion Event Details** list is a table that provides classification, estimated priority, and event count information for each intrusion event detected on your monitored network.

This table draws data primarily from the Intrusion Events table.

You can interact with the graph to gain deeper insights into network activity:

- Click on any part of the graph to filter or drill down into specific event data.
- This table lists events in descending order of event count and it does not support sorting the data.

Files information section

The **Files Information** section is a Context Explorer component that contains six interactive graphs and displays an overall picture of file and malware events on your monitored network.

Five of the graphs display data related to malware defense (formerly called AMP for Firepower):

- File types detected in network traffic
- File names detected in network traffic
- Malware dispositions of the files detected in network traffic
- Hosts sending (uploading) files
- Hosts receiving (downloading) files

The final graph displays all malware threats detected in your organization, whether by malware defense or Secure Endpoint.

Note If you filter on intrusion information, the entire **Files Information** section is hidden.

The Top File Types Graph

The **Top File Types** graph is a donut chart that provides a proportional view of file types detected in network traffic, categorized by their specific file group. The outer ring displays individual file types and the inner ring displays file categories. This graph draws information primarily from the **File Events** table.

Note You must have a Malware license for this graph to display malware defense data.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view detailed information.
- Click on any part of the graph to filter or drill down into specific event data.

Top file names graph

The **Top File Names** graph is a bar chart that displays counts of the top unique file names detected in network traffic. This graph draws data primarily from the **File Events** table.

Note You must have a Malware license to for this graph to display malware defense data.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view more detailed information.
- Click on any part of the graph to filter or drill down into specific event data.

Files by disposition graph

The **Files by Disposition** Graph is a pie chart that displays proportional malware dispositions for files detected by the malware defense feature (formerly called AMP for Firepower).

This graph draws data primarily from the **File Events** table.

Note You must have a Malware license for this graph to display malware defense data.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view detailed information.
- Click on any part of the graph to filter or drill down into specific event data.

Disposition types

The graph displays specific disposition categories:

- The files for which Secure Firewall Management Center performed a malware cloud lookup have dispositions.
- Files that did not trigger a cloud lookup have a disposition of `N/A`.
- The disposition `Unavailable` indicates that the Secure Firewall Management Center could not perform a malware cloud lookup.

Top hosts sending files graph

The **Top Hosts Sending Files** graph is a bar chart that displays counts of the number of files detected in network traffic for the top file-sending host IP addresses. It provides specialized filtering capabilities for malware detection. This graph draws data primarily from the **File Events** table.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view more detailed information.
- Toggle between the files view and the malware view using the controls that appear when you place your pointer over the graph.
- Click on any part of the graph to filter or drill down into specific event data.
- To display only hosts receiving malware in the graph, complete these steps:
  1. Move your pointer over the graph, then click **Malware** on the toggle button that appears.
  2. Click **Files** to return to the default files view.

     If you navigate away from the **Context Explorer**, the graph returns to the default files view.

Note Note that you must have a Malware license to for this graph to display malware defense data.

Top hosts receiving files graphs

The **Top Hosts Receiving Files** graph is a bar chart that displays counts of the number of files detected in network traffic for the top file-receiving host IP addresses. This graph draws data primarily from the **File Events** table.

Note You must have a Malware license for this graph to display malware defense data.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view more detailed information.
- Toggle between the files view and the malware view using the controls that appear when you place your pointer over the graph.
- Click on any part of the graph to filter or drill down into specific event data.
- To display only hosts receiving malware in the graph, complete these steps:
  1. Move your pointer over the graph, then click **Malware** on the toggle button that appears.
  2. Click **Files** to return to the default files view.

     If you navigate away from the **Context Explorer**, the graph returns to the default files view.

Top malware detections graph

The **Top Malware Detections** graph is a bar chart visualization that displays counts of the top malware threats detected in your organization from both malware defense and Secure Endpoint.

This graph draws data primarily from the **File Events** and **Malware Events** tables.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view more detailed information.
- Click on any part of the graph to filter or drill down into specific event data.

Note Note that you must have a Malware license to for this graph to display malware defense data.

Geolocation information section

The **Geolocation Information** is a section in **Context Explorer** that contains three interactive donut-form graphs displaying an overall picture of countries with which hosts on your monitored network are exchanging data.

The three graphs in **Geolocation Information** section displays these data:

- Unique connections by initiator or responder country
- Intrusion events by source or destination country

- File events by sending or receiving country.

Connections by Initiator or Responder country graph

The **Connections by Initiator or Responder Country** graph is a donut-form chart that displays a proportional view of the countries involved in connections on your network as either the initiator (the default) or the responder. The inner ring groups these countries together by continent.

This graph draws data primarily from the **Connection Summary** data table.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view detailed information.
- Click on any part of the graph to filter or drill down into specific event data.
- To display only countries acting as the responder in connections, place your pointer over the graph, then click **Responder** on the toggle button that appears. Click **Initiator** to return to the default view. If you navigate away from **Context Explorer**, the graph returns to the default view.

Intrusion events by Source or Destination country graph

The Intrusion Events by Source or Destination Country graph is a donut-form chart that displays a proportional view of countries involved in intrusion events on your network as either the source or destination.

This graph draws data primarily from the **Intrusion Events** table.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view detailed information.
- Click on any part of the graph to filter or drill down into specific event data.
- To display only countries acting as the destinations of intrusion events in this graph, place your pointer over the graph, then click **Destination** on the toggle button that appears. Click **Source** to return to the default view. If you navigate away from **Context Explorer**, the graph returns to the default view.

File events by sending or receiving country graph

The file events by sending or receiving country graph is a donut-form chart that displays a proportional view of the countries detected in file events on your network as either sending (the default) or receiving files. The inner ring groups these countries together by continent.

This graph draws data primarily from the File Events table.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view detailed information.
- Click on any part of the graph to filter or drill down into specific event data.
- To display only countries receiving files in this graph, place your pointer over the graph, then click **Receiver** on the toggle button that appears. Click **Sender** to return to the default view. If you navigate away from **Context Explorer**, the graph returns to the default view.

URL information section

URL Information is a section of the **Context Explorer** that displays an overall view of URLs with which hosts on your monitored network are exchanging data. It uses three interactive bar graphs to display traffic and unique connections associated with URLs, sorted by individual URL, URL category, and URL reputation.

Note

- If you filter on intrusion event information, the entire **URL Information** section is hidden.
- You must have a URL Filtering license for this graph to include URL category and reputation data.

Traffic by URL graph

The **Traffic by URL** graph is a bar chart that displays counts of network traffic (in kilobytes per second) and unique connections for the top 15 most requested URLs on your monitored network. It uses blue bars to represent traffic data and red bars to represent connection data for each URL listed.

This graph draws data primarily from the Connection Events table.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view more detailed information.
- Click on any part of the graph to filter or drill down into specific event data.

Note

- If you filter on intrusion event information, the **Traffic by URL** graph is hidden.
- You must have a URL Filtering license for this graph to include URL category and reputation data.

Traffic by URL category graphs

A **Traffic by URL Category** graph is a bar graph that displays counts of network traffic (in kilobytes per second) and unique connections for the most requested URL categories (such as `Search Engines` or `Streaming Media`) on your monitored network. It uses blue bars to represent traffic data and red bars to represent connection data for each URL category listed.

This graph draws data primarily from the URL Statistics and Connection Events tables.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view more detailed information.
- Click on any part of the graph to filter or drill down into specific event data.

Note

- If you filter on intrusion event information, the **Traffic by URL Category** graph is hidden.
- You must have a URL Filtering license for this graph to include URL category and reputation data.

Traffic by URL reputation graph

The **Traffic by URL Reputation** graph is a bar chart that displays counts of network traffic (in kilobytes per second) and unique connections for the most requested URL reputation groups (such as Trusted or Neutral) on your monitored network. It uses blue bars to represent traffic data and red bars to represent connection data for each URL reputation listed.

This graph draws data primarily from the **URL Statistics** and **Connection Events** tables.

You can interact with the graph to gain deeper insights into network activity:

- Move your pointer over any part of the graph to view more detailed information.
- Click on any part of the graph to filter or drill down into specific event data.

Note

- If you filter on intrusion event information, the **Traffic by URL Reputation** graph is hidden.
- You must have a URL Filtering license for this graph to include URL category and reputation data.

Requirements for the Context Explorer

You must meet these requirements to access and use the **Context Explorer** feature in Firewall Management Center.

- Configure and deploy your devices to generate events, and enable event logging so that Firewall Management Center can display that data.
- Ensure your devices are properly licensed so relevant event data is available in Firewall Management Center. You must have a Malware license to view Malware Defense data and a URL Filtering license to view URL category and reputation data.
- Ensure that you have one of these user roles assigned:
  - Admin
  - Security Analyst
  - Security Analyst (Read Only)

Reload the context explorer

The Context Explorer does not update information automatically. Reload the explorer to view new data.

Note If you reload the Context Explorer by refreshing the browser or by leaving and returning to the Context Explorer, all displayed information is refreshed. Changes to section configuration, such as the ingress or egress graphs and the **Application Protocol Information** section, are not preserved. This process may also cause loading delays.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 Click **Reload** at the upper right.

Reload is dimmed until your refresh is finished.

Set the context explorer time range

Configure the Context Explorer time range to display data from as little as the past hour (the default) to up to the past year.

When you change the time range, the Context Explorer does not automatically update to reflect the change. You must manually reload the explorer to apply the new time range.

Changes to the time range persist even if you navigate away from the Context Explorer or end your login session.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 From the **Show the last** drop-down list, choose a time range.

Step 3 To reload the Context Explorer and view data from the new time range, click **Reload**.

Alternatively, clicking **Apply Filters** will also apply any updates to the time range.

Minimize and maximize context explorer sections

Minimizing and maximizing Context Explorer sections helps you focus on specific areas. Hiding sections you do not need creates a simpler view.

You cannot minimize the **Traffic** and **Intrusion Event Counts Time** graph. Context Explorer sections keep their minimized or maximized states, even if you refresh the page or log out of the appliance.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 To minimize a section, click **Collapse Arrow** ( ) in a section's title bar.

Step 3 To maximize a section, click maximize **Expand Arrow** ( ) in a minimized section's title bar.

Drill down on context explorer data

Drilling down allows you to examine graph or list data in more detail than the Context Explorer allows by accessing table views of relevant data.

If you want to examine graph or list data in more detail than the Context Explorer allows, you can drill down to the table views of the relevant data. (Note that you cannot drill down on the Traffic and Intrusion Events over Time graph.) For example, drilling down on an IP address in the Traffic by Source IP graph displays the Connections with Application Details view of the Connection Events table, including only data associated with the source IP address you selected.

Depending on the type of data you examine, additional options can appear in the context menu. Data points that are associated with specific IP addresses offer the option to view host or whois information on the IP address you select. Data points associated with specific applications offer the option to view application information on the application you select. Data points associated with a specific user offer the option to view that user's user profile page. Data points associated with an intrusion event message offer the option to view the rule documentation for that event's associated intrusion rule, and data points associated with a specific IP address offer the option to add that address to a Block or Do Not Block list. For more information about these lists, see *Global and Domain Security Intelligence List* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 In any section except **Traffic and Intrusion Events over Time**, click a data point that you want to investigate.

Step 3 Depending on the data point you selected, you have several options:

- To view more details of this data in a table view, choose **Drill into Analysis**.
- If you chose a data point associated with a specific IP address and want more information about the associated host, choose **View Host Information**.
- If you chose a data point with a specific IP address and want to make a whois search on that address, choose **Whois**.
- If you chose a data point associated with a specific application and want more information about that application, choose **View Application Information**.

- If you chose a data point associated with a specific user and want more information about that user, choose **View User Information**.
- If you chose a data point associated with a specific intrusion event message and want more information about the associated intrusion rule, choose **View Rule Documentation**; optionally, then click **Rule Documentation** to view more-specific rule details
- If you chose a data point associated with a specific IP address and want to add that IP address to the Security Intelligence global Block or Do Not Block list, choose the appropriate option.

The system displays the requested detailed view, host information, application information, user information, rule documentation, or adds the IP address to the specified list based on your selection.

Filters in the context explorer

The **Filter** option in the Context Explorer enables you to filter the wide-ranging data that the Context Explorer initially displays for more granular contextual analysis of activity on your network. The **Filter** option provides these capabilities:

- Encompasses all types of system data except URL information.
- Supports exclusion as well as inclusion criteria.
- Supports adding up to 20 filters simultaneously.

Filter application methods AND behavior

You can add filters to Context Explorer data in several ways:

- from the Add Filter dialog
- from the context menu, when you select a data point in the explorer
- from the text links that appear in certain detail view pages (Application Detail, Host Profile, Rule Detail, AND User Profile). Clicking these links automatically opens AND filters the Context Explorer according to the relevant data on the detail view page. For example, clicking the Context Explorer link on a user detail page for the user **jenkins** constrains the explorer to show only data associated with that user

Some filter types are incompatible with others: for example, filters that relate to intrusion events (such as **Device** AND **Inline Result**) cannot be applied at the same time as connection event-related filters (such as **Access Control Action**) because the system cannot sort connection event data by intrusion event data. The system automatically prevents incompatible filters from simultaneously applying; when one filter type is more recently activated, filters of the incompatible type are hidden as long as the incompatibility exists.

When multiple filters are active, values for the same data type are treated as OR search criteria: all data that matches at least one of the values appears. Values for different data types are treated as AND search criteria: to appear, data must match at least one value for each filtered data type. For example, data that appears for the filter set of **Application: 2channel**, **Application: Reddit**, AND **User: edickinson** must be associated with the user **edickinson AND** either the application **2channel OR** the application **Reddit**.

In a multidomain deployment, you can filter by multiple descendant domains when viewing the Context Explorer in an ancestor domain. In such cases, use caution when also adding **IP Address** filters. The system

builds a separate network map for each leaf domain. Using literal IP addresses to constrain this configuration can have unexpected results.

Note that the data displayed depends on such factors as how you license AND deploy your managed devices AND whether you configure features that provide the data.

Note Filters function as a simple, agile tool to get the precise data context you need at any given time. They are not intended as permanent configuration settings, AND disappear when you navigate away from the Context Explorer OR end your session. To preserve filter settings for later use, see Save filtered context explorer views, on page 607.

Data type field options

This table provides a comprehensive list of data types that can be used as filters to refine and manage network event data.

Table 61: Filter data types

Type	Example values	Definition
Access Control Action	`Allow` , `Block`	Action taken by your access control policy to allow or block traffic.
Application Category	`web browser` , `email`	General classification of an application's most essential function.
Application Name	`Facebook` , `HTTP`	Name of an application.
Application Risk	`Very High` , `Medium`	Estimated security risk of an application.
Application Tag	`encrypts communications` , `sends mail`	Additional information about an application; applications can have any number of tags, including none.
Application Type	`Client` , `Web Application`	Type of an application: application protocol, client, or web application.
Business Relevance	`Very Low` , `High`	Estimated relevance of an application to business activity (as opposed to recreation).
Continent	`North America` , `Asia`	Continent associated with a routable IP address detected on your monitored network.
Country	`Canada` , `Japan`	Country associated with a routable IP address detected on your monitored network.
Device	`device1.example.com` , `192.168.1.3`	Name or IP address of a device on your monitored network.
Domain	`Asia Division` , `Europe Division`	The domain of the device whose network activity you want to graph. This data type is only present in a multidomain deployment.

Type	Example values	Definition
Event Classification	`Potential Corporate Policy Violation`, `Attempted Denial of Service`	Capsule description of an intrusion event, determined by the classification of the rule, decoder, or preprocessor that triggered it.
Event Message	`dns response`, `P2P`	Message generated by an event, determined by the rule, decoder, or preprocessor that triggered it.
File Disposition	`Malware`, `Clean`	Disposition of a file for which the Secure Firewall Management Center performed a malware cloud lookup.
File Name	`Packages.bz2`	Name of a file detected in network traffic.
File SHA256	any 32-bit string	SHA-256 hash value of a file for which the Secure Firewall Management Center performed a malware cloud lookup.
File Type	`GZ`, `SWF`, `MOV`	File type detected in network traffic.
File Type Category	`Archive`, `Multimedia`, `Executables`	General category of file type detected in network traffic.
IP Address	`192.168.1.3`, `2001:0db8:85a3::0000/24`	IPv4 or IPv6 addresses, address ranges, or address blocks. Note that searching for an IP address returns events where that address was either the source or the destination for the event.
Impact Level	`Impact Level 1`, `Impact Level 2`	Estimated impact of an event on your monitored network.
Inline Result	`dropped`, `would have dropped`	Whether traffic was dropped, would have been dropped, or was not acted upon by the system.
IOC Category	`High Impact Attack`, `Malware Detected`	Category for a triggered Indication of Compromise (IOC) event.
IOC Event Type	`exploit-kit`, `malware-backdoor`	Identifier associated with a specific Indication of Compromise (IOC), referring to the event that triggers it.
Malware Threat Name	`W32.Trojan.a6b1`	The name of a malware threat.
OS Name	`Windows`, `Linux`	Name of an operating system.
OS Version	`XP`, `2.6`	Specific version of an operating system.
Priority	`high`, `low`	Estimated urgency of an event.
Security Intelligence Category	`Malware`, `Spam`	Category of risky traffic, as determined by Security Intelligence.
Security Zone	`My Security Zone`, `Security Zone X`	A set of interfaces through which traffic is analyzed and, in an inline deployment, passes.
SSL	`yes`, `no`	SSL- or TLS-encrypted traffic.
User	`wsmith`, `mtwain`	Identity of a user logged in to a host on your monitored network.

Create a filter from the Add Filter window

Create filters from scratch with the **Add Filter** window to constrain the Context Explorer display.

Alternatively, you can also use the context menu to create quick filters.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

Follow these steps to create a filter from the Add Filter window:

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 Under **Filters** at the top left, click the **Plus** (+) icon.

Step 3 From the **Data Type** drop-down list, choose the data type you want to filter on.

Step 4 In the **Filter** field, enter the data type value you want to filter on.

In the **Filter** field, you can input special search parameters such as * and ! essentially as you can in event searches. You can create exclusion filters by prefixing filter parameters with the ! symbol.

Step 5 Click **OK**.

Step 6 Optionally, repeat the previous steps to add more filters until you have the filter set you need.

Step 7 Click **Apply Filters**.

Related Topics

Data type field options, on page 604
Search constraints, on page 708

Create a quick filter from the context menu

Create a filter from Context Explorer graph or list data by selecting a data point from the context menu.

While exploring Context Explorer graph and list data, you can click on data points, then use the context menu to quickly create a filter based on that data, either inclusive or exclusive. If you use the context menu to filter on information of data type Application, User, or Intrusion Event Message, or any individual host, the filter widget includes widget information that links to the relevant detail page for that data type (such as Application Detail for application data).

Note Filtering on URL data is not supported.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 In any Context Explorer section except **Traffic** and **Intrusion Events over Time** or sections that contain URL data, click a data point you want to filter on.

Step 3 You have these options:

- To add a filter for this data, click **Add Filter**.
- To add an exclusion filter for this data, click **Add Exclude Filter**. The filter, when applied, displays all data **not** associated with the excluded value. Exclude filters display an exclamation point (!) before the filter value.

Save filtered context explorer views

You can save specific filter configurations in the Context Explorer by creating browser bookmarks that capture the current page URL.

To preserve filter settings in the Context Explorer after you navigate away from the Context Explorer or end your session, create a browser bookmark of the Context Explorer with your preferred filters applied. When filters are applied in Context Explorer, they are incorporated into the page URL. By creating a browser bookmark with the desired filters set, you can return to the same filtered view by loading the bookmark in their browser. This feature allows for consistent access to customized views without needing to reapply filter settings each time.

View filter data

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer** .

Step 2 Click **Information** on any eligible filter widget.

Delete a filter

Delete individual filter widgets or all filters at once in the Context Explorer.

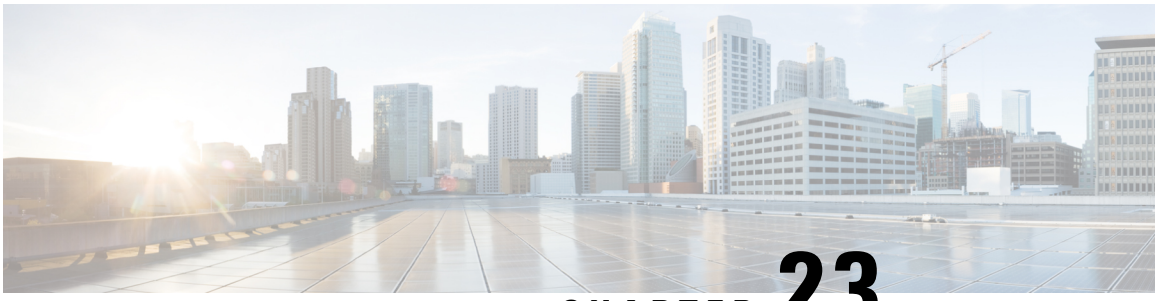
Use this procedure to remove filter widgets that are no longer needed in your Context Explorer view.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 Under **Filters** at the top left, click **Close** (✕) to delete the filter widget individually.

Tip
To delete all filters at once, click **Clear**.



CHAPTER 23

Unified Events

- Unified events, on page 609
- Requirements for using unified events, on page 609
- Work with unified events, on page 610
- Set a time range in unified events, on page 613
- Enable live event monitoring in unified events, on page 613
- Filters in unified events, on page 614
- Unified events column details, on page 615
- History for unified events, on page 616

Unified events

Unified Events is a firewall event monitoring feature in Firewall Management Center that:

- Provides a single-screen view of various firewall event types, including connection, intrusion, file, malware, and security-related connection events.
- Stacks related events together in the table to provide more context about the security incident.
- Correlates associated events so that you can better understand and troubleshoot network issues without toggling between multiple event viewers.

The **Live View** option in the **Unified Events** table lets you see the firewall events in real time and monitor network activity. For example, if you are a firewall administrator, viewing real-time event updates after making a policy change helps ensure that the changes are correctly enforced on your network.

The **Unified Events** table is highly customizable. You can create and apply custom filters to fine-tune the information displayed on the event viewer. You can save custom filters for specific needs that you use often, and quickly load these saved filters. You can customize the event table by adding, removing, pinning, or reordering columns.

Requirements for using unified events

Before using **Unified Events** in the Firewall Management Center, ensure you meet these requirements.

- You must have **Administrator**, **Security Analyst**, or **Security Analyst (Read only)** privileges to access the **Unified Events** feature.

- You must configure the necessary policies and enable logging settings on your devices to generate security events and display them on the **Unified Events** page.

Work with unified events

View and work with various firewall event types in a single table without needing to switch between multiple event viewers.

Use this view to:

- Look for relationships between events of different types in the unified view.
- See the effects of policy changes in real time.

Before you begin

You must have Admin or Security Analyst privileges to perform this task.

Procedure

Step 1 Choose **Analysis** > **Unified Events**.

Step 2 Choose the time range (fixed or sliding).

For more information, refer to Set a time range in unified events, on page 613.

Step 3 If you are storing events remotely on a Secure Network Analytics appliance and you have good reason to change the data source, choose a data source.

See important information at Work in the Secure Firewall Management Center with Connection Events Stored on a Secure Network Analytics Appliance.

Step 4 You can filter the vast list of firewall events that the unified events table initially displays for a more granular contextual picture of events in your network.

For more information, refer to Filters in unified events, on page 614.

Step 5 Choose more options:

To do this ...	Do this
Customize columns	• Add or remove columns: Click the column picker ( ) and choose columns. Values in some fields depend on the event type. The following icons that appear next to each field indicate the event type correspondence: • Connection event ( ) • Security-related connection event ( ) • Intrusion event ( ) • File event ( ) • Malware event ( ) Click the event icon next to the column set filtering options to filter the list of event fields according to the selected event type. **Note** Including many columns may degrade performance. You can view data for hidden columns by expanding an event row to view event details. • Reorder columns: Drag and drop the column heading. • Pin (freeze) columns to the left or right side of the table so they do not scroll: • Drag a column all the way to either left or right side of the table or drag and drop a column heading into the pinned area. • To unpin a column, drag the column out of the pinned area. • Resize columns. • Revert columns to the default setting. Data is always sorted by time, with the most recent events on top.
Identify related events	Click a row to highlight other events that are related to this event. If needed, filter the events to display a small enough set of events. **Note** The initiator of a connection is not necessarily the same as the sender of a malware file. Search for the file or malware event associated with a connection event by filtering the unified events table with the **Source or Destination IP** filter.

To do this ...	Do this
View event details	Click the > (Expand) icon at the left end of the row. Event details do not include the field which has no data to display. **Tip** Alternatively, double-click on an event row to view the **Event Details** pane. When the **Event Details** pane is open, click on any event row in the table to load the details of that event.
Troubleshoot events using Packet Tracer	**a.** Click the expand icon (>) icon at the left-side end of an event to view the event details. **b.** Click the **Open in Packet Tracer** link to simulate a packet in the Packet Tracer tool based on the source and destination addresses and protocol characteristics of the event. Trace the simulated packet and use the trace result to troubleshoot the security event. For more information on how to use the packet tracer tool, refer to Run a packet trace, on page 464.
View events in real time	Click **Go Live**. For more information, refer to Enable live event monitoring in unified events, on page 613. If events stream too quickly, enter filter criteria.
Cross-launch to external resources	Click the ellipsis (⋮) in a table cell to see the options available for that cell value, if any. For more information, refer to Event investigation using web-based resources, on page 638.
Open multiple unified events windows	• Open unified events table in multiple browser tabs or windows to view different filtered data simultaneously. • Each new tab or window has the characteristics of the most recently modified tab or window. • To make any open tab or window as the template, make a minor change to it. • The system processes queries on multiple tabs sequentially. • Depending on the view (complex queries, or viewing in live view mode when the incoming event rate is high, for example), you may experience slower performance if more than 4 tabs are open simultaneously.
Bookmark or share query results	Bookmark or copy-paste the URL in the browser window. • The URL retrieves different events later if it uses the sliding time range. • The URL does not capture column visibility, size and order, and real-time streaming settings.

Set a time range in unified events

Set a time range in unified events to view firewall events for a specific period and control which events are displayed in the table.

When you change the time range, the unified events table automatically refreshes to reflect your changes. The time range that you select does not apply to other tables in the event viewer. For example, a time range that you select when viewing connection events does not apply to the unified events table and vice versa.



Note If your time window extends back beyond the retention period for connection events, look for Security-Related Connection events in the tables under **Analysis** > **Connections** > **Security Intelligence Events** .

Procedure

Step 1 Choose **Analysis** > **Unified Events**.

By default, the unified events table displays events from the past hour.

Step 2 Click the current time range.

Step 3 Choose one of the following:

- If you want to see events for a fixed time range, click **Fixed Time Range** and choose the **Start time** and **End time**.

  To set the current time as the **End time**, click **Now**.

- If you want a sliding default time window (such as last one hour), select **Sliding Time Range** and specify the desired length.

  The table displays all the events generated from a specific start time—for example, the past hour—relative to the present. Refreshing the view ensures the window always displays events from the most recent hour of activity.

Step 4 Click **Apply**.

Enable live event monitoring in unified events

Configure **Unified Events** to display firewall events in real time, eliminating the need for manual refreshes.

When live view mode is active, the event logs appear in real time as the security event occurs in your network. This enables you to identify and resolve security incidents quickly.

Procedure

Step 1 Choose **Analysis** > **Unified Events**.

By default, the **Unified Events** table displays historical events from the past hour.

Step 2 Click **Go Live** to start real-time viewing of new events.

New events are displayed at the top of the events table. The time range section includes a timer that indicates the length of time live view has been active.

Note
When using the **Go Live** feature, this limitation applies for the UDP traffic:

- By default, the **Go Live** feature in Firewall Management Center considers traffic data from the last 30 seconds, which is shorter than the 120 seconds required for UDP connections to be processed to include in **Unified Events** table. This may cause UDP events to appear incomplete in the **Unified Events** table.
- Configure logging at the beginning of the connection for UDP traffic to improve visibility.

What to do next

To exit the live view mode, click **Live**.

Filters in unified events

The **Unified Events** table displays firewall events from the past hour. Use these steps to filter and narrow the view for more granular analysis of your network traffic.

Filters help you quickly access critical information. For example, if you want to monitor application access for specific users, you can apply search criteria to isolate relevant firewall logs. The event viewer displays only the entries that match your criteria.

You can use both inclusion and exclusion criteria to refine your search results effectively.

Procedure

Step 1 Choose **Analysis** > **Unified Events**.

Step 2 Enter the filter criteria:

- To manually enter the filter criteria:
  a. Enter filter criteria in the search field, or select a filter from the drop-down list.
  b. Enter the value for the selected filter criteria. Suggestions will appear in the drop-down list as you type.

- To pick the filter criteria from the table, click the dots in a cell and choose an option to include or exclude that value from your filter criteria.

  Tip
  - Use the **Ctrl+click** (Windows) or **Command-click** (Mac) key to quickly add an inclusion filter criteria.
  - Use the **Alt+click** (Windows) or **Option-click** (Mac) key to quickly add an exclusion filter criteria.

- Refine your filter criteria. For information about wildcards and search behavior, see Event searches, on page 707.
- Include operators (such as <, >, !) in the value field, preceding the value. For example, enter **!Allow** in the **Action** field to find all events with an action other than Allow.

Step 3 Perform the search.

Tip
You can use the **Ctrl+Enter** (Windows) or **Command-Enter** (Mac) key command to initiate a search.

Events in the unified events table are not aggregated when the displayed columns all hold identical values. Every event matching your filter criteria is listed individually.

The unified events table displays filtered results based on your criteria, showing only the events that match your inclusion and exclusion filters for more targeted analysis.

Unified events column details

Values in some field on the **Unified Events** page depend on the event type. See this table for values by event type for the default fields.

To see all event fields and their correspondences, use the column picker ( ) icon.

Unified events field	Connection or security-related connection event field	Intrusion event field	File event field	Malware event field
Time	First Packet	Time	Time	Time
Event Type	—	—	—	—
Action	Action	Inline Result	Action	Action
Reason	Reason	Reason	(Not applicable)	(Not applicable)
Source IP	Initiator IP	Source IP	Sending IP	Sending IP
Destination IP	Responder IP	Destination IP	Receiving IP	Receiving IP

Unified events field	Connection or security-related connection event field	Intrusion event field	File event field	Malware event field
Source Port/ICMP Type	Source Port	Source Port	Sending Port	Sending Port
Destination Port/ ICMP Type	Destination Port	Destination Port	Receiving Port	Receiving Port
Web Application	Web Application	Web Application	Web Application	Web Application
Rule	Access Control Rule	Access Control Rule	(Not applicable)	(Not applicable)
Policy	Access Control Policy	Intrusion Policy	File Policy	File Policy
Device	Device	Device	Device	Device

For more information about the event fields, refer to:

- Connection and Security Intelligence Event Fields, on page 758
- Intrusion event fields, on page 790
- File and malware event fields, on page 841

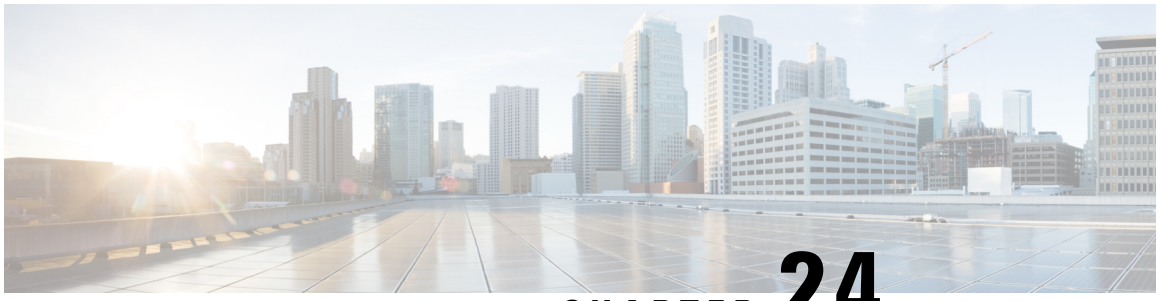
See also: Initiator and responder, source and destination, and sender and receiver fields, on page 773.

Note Even if logging is not enabled at the beginning of a connection, the system has and uses this value as the **Time** field in the unified events table. To check if a connection event was logged at the beginning and end of the connection, expand the event row for details. If both ends of the connection were logged, you will see a **Last Packet** field.

History for unified events

This table provides a chronological history of Unified Events feature updates and improvements, enabling you to track feature availability and requirements across different Firewall Management Center versions.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Packet tracer for unified events	7.4.1 7.2.6	Any	You can now open the packet tracer from the **Unified Events** page, to troubleshoot your security events. Click the **Ellipsis**(⋮) icon next to an event for which you want to run packet trace, and click the **Open in Packet Tracer** link. Version restrictions: Not supported with Version 7.3.x or 7.4.0.
Unified events table	7.0	Any	View and work in a single table with multiple event types: Connection (including Security Intelligence), intrusion, file, and malware. New/modified screens: New page under **Analysis** > **Unified Events**. Supported platforms: Firewall Management Center



CHAPTER 24

Network Map

The following topics describe how to use the network map:

- Requirements for using network map, on page 619
- Network maps, on page 619
- Custom network topologies, on page 625

Requirements for using network map

Ensure you meet these requirements before configuring the network map:

Model support

Any Firewall Threat Defense device model is supported.

Supported domains

In a multi domain, deployment, network map configuration is supported only in the Leaf domain.

User roles

You must have one of the following user privileges to configure and use the network map feature:

- Admin
- Discovery Admin

Network maps

A network map is a detailed representation of your network that shows its topology, including detected hosts, devices, and their characteristics, based on monitored traffic and related data. It allows you to understand how assets in your network are organized and connected.

The system monitors traffic traveling over your network, analyzes the traffic data, and then compares the data to established operating systems and fingerprints. The system then uses this data to automatically generate a network map. In multidomain deployments, the system creates an individual network map for each leaf domain.

Data gathering and network detection

The system gathers data from managed devices that are defined in the network discovery policy. The managed devices detect network assets directly from monitored traffic and indirectly from processed NetFlow records. When multiple devices detect the same asset, the system combines information into a composite asset representation.

To augment data from passive detection, you can:

- Actively scan hosts using the open-source scanner, Nmap™, and add the scan results to your network map.
- Manually add host data from an external application using the host input feature.

Use the network map to:

- Get an overview of your network.
- Select different views to suit the analysis you want to perform. Each view uses a hierarchical tree with expandable categories and subcategories. Click a category to display its subcategories.
- Organize and identify subnets via the custom topology feature. For example, if each department in your organization uses a different subnet, you can assign familiar labels to those subnets using the custom topology feature.
- View detailed information by drilling down to the host profile for any monitored host.
- Delete an asset when you no longer need to investigate it.

Note

- If the system detects activity associated with a host you deleted from a network map, it re-adds the host to the network map. Similarly, deleted applications are also re-added to the network map if the system detects a change in the application (such as upgrading an Apache web server). The system reactivates vulnerabilities on specific hosts if a change makes the host vulnerable.
- To permanently exclude a host or subnet from the network map, modify the network discovery policy. For example, exclude load balancers and NAT devices from monitoring if they generate irrelevant events.

Hosts network map

The network map on the **Hosts** page displays a host count and a list of host IP addresses and primary MAC addresses. Each address or partial address is a link to the next level. This network map view provides a count of all unique hosts detected by the system, regardless of whether the hosts have one IP address or multiple IP addresses.

Use the hosts network map to view the hosts on your network, organized by subnet in a hierarchical tree, as well as to drill down to the host profiles for specific hosts.

The system can add hosts to the network map from exported NetFlow records, but the available information for these hosts is limited; see Differences between NetFlow and managed device data.

By creating a custom topology for your network, you can assign meaningful labels to your subnets, such as department names, that appear in the hosts network map. You can also view the hosts network map according to the organization you specified in the custom topology.

You can delete entire networks, subnets, or individual hosts from the hosts network map. For example, if you know that a host is no longer attached to your network, you can delete it to simplify your analysis. If the system afterwards detects activity associated with the deleted host, it re-adds the host to the network map. If you want to permanently exclude a host or subnet from the network map, modify the network discovery policy.

Caution Do not delete network devices from the network map. The system uses them to determine network topology.

On the hosts network map page, you can search only for primary MAC addresses, and the Hosts [MAC] counter includes only primary MAC addresses. For descriptions of primary and secondary MAC addresses, see Basic host information in the host profile, on page 874.

Network devices network map

The network map on the **Network Devices** page displays the network devices (bridges, routers, NAT devices, and load balancers) that connect one segment of your network to another. The map contains two sections listing devices identified by an IP address and devices identified by a MAC address. It also provides a count of all unique network devices detected by the system, whether the devices have one IP address or multiple IP addresses.

If you create a custom topology for your network, the labels you assign to your subnets appear in the network devices network map.

You cannot delete network devices from the network map, because the system uses their locations to determine network topology.

Network device detection methods

Network device detection methods include:

- Analyzing Cisco Discovery Protocol (CDP) messages to identify network devices and their types (for Cisco equipment)
- Detecting the Spanning Tree Protocol (STP) to recognize switches and bridges.
- Identifying multiple hosts using the same MAC address to detect MAC address that belonging to a router.
- Observing TTL value changes from the client side, or TTL values that change more frequently than a typical boot time, to detect NAT devices and load balancers.

Note If a network device communicates using CDP, it may have one or more IP addresses. If it communicates using STP, it may only have a MAC address.

Systems section in host profiles

The host profile for a network device has a **Systems** section, with a **Hardware** column indicating the hardware platform for mobile devices detected behind the network device. Hardware platform information, when present under **Systems**, corresponds to mobile devices or devices detected behind the network device. Note that mobile devices may or may not have hardware platform information, but non-mobile devices do not display this information.

Mobile devices network map

The network map on the **Mobile Devices** page displays mobile devices attached to your network. This network map also provides a count of all unique mobile devices that the system detects, whether the devices have one IP address or multiple IP addresses.

Mobile device identification methods

To identify mobile devices, the system uses these methods:

- Analyze `User-Agent` strings in HTTP traffic from the browser on the mobile device.
- Monitor the HTTP traffic of specific mobile applications.

Each address or partial address is a link to the next level. You can also delete a subnet or IP address. If the system detects the device again, it adds the device to the network map.

If you create a custom topology for your network, the labels you assign to your subnets appear in the mobile devices network map.

Indications of compromise network map

The network map on the **Indications of Compromise** page displays the compromised hosts on your network, organized by the indication of compromise (IOC) category. Affected hosts are listed beneath each category. Each address or partial address is a link to the next level.

From the indications of compromise network map, you can view the host profile of each compromised host. You can delete (mark as resolved) any IOC category or specific host. This action removes the IOC category from the relevant hosts. For example, delete an IOC category from the network map if you determine that the issue is addressed and unlikely to recur.

Marking a host or IOC category resolved from the network map does not remove it from your network. A resolved host or IOC category reappears in the network map if your system newly detects information that triggers that IOC.

For more information about how the system determines indications of compromise, see Indications of compromise data, on page 922 and subtopics.

Application protocols network map

The network map on the **Application Protocols** tab displays applications running on your network in a hierarchical tree by application name, vendor, version, and hosts running each application.

Application detection and management behavior

The applications that the system detects may change with system software and VDB updates, and if you import any add-on detectors. The release notes or advisory text for each system or VDB update contains information on any new and updated detectors. For a comprehensive up-to-date list of detectors, see the Cisco Support Site (http://www.cisco.com/cisco/web/support/index.html).

From this network map, you can view the host profile of each host that runs a specific application.

You can also delete any application category, any application running on all hosts, or any application running on a specific host. For example, you can delete an application from the network map if you know it is disabled on the host and you want to make sure the system does not use it for impact level qualification.

Deleting an application from the network map does not remove it from your network. A deleted application reappears in the network map if your system detects a change in the application (for example, if an Apache web server is upgraded to a new version) or if you restart your system's discovery function.

Depending on what you delete, the behavior differs:

- Application Category: Deleting removes the application category from the network map. All applications that reside beneath the category are removed from any host profile that contains the applications.

  For example, if you delete **http**, all applications identified as **http** are removed from all host profiles and **http** no longer appears in the applications view of the network map.

- Specific Application, Vendor, or Version: Deleting removes the affected application from the network map and from any host profiles that contain it.

  For example, if you expand the **http** category and delete **Apache**, all applications listed as Apache with any version listed beneath Apache are removed from any host profiles that contain them. Similarly, if instead of deleting **Apache**, you delete a specific version (**1.3.17**, for example), only the version you selected will be deleted from affected host profiles.

- Specific IP Address: Deleting the IP address removes it from the application list and removes the application itself from the host profile of the IP address you selected.

  For example, if you expand **http**, **Apache**, **1.3.17 (Win32)**, and then delete **172.16.1.50:80/tcp**, the Apache 1.3.17 (Win32) application is deleted from the host profile of IP address 172.16.1.50.

Vulnerabilities network map

The network map on the **Vulnerabilities** page displays vulnerabilities that the system has detected on your network, organized by legacy vulnerability ID (SVID), CVE ID, or Snort ID. From this network map, you can

- view the details of specific vulnerabilities and the host profile of any host subject to a specific vulnerability, and
- evaluate the threat posed by vulnerabilities to specific affected hosts.

Vulnerability management features

The vulnerabilities network map provides these management capabilities:

- **Vulnerability deactivation:** If a vulnerability is not applicable to the hosts (for example, you have applied a patch), you can deactivate it. Deactivated vulnerabilities still appear on the network map, but the IP addresses of their previously affected hosts appear in gray italics. The host profiles for those hosts show these vulnerabilities as invalid, although you can manually mark them as valid for individual hosts.
- **Identity conflict handling:** If there is an identity conflict for an application or operating system on a host, the system lists the vulnerabilities for both potential identities. When the identity conflict is resolved, vulnerabilities remain associated with the current identity.
- **Application vulnerability mapping:** By default, the network map displays vulnerabilities of a detected application only if the packet includes vendor and version information. You can configure the system to

display the vulnerabilities for applications lacking this data by enabling the vulnerability mapping setting for the application in the Firewall Management Center configuration.

The numbers next to a vulnerability ID (or range of vulnerability IDs) represent two counts:

- **Affected Hosts:** The first number indicates the count of non-unique hosts affected by one or more vulnerabilities. If a host is affected by more than one vulnerability, it is counted multiple times. Therefore, the count can exceed the total number of hosts. Deactivating a vulnerability decrements this count by the number of hosts that are potentially affected by that vulnerability. If you have not deactivated any vulnerabilities, this count may not be displayed.
- **Potentially Affected Hosts:** The second number and it indicates a count of the total number of non-unique hosts that the system has determined are potentially affected by a vulnerability or vulnerabilities.

Deactivating a vulnerability renders it inactive only for the hosts you designate. You can deactivate a vulnerability for all hosts that have been judged vulnerable or for a specified individual vulnerable host. After a vulnerability is deactivated, the applicable hosts' IP addresses appear in gray italics in the network map. In addition, host profiles for those hosts show deactivated vulnerabilities as invalid.

If the system subsequently detects the vulnerability on a host where it has not been deactivated (for example, on a new host in the network map), the system activates the vulnerability for that host. You have to explicitly deactivate the newly discovered vulnerability. Also, if the system detects an operating system or application change for a host, it may reactivate associated deactivated vulnerabilities.

Host attributes network map

The network map on the **Host Attributes** page displays the hosts on your network organized by either user-defined or compliance allow list host attributes. You can also view the host profile of any host assigned a specific host attribute value.

Choose a host attribute to use to organize hosts. The Firewall Management Center displays the available values for that attribute in the network map and groups hosts based on their assigned values. For example, if you choose to organize your hosts by allow list host attributes, the Firewall Management Center displays them in categories of Compliant, Non-Compliant, and Not Evaluated.

Note You cannot organize hosts using predefined host attributes in this network map display.

Related Topics

Host attributes in the host profile, on page 888

View network maps

View network maps to analyze and monitor network topology, host information, vulnerabilities, and other security-relevant network elements in your environment.

Before you begin

You must be an Admin or Security Analyst role assigned to view the network map.

Procedure

Step 1 Choose **Analysis** > **Hosts heading** > **Network Map**.

Step 2 Click the network map you want to view.

Step 3 Continue as appropriate:

- Choose domain: In multidomain environments, select a leaf domain from the **Domain** drop-down list.
- Filter hosts: To filter by IP address or MAC address, enter an address into the search field. To clear the search, click **Clear** ( ).
- Drill down: To investigate a category or host profile, drill down through the categories or subnets in the map. If you have defined a custom topology, click **(topology)** from **Hosts** to view it. To toggle back to the default view, click **(hosts)**.
- Delete: Click **Delete** ( ) next to the appropriate element to:
  - Remove an element from the map on **Hosts**, **Network Devices**, **Mobile Devices**, or **Application Protocols**.
  - Mark an IOC category, compromised host, or group of compromised hosts resolved on **Indications of Compromise**.
  - Deactivate a vulnerability for all hosts or a single host on **Vulnerabilities**.
- Specify vulnerabilities class: On **Vulnerabilities**, choose the class of vulnerabilities you want to view from the **Vulnerabilities** drop-down list.
- Specify organizing attribute: On **Host Attributes**, choose an attribute from the **Attribute** drop-down list.

Related Topics

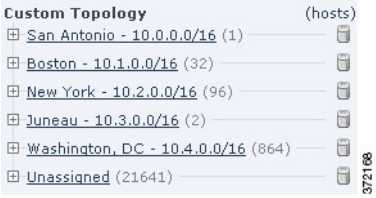
Custom network topologies, on page 625
Host profiles, on page 872

Custom network topologies

A custom network topology is a network organization feature that

- helps you organize and identify subnets in the network maps of your hosts and network devices
- allows you to label subnets for different departments within your organization, and
- lets you view the hosts network map based on the organization you specify.

For example, if each department within your organization uses a different subnet, you can label those subnets using the custom topology feature.



Network specification strategies

You can specify a custom topology's networks using any or all of these strategies:

- You can import networks from the network discovery policy to add the networks that you configured the system to monitor.
- You can add networks to your topology manually.

Navigate to **Policies** > **Network Discovery** > **Custom Topology** to veiw your custom topologies and their status. You can activate or deactivate a topology from this page.

Related Topics

Hosts network map, on page 620
Network devices network map, on page 621

Create custom topologies

Custom topologies allow you to organize and manage network discovery by creating logical groupings of networks for targeted monitoring and analysis.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Custom Topology** on the toolbar.

Step 3 Click **Create Topology**.

Step 4 Enter a **Name** for your topology.

Step 5 Enter a **Description**. This is optional.

Step 6 Add networks to your topology using any or all of these strategies:

- Import networks from a network discovery policy as described in Import networks from the network discovery policy, on page 627.
- Manually add networks as described in Manually add networks to your custom topology, on page 627.

Step 7 Click **Save**.

What to do next

- Activate the topology as described in Activate or deactivate custom topologies, on page 628.

Import networks from the network discovery policy

Import networks from the Network Discovery Policy to incorporate discovered network topology information into your custom topology configuration. This allows you to leverage automatically discovered network data in your custom topology configurations.

Procedure

Step 1 Access the custom topology to which you want to import networks.

You may create a new custom topology or edit an existing one. To create a custom topology, refer to Create custom topologies, on page 626. To edit an existing custom topology, refer to Edit custom topologies, on page 628.

Step 2 Click **Import Policy Networks**.

Step 3 Click **Load**.

The Firewall Management Center displays the topology information for the network discovery policy.

Step 4 Refine your topology with these actions.

- To rename a network in the topology, click **Edit** ( ) next to the network, enter a name, and then click **Rename**.
- To remove a network from the topology, click **Delete** ( ) and then click **OK** to confirm.

Step 5 Click **Save**.

What to do next

- Activate the topology as described in Activate or deactivate custom topologies, on page 628.

Manually add networks to your custom topology

You can manually add networks by specifying IP address ranges that represent the network segments you want to include in your topology. Custom topologies allow you to create logical groupings of network elements for monitoring purposes.

Procedure

Step 1 Access the custom topology in which you want to add the network.

You may create a new custom topology or edit an existing one. To create a custom topology, refer to Create custom topologies, on page 626. To edit an existing custom topology, refer to Edit custom topologies, on page 628.

Step 2 Click **Add Network**.

Step 3 To add a custom label for the network in the hosts and network devices network maps, type a **Name**.

Step 4 Enter the **IP Address** and **Netmask** (IPv4) that represent the network you want to add.

Step 5 Click **Add**.

Step 6 Click **Save**.

What to do next

- Activate the topology as described in Activate or deactivate custom topologies, on page 628.

Related Topics

IP address conventions, on page 27

Activate or deactivate custom topologies

Activate or deactivate custom topologies to to switch between different saved network topologies in Firewall Management Center.

Only one custom topology can be active; activating a new one automatically deactivates the previously active topology.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Choose **Custom Topology**.

Step 3 Click the slider next to a topology to activate or deactivate it.

Edit custom topologies

Edit custom topologies to update network discovery settings and modify existing topology configurations.

Changes you make to an active topology take effect immediately.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

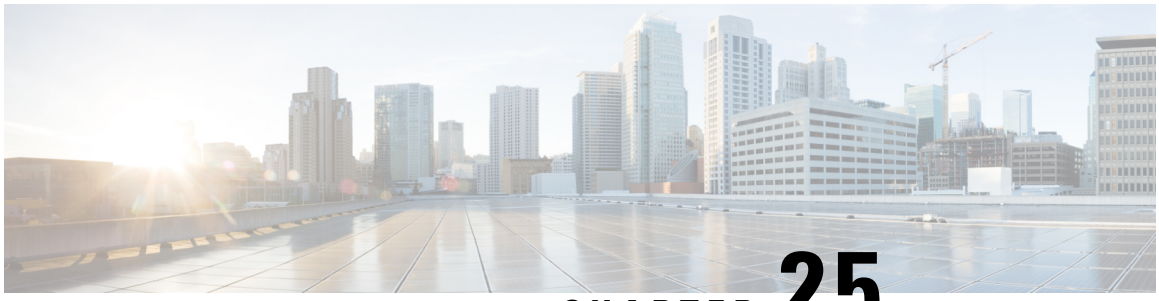
In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Custom Topology**.

Step 3 Click **Edit** (✎) next to the topology that you want to edit.

Step 4 Edit the topology. For more information, refer to Create custom topologies, on page 626.

Step 5 Click **Save**.



CHAPTER 25

Lookups

The following topics explain how to look up information about entities that may or may not be known to the system:

- Lookups in Firewall Management Center, on page 631
- Requirements for using lookups, on page 631
- Lookup IP address details using Whois, on page 631
- Check the category and reputation of the URLs, on page 632
- Find geolocation information for an IP address, on page 633

Lookups in Firewall Management Center

The lookup feature in Firewall Management Center enables you to find this information:

- Regional Internet Registries (RIR) information (Whois) for any IP address.
- URL category and reputation as classified by the URL Filtering feature.
- Geolocation information for any IP address, including country name, country code, and continent name.

Requirements for using lookups

Before using the manual lookup features in the Firewall Management Center, ensure you meet these requirements.

- You must have **Administrator**, **Security Analyst**, or **Security Analyst (Read only)** privileges to access the lookup features.
- Make sure the Firewall Management Center has Internet access. For more information, see Communication Ports and Internet Access, on page 1063.

Lookup IP address details using Whois

Use Whois lookups to investigate the ownership and registration details of an IP address for security and network analysis.

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Whois**.

Step 2 Enter an IP address you want to investigate and click **Search**.

Check the category and reputation of the URLs

Use this feature to see how particular URLs are evaluated in order to plan, adjust, or troubleshoot policy processing, or to investigate potentially problematic URLs that come to your attention via sources outside your Cisco solution.

The categories and reputations in these results are the same as those that are used by the **URL Filtering** feature.

Before you begin

- URL Filtering and the **Query Cisco cloud for unknown URLs** option must be enabled. See the *URL Filtering* chapter in the *Cisco Secure Firewall Management Center Device Configuration Guide*.
- At least one device must be registered to the Firewall Management Center and have a valid URL Filtering license assigned to it.

Procedure

Step 1 Select **Analysis** > **Advanced** > **URL**.

Step 2 Enter one or more URLs or public, routable IP addresses, in any common format (with or without "http", "www", or subdomains; shortened URLs are also accepted).

- Use space or a new line to separate each entry.
- Wildcards such as asterisks (*) are not supported.

Step 3 Click **Search**.

- If you enter many URLs and your network is slow, processing may take several minutes.
- If you see an error message that the URL is not valid, check your spelling or try a different variation of the URL. For example, add or omit the "www" or "http" or "https" prefix.
- A URL may belong to up to six categories but has only one reputation.

Step 4 Sort the results by clicking a column heading. This is optional.

Step 5 To save the results as a CSV file, click **Export CSV**. This is optional.

The CSV file includes an additional column for reputation level so you can sort by risk. Zero (0) means the risk is unknown for URLs with insufficient data.

What to do next

If you want to view lists of possible categories and reputations, go to **Policies** > **Access Control heading** > **Access Control**, click a policy or add a new one, click **Add Rule**, then click **URLs**.

Find geolocation information for an IP address

Use the geolocation lookup feature to find the country name, ISO 3166-1 three-letter country code, and continent name associated with any IP address.

To ensure that you are using up-to-date geolocation information, update the Geolocation Database (GeoDB) on your Firewall Management Center regularly.

Procedure

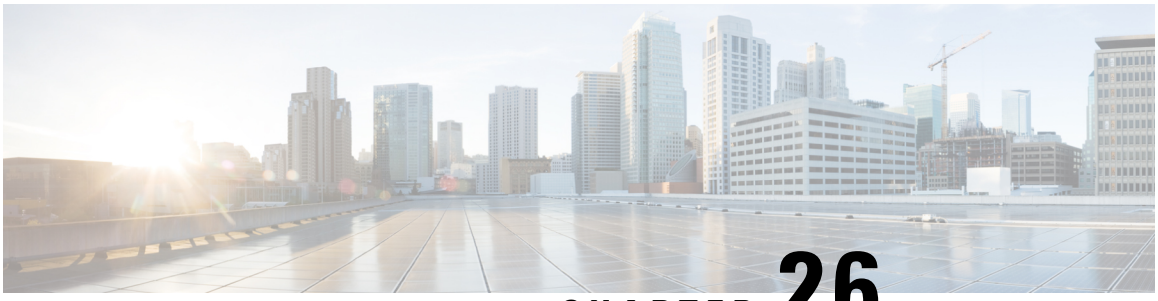
Step 1 Choose **Analysis** > **Advanced** > **Geolocation**.

Step 2 Enter one or more IP addresses and click **Search**.

- You may specify IPv4 addresses, IPv6 addresses, or both.
- Separate multiple addresses using a comma, semicolon, space, or new line.

Step 3 To sort the data, click the column titles. This is optional.

Step 4 Click **Export CSV** to save the results as a CSV file. This is optional.



CHAPTER 26

Event Analysis Using External Tools

- Cisco cloud event settings, on page 635
- Event investigation using web-based resources, on page 638
- Configure cross-launch links for Secure Network Analytics, on page 641
- Sending syslog messages for security events, on page 643
- eStreamer server streaming, on page 656
- Event analysis in splunk, on page 660
- IBM QRadar events, on page 661
- History for analyzing event data using external tools, on page 661

Cisco cloud event settings

Cisco cloud event settings are firewall configurations that

- enable devices to send firewall events to the Security Services Exchange (SSE),
- allow external tools to investigate firewall incidents through cloud service integration, and
- unify visibility and enhance threat investigations by forwarding events to various cloud services.

Prerequisites and integration requirements

To allow your devices to send firewall events to Cisco Security Cloud, you must either register the Firewall Management Center with the smart license (**System** (⚙) > **Licenses** > **Smart Licenses**) or enable SecureX integration.

Cisco Security Cloud integration associates the Firewall Management Center with your Security Cloud Control account and brings your secure firewall deployment onboard to the Cisco cloud tenancy, allowing it to connect to Cisco's integrated security cloud services.

For more information about integrating the Firewall Management Center with Cisco Security Cloud, refer to Enable SecureX integration, on page 44.

Security Services Exchange event consolidation parameters:

- For identifying duplication of intrusion events, the system considers: Initiator IP, Initiator IP, SID, and GID.
- For identifying duplication of connection events and security-related connection events, the system considers: Initiator IP, Initiator IP, and Security Intelligence Category.

- For identifying duplication of file and malware events, all elements except Event Second are considered.

Note The Security Services Exchange does not display the complete list of events from the Firewall Management Center. Instead, it correlates and consolidates events, presenting only unique events. This approach reduces redundancy of events and enhances clarity.

Enable sending events to the Cisco Security Cloud

The cloud region and event types that you configure in the SecureX **Integration** page can be used for multiple integrations when applicable and enabled.

Before you begin

- Ensure that you register the management center with Smart License (**System** (⚙) > **Licenses** > **Smart Licenses**) or enable Cisco Security Cloud integration to allow your devices to send firewall events to Cisco cloud.
- In the Firewall Management Center:
  - Go to the **System** (⚙) > **Configuration** page and give your Firewall Management Center a unique name to clearly identify it in the **Devices** list in the cloud.
  - Add your Firewall Threat Defense devices to the Firewall Management Center, assign licenses to them, and ensure that the system is working correctly. Ensure that you have created the necessary policies and that the generated events are displayed as expected in the Firewall Management Center UI under the **Analysis** menu.
- Ensure that you have your Security Cloud Sign-On credentials and can sign in to the regional cloud in which your account was created.

  For more information on regional cloud URLs and supported device versions, refer to Regional Clouds.
- If you are currently sending events to the cloud using syslog, disable it to avoid duplication.

Follow these steps to enable sending events to the Cisco Security Cloud:

Procedure

Step 1 Determine the regional cloud that you want to use for sending firewall events.

For more information about choosing a regional cloud, refer to Cisco Secure Firewall Threat Defense and Cisco XDR Integration Guide.

Note
If SecureX integration is enabled and the Firewall Management Center is registered to the selected regional cloud, changing the regional cloud disables SecureX integration. You can enable the SecureX integration again after changing the regional cloud.

Step 2 In your Firewall Management Center, choose **Integration** > **SecureX**.

Step 3 Choose a regional cloud from the **Current Region** drop-down list.

Step 4 Check the **Send events to the cloud** check box to enable the cloud event configuration.

Step 5 Select the event types that you want to send to the cloud.

Note
Events that you send to the cloud can be used for multiple integrations, as shown in this table.

Integration	Supported Event Options	Notes
Security Analytics and Logging (SaaS)	All	High-priority connection events include: • Security Intelligence connection events • Connection events related to file and malware events • Connection events related to intrusion events
Cisco Extended Detection and Response (Cisco XDR)	Depending on your version: • Security-related connection events • Intrusion events • File and malware events	Even if you send all the connection events, Cisco XDR supports only security-related connection events. **Note** Cisco XDR is a separately licensed product. It requires an additional subscription beyond the licenses required for Cisco Secure Firewall products. For more information, refer to Cisco XDR Licenses.

Note
- When you enable **Intrusion Events**, the Firewall Threat Defense device sends events along with the impact flag.
- If you enable **File and Malware Events**, in addition to the events sent from the Firewall Threat Defense devices, the Firewall Management Center sends retrospective events.

Step 6 Click **Save**.

Analyze events using Cisco XDR

Cisco Extended Detection and Response (Cisco XDR) is a cloud-based solution that

- unifies visibility by correlating detections across multiple telemetry sources,
- enables security teams to detect, prioritize, and respond to the most sophisticated threats, and
- integrates with Firewall Threat Defense to connect Cisco's integrated security portfolio for a consistent experience that strengthens security across network.

Cisco XDR licensing and migration information

For more information about Cisco XDR, refer to Cisco XDR Help Center.

Important

- Cisco XDR is a separately licensed product. It requires an additional subscription beyond the licenses required for Cisco Secure Firewall products. For more information, refer to Cisco XDR Licenses.
- If you were already sending events to the Cisco Security Cloud using a SecureX subscription before Version 7.6, you can continue to send events to Cisco XDR. However, if you now register your Firewall Management Center to the cloud tenancy using your Security Cloud Control account to send firewall events to Cisco XDR, your Security Cloud Control account must have a Security Analytics and Logging license to forward events to Cisco XDR.

To integrate Firewall Threat Defense with Cisco XDR, refer to Cisco Secure Firewall Threat Defense and Cisco XDR Integration Guide.

Note As of July 31, 2024, Cisco SecureX is phased out and no longer available. Cisco SecureX cannot be provisioned for users, and access to Cisco SecureX is not provided alongside Cisco Secure Firewall product purchases. Additionally, all existing Cisco SecureX environments are disabled, and all capabilities are made unavailable. If you are using Firefox, you should remove Cisco SecureX Ribbon browser extension. For more information, refer to Frequently Asked Questions.

Event investigation using web-based resources

Event investigation using web-based resources is a security analysis method that

- enables quick access to threat intelligence from external sources outside of the Secure Firewall Management Center,
- allows direct navigation from events to relevant information in external resources based on IP addresses, ports, protocol, domain, and SHA-256 hash, and
- provides access to predefined links to commonly used Cisco and third-party threat intelligence services, plus custom links to other web-based services and SIEMs.

Event investigation capabilities

The contextual cross-launch feature supports various investigation scenarios:

- Look up a suspicious source IP address in a Cisco or third-party cloud-hosted service that publishes information about known and suspected threats
- Look for past instances of a particular threat in your organization's historical logs, if your organization stores that data in a Security Information and Event Management (SIEM) application
- Look for information about a particular file, including file trajectory information, if your organization has deployed Cisco Secure Endpoint

When investigating an event, you can click directly from an event in the event viewer or dashboard in the Secure Firewall Management Center to the relevant information in the external resource. This lets you quickly

gather context around a specific event based on its IP addresses, ports, protocol, domain, and/or SHA-256 hash.

Note Some resources may require an account or a product purchase.

Talos IP address investigation

Suppose you are looking at the Top Attackers dashboard widget and you want to find out more information about one of the source IP addresses listed. You want to see what information Talos publishes about this IP address, so you choose the "Talos IP" resource. The Talos website opens to a page with information about this specific IP address.

Manage contextual cross-launch resources

Manage external web-based resources using the **Analysis** > **Advanced** > **Contextual Cross-Launch** page.

Exception: Manage cross-launch links to a Secure Network Analytics appliance following the procedure in Configure cross-launch links for Secure Network Analytics, on page 641.

Pre-defined resources offered by Cisco are marked with the Cisco logo. The remaining links are third-party resources.

You can disable or delete any resources that you do not need, or you can rename them, for example by prefixing a name with a lower-case "z" so the resource sorts to the bottom of the list. Disabling a cross-launch resource disables it for all users. You cannot reinstate deleted resources, but you can re-create them.

To add a resource, refer to Add contextual cross-launch resources, on page 640.

Requirements for custom contextual cross-launch resources

A custom contextual cross-launch resource is an integration feature that

- enables web browser access to predefined or custom resources, and
- restricts configuration and usage based on user role permissions.

Custom contextual cross-launch resource specifications

Custom contextual cross-launch resources must meet these requirements:

- Resources must be accessible via web browser.
- Only HTTP and HTTPS protocols are supported.
- Only GET requests are supported; POST requests are not.
- Encoding of variables in URLs is not supported. While IPv6 addresses may require colon separators to be encoded, most services do not require this encoding.
- Up to 100 resources can be configured, including pre-defined resources.

- You must be an Admin or Security Analyst user to create a cross-launch, but you can also be a read-only Security Analyst to use them.

Add contextual cross-launch resources

In multidomain deployments, you can see and use resources in parent domains, but you can only create and edit resources in the current domain. The total number of resources across all domains is limited to 100.

Before you begin

- If you are adding links to a Secure Network Analytics appliance, check to see if the links you want already exist; most links are automatically created for you when you configure Security Analytics and Logging (On Premises).
- Refer to Requirements for custom contextual cross-launch resources, on page 639.
- If needed for the resource you will link to, create or obtain an account and the credentials needed for access. Optionally, assign and distribute credentials for each user who needs access.
- Determine the syntax of the query link for the resource that you will link to:

  Access the resource via browser and, using the documentation for that resource as needed, formulate the query link needed to search for a specific sample of the type of information you want your query link to find, such as an IP address.

  Run the query, then copy the resulting URL from the browser's location bar.

  For example, you might have the query URL `https://www.talosintelligence.com/reputation_center/lookup?search=10.10.10.10`.

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Contextual Cross-Launch**.

Step 2 Click **New Cross-launch**.

In the form that appears, all fields marked with an asterisk require a value.

Step 3 Enter a unique resource name.

Step 4 Paste the working URL string from your resource into the **URL Template** field.

Step 5 Replace the specific data (such as an IP address) in the query string with an appropriate variable: Position your cursor, then click a variable (for example, **IP**) once to insert the variable.

In the example from the "Before You Begin" section above, the resulting URL might be `https://www.talosintelligence.com/reputation_center/lookup?search={IP}`. When the contextual cross-launch link is used, the {IP} variable in the URL will be replaced by the IP address that the user right-clicks on in the event viewer or dashboard.

For a description of each variable, hover over the variable.

You can create multiple contextual cross-launch links for a single tool or service, using different variables for each.

Step 6 Click **Test with example data** ( ) to test your link with example data.

Step 7 Fix any problems.

Step 8 Click **Save**.

Investigate events using contextual Cross-Launch

Contextual cross-launch allows you to right-click on event data and immediately access external analysis tools and threat intelligence resources. The available options depend on the data type you select, ensuring you only see relevant investigation resources.

Before you begin

If the resource you will access requires credentials, make sure you have those credentials.

Follow these steps to investigate events using contextual cross-launch:

Procedure

Step 1 Navigate to one of the following pages in the Secure Firewall Management Center that shows events:

- A dashboard (**Overview** > **Dashboards heading** > **Dashboard**), or
- An event viewer page (any menu option under the **Analysis** menu that includes a table of events.)

Step 2 Right-click the event of interest and choose the contextual cross-launch resource to use.

If necessary, scroll down in the context menu to see all available options.

The data type you right-click on determines the options you see; for example, if you right-click an IP address, you will only see contextual cross-launch options that are relevant to IP addresses.

For example, to get threat intelligence from Cisco Talos about a source IP address in the intrusion event, choose **Talos SrcIP** or **Talos IP**.

If a resource includes multiple variables, the option to choose that resource is available only for events that have a single possible value for each included variable.

The contextual cross-launch resource opens in a separate browser window.

It may take some time for the query to be processed, depending on the amount of data to be queried, speed of and demand on the resource, and so on.

Step 3 Sign in to the resource if necessary.

Configure cross-launch links for Secure Network Analytics

You can cross-launch from event data in Firewall Threat Defense to related data in your Secure Network Analytics appliance. For more information about the Secure Network Analytics product, visit the Cisco Security Analytics and Logging product page.

For general information about contextual cross-launching, refer to Investigate events using contextual Cross-Launch, on page 641.

Before you begin

- Ensure that a Secure Network Analytics appliance is available for use
- If your device sends events directly to Secure Network Analytics, disable syslog for those devices or assign them an access control policy without syslog configurations. This avoids duplicate events on the remote volume.
- Ensure you have these items:
  - Hostname or IP address of your manager.
  - Credentials for an account on your Secure Network Analytics appliance that has administrator privileges.

If you want to send Firewall Threat Defense data to your Secure Network Analytics appliance using Security Analytics and Logging (On Premises), refer to Remote data storage on a Secure Network Analytics appliance, on page 525.

Follow these steps to configure a set of cross-launch links to your Secure Network Analytics appliance:

Procedure

Step 1 Choose **Integration** > **Security Analytics & Logging**.

Step 2 You have two options for Secure Network Analytics deployment:

- Manager Only—Deploy a standalone Manager to receive and store events, and from which you can review and query events.
- Data Store—Deploy a Cisco Secure Network Analytics Flow Collector to receive events, a Secure Network Analytics Data Store to store events, and a Manager from which you can review and query events.

Choose the deployment option and click **Start**.

Step 3 Complete the wizard. For more information, refer to the Firewall Management Center Configuration section of Cisco Security Analytics and Logging Firewall Event Integration Guide.

Step 4 Choose **Analysis** > **Advanced** > **Contextual Cross-Launch** to verify your new cross-launch links.

If you want to make changes, return to this procedure; you cannot make changes directly on the contextual cross-launch listing page.

Note

- If you want to change these links later, return to this procedure; you cannot make changes directly on the contextual cross-launch listing page.

- You can manually create additional links to cross-launch into your Secure Network Analytics appliance using the procedure in Add contextual cross-launch resources, on page 640, but those links remain independent of the auto-created resources and you must manage them manually.

What to do next

Use your Secure Network Analytics credentials to cross-launch from an event into the Secure Network Analytics event viewer.

To cross launch from an event in the Firewall Management Center event viewer or dashboard, right-click a relevant event's table cell and choose the appropriate option.

Processing queries may take some time, depending on the amount of data, data transfer speed on the Secure Network Analytics Manager, and other factors.

Sending syslog messages for security events

Sending syslog messages for security events is a data transmission capability that

- sends data related to connection, security intelligence, intrusion, and file and malware events via syslog to a Security Information and Event Management (SIEM) tool or another external event storage and management solution,
- supports events displayed in the Firewall Management Center web interface in the menu options under the Analysis menu, which are sometimes collectively referred to as "the event viewer", and
- handles events that are also sometimes referred to as Snort® events.

Syslog forwarding behavior changes

Note In version 7.2.1, syslog traffic was allowed to be forwarded using route lookup. This enabled the traffic to be forwarded regardless of the interface specified in the logging host configuration. However, in versions 7.2.5.1 and higher, the changes introduced in 7.2.1 were removed.

Thus, from 7.2.5.1 and higher versions, the configuration specified in logging host configuration precedes over the route lookup and the syslog traffic is forwarded from the specified interface.

Security event syslog platform settings

This topic lists the platform settings relevant to syslog messages for security events and notes on applying policy changes.

To find more about the platform settings that apply to syslog messages for security events, refer to *Threat Defense Platform Settings that Apply to Security Event Syslog Messages* in the *Cisco Secure Firewall Management Center Device Configuration Guide* .

If you make changes to syslog settings in any policy, you must redeploy for the changes to take effect.

Best practices for configuring security event syslog messaging

Best practices for configuring security event syslog messaging are recommended configuration approaches that

- ensure consistent syslog configuration across device types,
- optimize security event logging for external analysis tools, and
- avoid configuration conflicts between policy levels.

Configuration approaches by device type

Configuration requirements vary by device type and version.

Table 62: Security event syslog configuration by device

Device and Version	Configuration Location
All	If you will use syslog or store events externally, avoid special characters in object names such as policy and rule names. Object names should not contain special characters, such as commas, that the receiving application may use as separators.
Secure Firewall Threat Defense	**1.** Do the following to configure syslog in Firewall Threat Defense platform settings: **a.** Click **Devices** > **Platform Settings**. **b.** Edit the threat defense settings policy. **c.** In the left navigation pane, click **Syslog**. For more information, refer to *Threat Defense Platform Settings That Apply to Security Event Syslog Messages* in the *Cisco Secure Firewall Management Center Device Configuration Guide*. **2.** In your access control policy **Logging** tab, opt to use the Firewall Threat Defense platform settings. **3.** (For intrusion events) Configure intrusion policies to use the settings in your access control policy **Logging** tab. (This is the default.) Overriding any of these settings is not recommended. For essential details, refer to Send security event syslog messages from your Firewall Threat Defense devices , on page 645.
All other devices	**1.** Create an alert response. **2.** Configure access control policy Logging to use the alert response. **3.** (For intrusion events) Configure syslog settings in intrusion policies. For complete details, refer to Configure security event syslog messages for Classic devices, on page 647.

Send security event syslog messages from your Firewall Threat Defense devices

This task outlines the recommended procedure for sending syslog messages for connection, Security Intelligence, intrusion, file, and malware events from devices managed by Secure Firewall Management Center.

Note Many Firewall Threat Defense syslog settings are not applicable to security events. Configure only the options described in this procedure.

Before you begin

- In Secure Firewall Management Center, configure policies to generate security events and verify that the events you expect to see appear in the applicable tables under the Analysis menu.
- Gather the syslog server IP address, port, and protocol (UDP or TCP):
- Ensure that your devices can reach the syslog server(s).
- Confirm that the syslog server(s) can accept remote messages.
- For additional information about connection logging, refer to Connection Logging, on page 737.

Follow these steps to configure security event syslog messaging:

Procedure

Step 1 Configure syslog settings for your Firewall Threat Defense device:

a) Click **Devices** > **Platform Settings**.
b) Click **Edit** to edit the platform settings policy associated with your Firewall Threat Defense device.
c) Click **Syslog** from the navigation pane.
d) Click **Syslog Servers** and click **Add** (+) to enter server, protocol, interface, and related information.

If you have questions about options on this page, refer to *Cisco Secure Firewall Management Center Device Configuration Guide*.

e) Click **Syslog Settings** and configure the following settings:

- **Enable timestamp on syslog messages**
- **Timestamp Format**
- **Enable syslog device ID**

f) Click **Logging Setup**.
g) On the **Basic Logging Settings**, select whether or not to **Send syslogs in EMBLEM format**.
h) Click **Save** to save your settings.

Step 2 Configure general logging settings for the access control policy (including file and malware logging):

a) Click **Policies** > **Access Control heading** > **Access Control**.
b) Edit the applicable access control policy.
c) From the **More** drop-down, click **Logging**.

d) Firewall Threat Defense 6.3 and later: Select **Use the syslog settings configured in the Threat Defense Platform Settings policy deployed on the device**.
e) (Optional) Select a **Syslog Severity**.
f) If you want to send file and malware events, select **Send Syslog messages for File and Malware events**.
g) Click **Save**.

Step 3 Enable logging for Security Intelligence events for the access control policy:

a) In the same access control policy, click the **Security Intelligence** tab.
b) In each of the following locations, click **Logging** ( ) and enable beginning and end of connections and **Syslog Server**:
   - Beside **DNS Policy**.
   - In the **Block List** box, for **Networks** and for **URLs**.
c) Click **Save**.

Step 4 Enable syslog logging for each rule in the access control policy:

a) In the same access control policy, click **Access Control**, and then choose **Add Rule**.
b) Select a rule to edit, and select the **Logging** tab in the rule.
c) Choose whether to log the beginning or end of connections, or both.

   (Connection logging generates a lot of data; logging both beginning and end generates roughly double that much data. Not every connection can be logged both at beginning and end.)

d) If you want to log file events, select **Log Files**.
e) Enable **Syslog Server**.
f) Verify that the rule is " **Using default syslog configuration in Access Control Logging**."
g) Click **Confirm**.
h) Repeat for each rule in the policy.

Step 5 If you send intrusion events:

a) Navigate to the intrusion policy associated with your access control policy.
b) In your intrusion policy, from **Advanced Settings** drop-down menu, click **Syslog Alerting**, and then click **Enabled**.
c) If necessary, click **Edit**
d) Enter options:

Option	Value
Logging Host	Unless you will send intrusion event syslog messages to a different syslog server than you will send other syslog messages, leave this blank to use the settings you have configured above.
Facility	This setting is applicable only if you specify a Logging Host on this page. For descriptions, refer to SYSLOG alert facilities, on page 565. .
Severity	This setting is applicable only if you specify a Logging Host on this page. For descriptions, refer to Syslog severity levels, on page 566. .

e) Click **Back**.
f) Click **Policy Information** in the left navigation pane.
g) Click **Commit Changes**.

What to do next

- (Optional) Configure different logging settings for individual policies and rules.

  Refer to the applicable table rows in Configuration Locations for Syslogs for Connection and Security Intelligence Events (All Devices), on page 648 .

  These settings will require syslog alert responses, which are configured as described in Create a syslog alert response, on page 566. They do not use the platform settings you configured in this procedure.

- To configure security event syslog logging for Classic devices, refer to Configure security event syslog messages for Classic devices, on page 647.

- If you are done making changes, deploy your changes to managed devices.

Configure security event syslog messages for Classic devices

Syslog messages generated by Classic devices contain critical security events that help track incidents and maintain compliance. To ensure proper logging, configure both the alert response and relevant policies.

Before you begin

- Configure policies to generate security events.
- Ensure that your devices can reach the syslog server(s).
- Confirm that the syslog server(s) can accept remote messages.
- For important information about connection logging, refer to Connection Logging, on page 737.

Follow these steps to configure security event syslog messages for Classic devices:

Procedure

Step 1 Configure an alert response for your Classic devices.

For more information, refer to Create a syslog alert response, on page 566.

Step 2 Configure syslog settings in the access control policy:

a) Click **Policies** > **Access Control heading** > **Access Control**.
b) Edit the applicable access control policy.
c) Click **Logging**.
d) Select **Send using specific syslog alert**.
e) Select the **Syslog Alert** you created above.
f) Click **Save**.

Step 3 If sending file and malware events:

a) Select **Send Syslog messages for File and Malware events**.
b) Click **Save**.

Step 4 If sending intrusion events:

a) Navigate to the intrusion policy associated with your access control policy.
b) In your intrusion policy, from **Advanced Settings** drop-down menu, click **Syslog Alerting**, and then click **Enabled**.
c) If necessary, click **Edit**
d) Enter options:

Option	Value
Logging Host	Unless you will send intrusion event syslog messages to a different syslog server than you will send other syslog messages, leave this blank to use the settings you have configured above.
Facility	This setting is applicable only if you specify a Logging Host on this page. For more information, refer to SYSLOG alert facilities, on page 565.
Severity	This setting is applicable only if you specify a Logging Host on this page. For more information, refer to Syslog severity levels, on page 566.

e) Click **Back**.
f) Click **Policy Information** in the navigation pane.
g) Click **Commit Changes**.

What to do next

- (Optional) Configure different logging settings for individual access control rules. For more information, refer to the applicable table rows in Configuration Locations for Syslogs for Connection and Security Intelligence Events (All Devices), on page 648. These settings will require syslog alert responses, which are configured as described in Create a syslog alert response, on page 566. They do not use the settings you configured above.
- To configure security event syslog logging for Firewall Threat Defense devices, refer to Send security event syslog messages from your Firewall Threat Defense devices , on page 645.

Configuration locations for security event syslogs

These topics describe how to configure location for security event syslogs.

Configuration Locations for Syslogs for Connection and Security Intelligence Events (All Devices)

There are many places to configure logging settings. Use the table below to ensure that you set the options you need.

Important

- Pay careful attention when configuring syslog settings, especially when using inherited defaults from other configurations. Some options may not be available to all managed device models and software versions, as noted in the table below.
- For important information when configuring connection logging, see the chapter on Connection Logging, on page 737.

Configuration Location	Description and More Information
Devices > **Platform Settings**, Threat Defense Settings policy, **Syslog**	This option applies only to Firewall Threat Defense devices. Settings you configure here can be specified in the Logging settings for an Access Control policy and then used or overridden in the remaining policies and rules in this table. See *Cisco Secure Firewall Management Center Device Configuration Guide*.
Policies > **Access Control heading** > **Access Control**, <each policy>, **Logging**	Settings you configure here are the default settings for syslogs for all connection and security intelligence events, unless you override the defaults in descendant policies and rules at the locations specified in the remaining rows of this table. Recommended setting for Firewall Threat Defense devices: Use Threat Defense Platform Settings. For information, see *Cisco Secure Firewall Management Center Device Configuration Guide*. Required setting for all other devices: Use a syslog alert. If you specify a syslog alert, see Create a syslog alert response, on page 566. For more information about the settings on the Logging tab, see *Cisco Secure Firewall Management Center Device Configuration Guide*.
Policies > **Access Control heading** > **Access Control**, <each policy>, **Rules**, **Default Action** row, **Logging** ( )	Logging settings for the default action associated with an access control policy. See information about logging in *Cisco Secure Firewall Management Center Device Configuration Guide* and Logging connections with a policy default action, on page 753.
Policies > **Access Control heading** > **Access Control**, <each policy>, **Rules**, <each rule>, **Logging**	Logging settings for a particular rule in an access control policy. See information about logging in *Cisco Secure Firewall Management Center Device Configuration Guide*.

Configuration Location	Description and More Information
Policies > **Access Control heading** > **Access Control**, <each policy>, **Security Intelligence**, **Logging** ( )	Logging settings for Security Intelligence Block lists. Click these buttons to configure: • DNS Block List Logging Options • URL Block List Logging Options • Network Block List Logging Options (for IP addresses on the blocked list) See *Cisco Secure Firewall Management Center Device Configuration Guide*
Policies > **Access Control heading** > **SSL**, <each policy>, **Default Action** row, **Logging** ( )	Logging settings for the default action associated with an SSL policy. See Logging connections with a policy default action, on page 753.
Policies > **Access Control heading** > **SSL**, <each policy>, <each rule>, **Logging**	Logging settings for SSL rules. See *Cisco Secure Firewall Management Center Device Configuration Guide*.
Policies > **Access Control heading** > **Prefilter**, <each policy>, **Default Action** row, **Logging** ( )	Logging settings for the default action associated with a prefilter policy. See Logging connections with a policy default action, on page 753.
Policies > **Access Control heading** > **Prefilter**, <each policy>, <each prefilter rule>, **Logging**	Logging settings for each prefilter rule in a prefilter policy. See *Cisco Secure Firewall Management Center Device Configuration Guide*
Policies > **Access Control heading** > **Prefilter**, <each policy>, <each tunnel rule> , **Logging**	Logging settings for each tunnel rule in a prefilter policy. See *Cisco Secure Firewall Management Center Device Configuration Guide*
Additional syslog settings for Firewall Threat Defense cluster configurations:	The *Cisco Secure Firewall Management Center Device Configuration Guide* has multiple references to syslog; search the chapter for "syslog."

Syslog configuration locations for intrusion events (Firewall Threat Defense devices)

Syslog settings for intrusion policies can be specified in various locations. You may also inherit settings from the access control policy or Threat Defense Platform Settings.

The table summarizes each configuration location and its purpose.

Configuration Location	Description and More Information
Devices > **Platform Settings**, Threat Defense Settings policy, **Syslog**	Syslog destinations that you configure here can be specified in the Logging tab of an access control policy which can be the default for an intrusion policy. For more information, refer to *Cisco Secure Firewall Management Center Device Configuration Guide*.
Policies > **Access Control heading** > **Access Control**, <each policy>, **Logging**	Default setting for syslog destination for intrusion events, if the intrusion policy does not specify other logging hosts. See *Cisco Secure Firewall Management Center Device Configuration Guide*.
Policies > **Access Control heading** > **Intrusion**, <each policy>, **Advanced Settings**, enable **Syslog Alerting**, click **Edit**	To specify syslog collectors other than the destinations specified in the access control policy Logging tab, and to specify facility and severity, refer to Configure syslog alerting for intrusion events, on page 576. If you want to use the **Severity** or **Facility** or both as configured in the intrusion policy, you must also configure the logging hosts in the policy. If you use the logging hosts specified in the access control policy, the severity and facility specified in the intrusion policy will not be used.
Policies > **Access Control heading** > **Intrusion** > **Logging** > **IPS settings**	If you want to send syslog messages for IPS events. Default syslog settings configured are used for syslog destinations for IPS events.

Configuration Locations for Syslogs for Intrusion Events (Devices Other than Firewall Threat Defense)

- (Default) Access control policy *Cisco Secure Firewall Management Center Device Configuration Guide*, IF you specify a syslog alert (See Create a syslog alert response, on page 566.)
- Or see Configure syslog alerting for intrusion events, on page 576.

By default, the intrusion policy uses the settings in the Logging tab of the access control policy. If settings applicable to devices other than Firewall Threat Defense are not configured there, syslogs will not be sent for devices other than Firewall Threat Defense and no warning appears.

Configuration locations for syslogs for file and malware events

This reference provides information about where to configure syslogs for file and malware events in the system.

Configuration Location	Description and More Information
In an access control policy: **Policies > Access Control**, <each policy>, **Logging**	This is the main location for configuring the system to send syslogs for file and malware events. If you do not use the syslog settings in Firewall Threat Defense Platform Settings, you must also create an alert response. For more information, refer to Create a syslog alert response, on page 566.
In Firewall Threat Defense Platform Settings: **Devices > Platform Settings**, Threat Defense Settings policy, **Syslog**	These settings apply only to Firewall Threat Defense devices running supported versions, and only if you configure the Logging tab in the access control policy to use Firewall Threat Defense platform settings. See *Cisco Secure Firewall Management Center Device Configuration Guide*.
In an access control rule: **Policies > Access Control**, <each policy>, <each rule>, **Logging**	If you do not use the syslog settings in Firewall Threat Defense Platform Settings, you must also create an alert response. For more information, refer to Create a syslog alert response, on page 566.

Components of security event syslog messages

Describes the structure and components of security event syslog messages, including examples and detailed field descriptions.

Example security event message from Firewall Threat Defense (Intrusion Event)

```
  0       1              2              3       4  5    6

<37>2018-06-27 192.168.0.81 SFIMS : %FTD-5-430001:
192.168.1.10, DstIP: 192.168.1.102, SrcPort: 33994
Protocol: tcp, Priority: 2, GID: 133, SID: 17, Rev
Message: "DCE2_EVENT  SMB_INVALID_DSIZE", Classif
Potentially Bad Traffic, User: No Authentication R
Client: NetBIOS-ssn (SMB) client, ApplicationProto
(SMB), ACPolicy: test, NAPPolicy: Balanced Securit
Connectivity, InlineResult: Blocked
```

Table 63: Components of security event syslog messages

Item Number in Sample Message	Header Element	Description
0	PRI	The priority value that represents both Facility and Severity of the alert. The value appears in the syslog messages only when you enable logging in EMBLEM format using Firewall Management Center platform settings. If you enable logging of intrusion events through access control policy Logging tab, the PRI value is automatically displayed in the syslog messages. For information on how to enable the EMBLEM format, refer to *Cisco Secure Firewall Management Center Device Configuration Guide*. For information on PRI, refer to RFC5424.
1	Timestamp	Date and time the syslog message was sent from the device. • (Syslogs sent from Firewall Threat Defense devices) For syslogs sent using settings in the access control policy and its descendants, or if specified to use this format in the Threat Defense Platform Settings, the date format is the format defined in the ISO 8601 timestamp format as specified in RFC 5424 (yyyy-MM-ddTHH:mm:ssZ), where the letter Z indicates the UTC time zone. • (Syslogs sent from all other devices) For syslogs sent using settings in the access control policy and its descendants, the date format is the format defined in the ISO 8601 timestamp format as specified in RFC 5424 (yyyy-MM-ddTHH:mm:ssZ), where the letter Z indicates the UTC time zone. • Otherwise, it is the month, day, and time in UTC time zone, though the time zone is not indicated. To configure the timestamp setting in Threat Defense Platform Settings, refer to *Cisco Secure Firewall Management Center Device Configuration Guide*.
2	Device or interface from which the message was sent. This can be: • IP address of the interface • Device hostname • Custom device identifier	(For syslogs sent from Firewall Threat Defense devices) If the syslog message was sent using the Threat Defense Platform Settings, this is the value configured in **Syslog Settings** for the **Enable Syslog Device ID** option, if specified. Otherwise, this element is not present in the header. To configure this setting in Threat Defense Platform Settings, refer to *Cisco Secure Firewall Management Center Device Configuration Guide*.

Item Number in Sample Message	Header Element	Description
3	Custom value	If the message was sent using an alert response, this is the **Tag** value configured in the alert response that sent the message, if configured. (Refer to Create a syslog alert response, on page 566.) Otherwise, this element is not present in the header.
4	%FTD	Type of device that sent the message. %FTD is Secure Firewall Threat Defense.
5	Severity	The severity specified in the syslog settings for the policy that triggered the message. For severity descriptions, refer to *Severity Levels* in the *Cisco Secure Firewall Management Center Device Configuration Guide* or Syslog severity levels, on page 566.
6	Event type identifier	• 430001: Intrusion event • 430002: Connection event logged at beginning of connection • 430003: Connection event logged at end of connection • 430004: File event • 430005: File malware event
--	Facility	Refer to Facility values and locations for security event syslog messages, on page 655.
--	Remainder of message	Fields and values separated by colons. Fields with empty or unknown values are omitted from messages. For field descriptions, refer to: • Connection and Security Intelligence Event Fields, on page 758 • Intrusion event fields, on page 790 • File and malware event fields, on page 841 **Note** Field description lists include both syslog fields and fields visible in the event viewer (menu options under the Analysis menu in the Firewall Management Center web interface.) Fields available via syslog are labeled as such. Some fields visible in the event viewer are not available via syslog. Also, some syslog fields are not included in the event viewer (but may be available via search), and some fields are combined or separated.

Facility values and locations for security event syslog messages

Facility values are not generally relevant to syslog messages generated for security events. However, some scenarios may require facility values to be specified.

The table summarizes how facility information is included or derived for different device types and security event categories.

Device	To Include Facility in Connection Events	To Include Facility in Intrusion Events	Location in Syslog Message
Firewall Threat Defense	Use the EMBLEM option in Threat Defense Platform Settings. Facility is always **ALERT** for connection events when sending syslog messages using Threat Defense Platform Settings.	Use the EMBLEM option in Threat Defense Platform Settings or configure logging using the syslog settings in the intrusion policy. If you use the intrusion policy, you must also specify the logging host in the intrusion policy settings. Enable syslog alerting and configure facility and severity on the intrusion policy. For more information, refer to Configure syslog alerting for intrusion events, on page 576.	Facility does not appear in the message header, but the syslog collector can derive the value based on RFC 5424, section 6.2.1.
Devices other than Firewall Threat Defense	Use an ALERT response.	Use the syslog setting in the intrusion policy advanced settings or an ALERT response identified in the access control policy Logging tab.	

For more information, refer to Facility and severity values for intrusion syslog alerts, on page 577 and Create a syslog alert response, on page 566.

Syslog message types supported by Secure Firewall devices

Secure Firewall devices can send multiple types of syslog messages.

The table summarizes each message type and provides references for configuration or further information.

Syslog Data Type	See
Audit logs from Firewall Management Center	Stream audit logs to syslog, on page 54 and the Audit and Syslog, on page 423 chapter
Device health and network-related logs from Firewall Threat Defense devices	*Cisco Secure Firewall Management Center Device Configuration Guide*
Connection, security intelligence, and intrusion event logs from Firewall Threat Defense devices	Security event syslog platform settings, on page 643.

Syslog Data Type	See
Connection, security intelligence, and intrusion event logs from Classic devices	Security event syslog platform settings, on page 643
Logs for file and malware events	Security event syslog platform settings, on page 643
IPS Settings	Send Syslog messages for IPS events. Syslog configuration locations for intrusion events (Firewall Threat Defense devices), on page 650

Limitations of syslog for security events

Syslog limitations for security events are system constraints that

- restrict certain object naming conventions and special character usage,
- introduce delays in event delivery to syslog collectors, and
- limit availability of specific file and malware event data types.

Syslog security event restrictions

The following limitations apply when using syslog for security events:

- If you will use syslog or store events externally, avoid special characters in object names such as policy and rule names. Object names should not contain special characters, such as commas, that the receiving application may use as separators.
- It may take up to 15 minutes for events to appear on your syslog collector.
- Data for the following file and malware events is not available via syslog:
  - Retrospective events
  - Events generated by Secure Endpoint

eStreamer server streaming

An eStreamer server is a streaming service that

- streams several kinds of event data from a Secure Firewall Management Center to custom-developed client applications,
- delivers event data either as binary data or as fully-qualified events in plaintext format, and
- requires configuration to specify which event types can be transmitted to requesting clients.

eStreamer server configuration requirements

Before the appliance you want to use as an eStreamer server can begin streaming eStreamer events to an external client, you must configure the eStreamer server to send events to clients, provide information about

the client, and generate a set of authentication credentials to use when establishing communication. You can perform all of these tasks from the appliance's user interface. Once your settings are saved, the events you selected will be forwarded to eStreamer clients when requested.

You can control which types of events the eStreamer server is able to transmit to clients that request them.



Note eStreamer Client has been deprioritized, and the eNcoreCLI is End of Life as of 7/1/2024.

For more information, refer to *Secure Firewall Management Center Event Streamer Integration Guide*.

Table 64: Event types transmittable by the eStreamer server

Event type	Description
Intrusion Events	intrusion events generated by managed devices
Intrusion Event Packet Data	packets associated with intrusion events
Intrusion Event Extra Data	additional data associated with an intrusion event such as the originating IP addresses of a client connecting to a web server through an HTTP proxy or load balancer
Discovery Events	Network discovery events
Correlation and Allow List Events	correlation and compliance allow list events
Impact Flag Alerts	impact alerts generated by the Firewall Management Center
User Events	user events
Malware Events	malware events
File Events	file events
Connection Events	information about the session traffic between your monitored hosts and all other hosts.

Comparison of syslog and eStreamer for security eventing

Generally, organizations that do not currently have significant existing investment in eStreamer should use syslog rather than eStreamer to manage security event data externally.

Syslog	eStreamer
No customization required	Significant customization and ongoing maintenance required to accommodate changes in each release
Standard	Proprietary
Syslog standard does not protect against data loss, especially when using UDP	Protection against data loss

Syslog	eStreamer
Sends directly from devices	Sends from Firewall Management Center, adding processing overhead
Support for file and malware events, connection events (including security intelligence events) and intrusion events.	Support for all event types listed in eStreamer server streaming, on page 656.
Some event data can be sent only from Firewall Management Center. For more information, refer to Data sent only via eStreamer, not via syslog, on page 658.	Includes data that cannot be sent via syslog directly from devices. For more information, refer to Data sent only via eStreamer, not via syslog, on page 658.

Data sent only via eStreamer, not via syslog

This reference provides information about data that is available only from Secure Firewall Management Center and thus cannot be sent via syslog from devices.

- Packet Logs
- Intrusion Event Extra Data events

  For description, refer to eStreamer server streaming, on page 656.
- Statistics and aggregate events
- Network Discovery events
- User activity and login events
- Correlation events
- For malware events:
  - retrospective verdicts
  - ThreatName and Disposition, unless information about the relevant SHAs has already been synchronized to the device
- These fields:
  - Impact and ImpactFlag fields

    For description, refer to eStreamer server streaming, on page 656.
  - the IOC_Count field
- Most raw IDs and UUIDs.

  Exceptions:
  - Syslogs for connection events do include these: FirewallPolicyUUID, FirewallRuleID, TunnelRuleID, MonitorRuleID, SI_CategoryID, SSL_PolicyUUID, and SSL_RuleID
  - Syslogs for intrusion events do include IntrusionPolicyUUID, GeneratorID, and SignatureID

- Extended metadata, including but not limited to:
  - User details provided by LDAP, such as full name, department, phone number, etc.

    Syslog only provides usernames in the events.
  - Details for state-based information such as SSL Certificate details.

    Syslog provides basic information like the certificate fingerprint, but will not provide other certificate details like the cert CN.
  - Detailed application information, such as App Tags and Categories.

    Syslog provides only Application names.

  Some metadata messages also include extra information about the objects.
- Geolocation information

Choose eStreamer event types

The **eStreamer Event Configuration** checkboxes control which events the eStreamer server can transmit. Your client must still specifically request the types of events you want it to receive in the request message it sends to the eStreamer server. For more information, refer to *Secure Firewall Management Center Event Streamer Integration Guide*.

In a multidomain deployment, you can configure eStreamer Event Configuration at any domain level. However, if an ancestor domain has enabled a particular event type, you cannot disable that event type in the descendant domains.

You must be an Admin user to perform this task, for Firewall Management Center.

Procedure

Step 1 Choose **Integration** > **Other Integrations**.

Step 2 Click **eStreamer**.

Step 3 Under **eStreamer Event Configuration**, check or clear the check boxes next to the types of events you want eStreamer to forward to requesting clients, described in eStreamer server streaming, on page 656.

Step 4 Click **Save**.

Configure eStreamer client communications

Before eStreamer can send eStreamer events to a client, you must add the client to the eStreamer server's peers database from the eStreamer page. You must also copy the authentication certificate generated by the eStreamer server to the client. After completing these steps you do not need to restart the eStreamer service to enable the client to connect to the eStreamer server.

In a multidomain deployment, you can create an eStreamer client in any domain. The authentication certificate allows the client to request events only from the client certificate's domain and any descendant domains. The

eStreamer configuration page shows only clients associated with the current domain, so if you want to download or revoke a certificate, switch to the domain where the client was created.

You must be an Admin or Discovery Admin user to perform this task, for Firewall Management Center.

Procedure

Step 1 Choose **Integration** > **Other Integrations**.

Step 2 Click **eStreamer**.

Step 3 Click **Create Client**.

Step 4 In the **Hostname** field, enter the host name or IP address of the host running the eStreamer client.

If you want to encrypt the certificate file, enter a password in the **Password** field.

Note
If you have not configured DNS resolution, use an IP address.

Step 5 Click **Save**.
The eStreamer server now allows the host to access port 8302 on the eStreamer server and creates an authentication certificate to use during client-server authentication.

Step 6 Click **Download** ( ) next to the client hostname to download the certificate file.

Step 7 Save the certificate file to the appropriate directory used by your client for SSL authentication.

Step 8 To revoke access for a client, click **Delete** ( ) next to the host you want to remove.

Note
You do not need to restart the eStreamer service; access is revoked immediately.

Event analysis in splunk

Event analysis in Splunk is a security investigation capability that

- enables you to use the Cisco Secure Firewall app for Splunk as an external tool to display and work with Secure Firewall event data,
- allows you to hunt and investigate threats on your network, and
- requires eStreamer for advanced functionality.

Implementation options

Use eStreamer with the Splunk App - This is an advanced functionality that requires eStreamer configuration. Refer to eStreamer server streaming, on page 656.

For more information about the Splunk App, refer to User Guide for Cisco Secure Firewall App for Splunk.

IBM QRadar events

IBM QRadar events are security functionalities that

- provide an alternate way to display event data using the Cisco Firepower app for IBM QRadar,
- help you analyze, hunt for, and investigate threats to your network, and
- require eStreamer for advanced functionality.

Additional reference information

eStreamer is required for advanced capabilities. For more information, refer to eStreamer server streaming, on page 656.

For more information, refer to https://www.cisco.com/c/en/us/td/docs/security/firepower/integrations/QRadar/integration-guide-for-the-cisco-firepower-app-for-ibm-qradar.html.

History for analyzing event data using external tools

Documents the history of features for analyzing event data using external tools, including enhancements and deprecations.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Deprecated: SecureX Ribbon	Any	Any	SecureX Ribbon is deprecated. If you have installed the **Cisco SecureX Ribbon** browser extension in your Firefox browser and are experiencing compatibility errors while using Firewall Management Center, remove the SecureX Ribbon extension. To remove the extension, open Firefox, go to the browser's add-ons or extensions manager, locate the **Cisco SecureX Ribbon** extension, and remove or disable it. Restart Firefox to apply the changes.
Register your firewall deployment to Cisco Security Cloud tenancy using your Cisco Security Cloud Sign-On account and your Security Cloud Control tenant.	Any	Any	You can now register your Firewall Management Center and its managed devices to the Cisco Security Cloud using your Cisco Security Cloud Sign-On account and your Security Cloud Control tenant. Cisco SecureX is phased out and no longer available. If you have an active Cisco SecureX environments, you will continue to receive support from the Cisco Technical Assistance Center (TAC) through the product end-of-support. For more information, refer to Frequently Asked Questions.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
SecureX ribbon	7.0	Any	The SecureX ribbon pivots into SecureX for instant visibility into the threat landscape across your Cisco security products. To display the SecureX ribbon in Firewall Management Center, refer to the *Firepower and SecureX Integration Guide* at https://cisco.com/go/firepower-securex-documentation. New/Modified screens: New page: **System** > **SecureX**.
Send all connection events to the Cisco cloud	7.0	Any	You can now send all connection events to the Cisco cloud, rather than just sending high-priority connection events. New/Modified screens: New option on the **System** > **Integration** > **Cloud Services**.
Cross-launch to view data in Secure Network Analytics	6.7	Any	This feature introduces a quick way to create multiple entries for your Secure Network Analytics appliance on the Analysis > Contextual Cross-Launch page. These entries allow you to right-click a relevant event to cross-launch Secure Network Analytics and display information related to the data point from which you cross-launched. New menu item: **System** > **Logging** > **Security Analytics and Logging**. New page to configure sending events to Secure Network Analytics.
Contextual cross-launch from additional field types	6.7	Any	You can now cross-launch into an external application using the following additional types of event data: • Access control policy • Intrusion policy • Application protocol • Client application • Web application • Username (including realm) New menu options: Contextual cross-launch options are now available when right-clicking the above data types for events in Dashboard widgets and event tables on pages under the Analysis menu. Supported platforms: Secure Firewall Management Center
Integration with IBM QRadar	6.0 and later	Any	IBM QRadar users can use a new Firepower-specific app to analyze their event data. Available functionality is affected by your Firepower version. Refer to IBM QRadar events, on page 661.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Enhancements to integration with SecureX Threat Response	6.5	Any	• Support for regional clouds: • United States (North America) • Europe • Support for additional event types: • File and malware events • High-priority connection events These are connection events related to the following: • Intrusion events • Security Intelligence events • File and malware events Modified screens: New options on **System** > **Integration** > **Cloud Services**. Supported Platforms: All devices supported in this release, either via direct integration or syslog.
Syslog	6.5	Any	The AccessControlRuleName field is now available in intrusion event syslog messages.
Integration with Cisco Security Packet Analyzer	6.5	Any	Support for this feature was removed.
Integration with SecureX Threat Response	6.3 (via syslog, using a proxy collector) 6.4 (direct)	Any	Integrate Firepower intrusion event data with data from other sources for a unified view of threats on your network using the powerful analysis tools in SecureX Threat Response. Modified screens (version 6.4): New options on **System** > **Integration** > **Cloud Services**. Supported Platforms: Secure Firewall Threat Defense devices running version 6.3 (via syslog) or 6.4.
Syslog support for File and Malware events	6.4	Any	Fully-qualified file and malware event data can now be sent from managed devices via syslog. Modified screens: **Policies** > **Access Control** > **Access Control** > **Logging**. Supported Platforms: All managed devices running version 6.4.

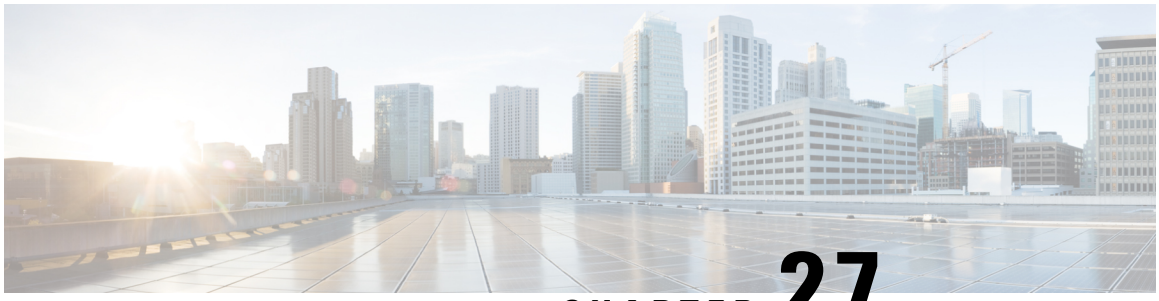
Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Integration with Splunk	Supports all 6.x versions	Any	Splunk users can use a new, separate Splunk app, Cisco Secure Firewall app for Splunk, to analyze events. Available functionality is affected by your Firepower version. Refer to Event analysis in splunk, on page 660.
Integration with Cisco Security Packet Analyzer	6.3	Any	Instantly query Cisco Security Packet Analyzer for packets related to an event, then click to examine the results in Cisco Security Packet Analyzer or download them for analysis in another external tool. New screens: **System** > **Integration** > **Packet Analyzer** **Analysis** > **Advanced** > **Packet Analyzer Queries**. New menu options: **Query** > **Packet Analyzer** menu item when right-clicking on an event on Dashboard pages and event tables on pages under the Analysis menu. Supported platforms: Secure Firewall Management Center
Contextual cross-launch	6.3	Any	Feature introduced: Right-click an event to look up related information in predefined or custom URL-based external resources. New screens: **Analysis** > **Advanced** > **Contextual Cross-Launch.** New menu options: Multiple options when right-clicking on an event on Dashboard pages and event tables on pages under the Analysis menu. Supported platforms: Secure Firewall Management Center
Syslog messages for connection and intrusion events	6.3	Any	Ability to send fully-qualified connection and intrusion events to external storage and tools via syslog, using new unified and simplified configurations. Message headers are now standardized and include event type identifiers, and messages are smaller because fields with unknown and empty values are omitted. Supported Platforms: • All new functionality: Firewall Threat Defense devices running version 6.3. • Some new functionality: Non- Firewall Threat Defense devices running version 6.3. • Less new functionality: All devices running versions older than 6.3. For more information, refer to the topics under Sending syslog messages for security events, on page 643 and subtopics.
eStreamer	6.3	Any	Moved eStreamer content from the Host Identity Sources chapter to this chapter and added a summary comparing eStreamer to syslog.



PART VII

Workflows and Tables

- Workflows, on page 667
- Event Search, on page 707
- Custom Workflows, on page 717
- Custom Tables, on page 725



CHAPTER 27

Workflows

The following topics describe how to use workflows:

- Workflows, on page 667
- Predefined workflows, on page 667
- Custom table workflows, on page 676
- Use workflows, on page 676
- Work with the unified event viewer, on page 702
- Bookmarks, on page 703
- History for workflows, on page 704

Workflows

A workflow is a tailored series of data pages on the Firewall Management Center web interface that analysts can use to evaluate events generated by the system.

Workflow types

These types of workflows are available in Firewall Management Center:

- Predefined workflows: Preset workflows delivered with the system. You cannot edit or delete a predefined workflow. You can, however, copy a predefined workflow and use it as the basis for a custom workflow.
- Saved custom workflows: Custom workflows based on saved custom tables delivered with the Firewall Management Center. You can edit, delete, and copy these workflows.
- Custom workflows: Workflows that you create and customize for your specific needs, or that the system generates automatically when you create custom tables. You can edit, delete, and copy these workflows.

The data displayed in a workflow often depends on such factors as how you license and deploy your managed devices, and whether you configure features that provide the data.

Predefined workflows

Predefined workflows are workflow templates that

- are delivered with the system and cannot be edited or deleted

- can be copied to serve as the basis for custom workflows, and
- provide standardized workflow structures for common use cases.

Predefined intrusion event workflows

This table describes the predefined intrusion event workflows included with the system.

Table 65: Predefined intrusion event workflows

Workflow name	Description
Destination Port	Because destination ports are usually tied to an application, this workflow can help you detect applications that are experiencing an uncommonly high volume of alerts. The Destination Port column can also help you identify applications that should not be present on your network.
Event-Specific	This workflow provides two useful features. Events that occur frequently may indicate: • false positives • a worm • a badly misconfigured network Events that occur infrequently are most likely evidence of a targeted attack and warrant special attention.
Events by Priority and Classification	This workflow lists events and their type in order of event priority, along with a count showing how many times each event has occurred.
Events to Destinations	This workflow provides a high-level view of which host IP addresses are being attacked and the nature of the attack; where available, you can also see information about the countries involved in attacks.
IP-Specific	This workflow shows which host IP addresses are generating the most alerts. Hosts with the greatest number of events are either public-facing and receiving worm-type traffic (indicating a good place to look for tuning) or require further investigation to determine the cause of the alerts. Hosts with the lowest counts also warrant investigation as they could be the subject of a targeted attack. Low counts may also indicate that a host may not belong on the network.
Impact and Priority	This workflow lets you find high-impact recurring events quickly. The reported impact level is shown with the number of times the event has occurred. Using this information, you can identify the high-impact events that recur most often, which might be an indicator of a widespread attack on your network.
Impact and Source	This workflow can help you identify the source of an attack in progress. The reported impact level is shown with the associated source IP address for the event. If, for example, events with a level 1 impact are coming from the same source IP address repeatedly, they may indicate an attacker who has identified vulnerable systems and is targeting them.
Impact to Destination	You can use this workflow to identify events repeatedly occurring on vulnerable computers, so you can address the vulnerabilities on those systems and stop any attacks in progress.
Source Port	This workflow indicates which servers are generating the most alerts. You can use this information to identify areas that require tuning, and to decide which servers require attention.

Workflow name	Description
Source and Destination	This workflow identifies host IP addresses sharing high levels of alerts. Pairs at the top of the list could be false positives, and may identify areas that require tuning. You can check pairs at the bottom of the list for targeted attacks, for users accessing resources they should not be accessing, or for hosts that do not belong on the network.

Predefined malware workflows

This table describes the predefined malware workflows included on the Firewall Management Center. All predefined malware workflows use the table view of malware events.

Table 66: Predefined malware workflows

Workflow name	Description
Malware Summary	This workflow provides a list of the malware detected in network traffic or by Secure Endpoint Connectors, grouped by individual threat.
Malware Event Summary	This workflow provides a quick breakdown of the different malware event types and subtypes.
Hosts Receiving Malware	This workflow provides a list of host IP addresses that have received malware, grouped by the malware files' associated dispositions.
Hosts Sending Malware	This workflow provides a list of host IP addresses that have sent malware, grouped by the malware files' associated dispositions.
Applications Introducing Malware	This workflow provides a list of host IP addresses that have received files, grouped by the associated malware dispositions for those files.

Predefined file workflows

This table describes the predefined file event workflows included in Firewall Management Center. All predefined file event workflows use the table view of file events.

Table 67: Predefined file workflows

Workflow name	Description
File Summary	This workflow provides a quick breakdown of the different file event categories and types, along with any associated malware dispositions.
Hosts Receiving Files	This workflow provides a list of host IP addresses that have received files, grouped by the associated malware dispositions for those files.
Hosts Sending Files	This workflow provides a list of host IP addresses that have sent files, grouped by the associated malware dispositions for those files.

Predefined captured file workflows

The predefined captured file workflows included on the Firewall Management Center use the table view of captured files.

Table 68: Predefined captured file workflows

Workflow name	Description
Captured File Summary	This workflow provides a breakdown of captured files based on type, category, and threat score.
Dynamic Analysis Status	This workflow provides a count of captured files based on whether they have been submitted for dynamic analysis.

Predefined connection data workflows

All predefined connection data workflows use the table view of connection data to provide various analytical perspectives of network activity and traffic patterns.

Table 69: Predefined connection data workflows

Workflow name	Description
Connection Events	This workflow provides a summary view of basic connection and detected application information, which you can then use to drill down to the table view of events.
Connections by Application	This workflow contains a graph of the 10 most active applications on the monitored network segment, based on the number of detected connections.
Connections by Initiator	This workflow contains a graph of the 10 most active host IP addresses on the monitored network segment, based on the number of connections where the host initiated the connection transaction.
Connections by Port	This workflow contains a graph of the 10 most active ports on the monitored network segment, based on the number of detected connections.
Connections by Responder	This workflow contains a graph of the 10 most active host IP addresses on the monitored network segment, based on the number of connections where the host IP was the responder in the connection transaction.
Connections over Time	This workflow contains a graph of the total number of connections on the monitored network segment over time.

Workflow name	Description
Traffic by Application	This workflow contains a graph of the 10 most active applications on the monitored network segment, based on the number of kilobytes transmitted. Application counts reflect each detector that matched against an application connection. The same application session may be represented more than once in the list depending on whether an application protocol, web application, client detector, or internal detector matched the traffic, as well as whether the traffic originated from a mobile device or was part of an encrypted session. If the application was seen in a client flow and no specific client detector exists, a generic client may be reported. For example, you may see the same session of YouTube traffic reported as **YouTube** (because it matched a YouTube web application detector) and as **YouTube client** (because an internal YouTube detector matched against characteristics typically seen in a client session). Use the information in the connection events and network map for your network to determine more context for specific application connections.
Traffic by Initiator	This workflow contains a graph of the 10 most active host IP addresses on the monitored network segment, based on the total number of kilobytes transmitted from each address.
Traffic by Port	This workflow contains a graph of the 10 most active ports on the monitored network segment, based on the number of kilobytes transmitted.
Traffic by Responder	This workflow contains a graph of the 10 most active host IP addresses on the monitored network segment, based on the total number of kilobytes received by each address.
Traffic over Time	This workflow contains a graph of the total kilobytes transmitted on the monitored network segment over time.
Unique Initiators by Responder	This workflow contains a graph of the 10 most active responding host IP addresses on the monitored network segment, based on the number of unique initiators that contacted each address.
Unique Responders by Initiator	This workflow contains a graph of the 10 most active initiating host IP addresses on the monitored network segment, based on the number of unique responders that the addresses contacted.

Predefined security intelligence workflows

This table describes the predefined Security Intelligence workflows included on the Firewall Management Center. All the predefined Security Intelligence workflows use the table view of Security Intelligence events.

Table 70: Predefined security intelligence workflows

Workflow name	Description
Security Intelligence events	This workflow provides a summary view of basic Security Intelligence and detected application information, which you can then use to drill down to the table view of events.
Security Intelligence summary	This workflow is identical to the Security Intelligence events workflow, but begins with the Security Intelligence Summary page, which lists security intelligence events by category and count only.

Workflow name	Description
Security Intelligence with DNS details	This workflow is identical to the Security Intelligence events workflow, but begins with the Security Intelligence with DNS Details page, which lists Security Intelligence events by category and DNS-related characteristics.

Predefined host workflows

This table describes the predefined workflows that you can use with host data.

Table 71: Predefined host workflows

Workflow name	Description
Hosts	This workflow contains a table view of hosts followed by the host view. Workflow views based on the Hosts table allow you to easily view data on all IP addresses associated with a host.
Operating System Summary	You can use this workflow to analyze the operating systems in use on your network.

Predefined indications of compromise workflows

This table describes the predefined workflows that you can use with indications of compromise (IOC) data.

Table 72: Predefined indications of compromise workflows

Workflow name	Description
Host Indications of Compromise	This workflow begins with a summary view of IOC data grouped by count and category, and provides a detail view that further subdivides the summary data by event type. Access this workflow via the **Analysis** > **Hosts heading** > **Hosts** menu.
Indications of Compromise by Host	You can use this workflow to gauge which hosts on your network are most likely to be compromised (based on IOC data). Access this workflow via the **Analysis** > **Hosts heading** > **Hosts** menu.
User Indications of Compromise	This workflow begins with a summary view of IOC data grouped by count and category, and provides a detail view that further subdivides the summary data by event type. Access this workflow via the **Analysis** > **Users heading** > **Users** menu.
Indications of Compromise by User	Use this workflow to gauge which users on your network are most likely to be involved in potential compromises (based on IOC data.) Access this workflow via the **Analysis** > **Users heading** > **Users** menu.

Predefined applications workflows

This table describes the predefined workflows that you can use with application data.

Table 73: Predefined applications workflows

Workflow name	Description
Application Business Relevance	You can use this workflow to analyze running applications of each estimated business relevance level on your network, so you can monitor appropriate use of your network resources.
Application Category	You can use this workflow to analyze running applications of each category (such as email, search engine, or social networking) on your network, so you can monitor appropriate use of your network resources.
Application Risk	You can use this workflow to analyze running applications of each estimated security risk level on your network, so you can estimate the potential risk of users' activity and take appropriate action.
Application Summary	You can use this workflow to obtain detailed information about the applications and associated hosts on your network, so you can closely examine host application activity.
Applications	You can use this workflow to analyze running applications on your network, so you can gain an overview of how the network is being used.

Predefined application details workflows

This table describes the predefined workflows that you can use with application detail and client data.

Table 74: Predefined application details workflows

Workflow Name	Description
Application Details	You can use this workflow to analyze the client applications on your network in more detail. The workflow then provides a table view of client applications, followed by the host view.
Clients	This workflow contains a table view of client applications, followed by the host view.

Predefined servers workflows

This table describes the predefined workflows that you can use with server data.

Table 75: Predefined servers workflows

Workflow name	Description
Network Applications by Count	You can use this workflow to analyze the most frequently used applications on your network.
Network Applications by Hit	You can use this workflow to analyze the most active applications on your network.
Server Details	You can use this workflow to analyze the vendors and versions of detected server application protocols in detail.
Servers	This workflow contains a table view of applications followed by the host view.

Predefined host attributes workflows

This table describes the predefined workflow that you can use with host attribute data.

Table 76: Predefined host attributes workflows

Workflow name	Description
Attributes	You can use this workflow to monitor IP addresses and status of the hosts on your network.

The predefined discovery events workflow

This table describes the predefined workflow that you can use to view discovery and identity data.

Table 77: Predefined discovery event workflows

Workflow name	Description
Discovery Events	This workflow provides a detailed list, in table view form, of discovery events, followed by the host view.

Predefined user workflows

This table describes the predefined workflow that you can use to view user discovery and user identity data.

Table 78: Predefined user workflows

Workflow name	Description
Active Sessions	This workflow provides a list of active sessions collected by user identity sources.
Users	This workflow provides a list of user information collected by user identity sources.

Predefined vulnerabilities workflows

This table describes the predefined vulnerabilities workflow included in Firewall Management Center.

Table 79: Predefined vulnerabilities workflows

Workflow name	Description
Vulnerabilities	You can use this workflow to review vulnerabilities in the database, including a table view of only those active vulnerabilities that apply to the detected hosts on your network. The workflow provides a vulnerability detail view, which contains a detailed description for every vulnerability that meets your constraints.

Predefined third-party vulnerabilities workflows

This table describes the predefined third-party vulnerabilities workflows included on the Firewall Management Center.

Table 80: Predefined Third-Party vulnerabilities workflows

Workflow name	Description
Vulnerabilities by IP Address	You can use this workflow to quickly see how many third-party vulnerabilities you have detected per host IP address on your monitored network.
Vulnerabilities by Source	You can use this workflow to quickly see how many third-party vulnerabilities you have detected per third-party vulnerability source, such as the QualysGuard Scanner.

Predefined correlation and Allow list workflows

There is a predefined workflow for each type of correlation data, allow list events, allow list violations, and remediation status events.

Table 81: Predefined correlation workflows

Workflow name	Description
Correlation Events	This workflow contains a table view of correlation events.
Allow List Events	This workflow contains a table view of allow list events.
Host Violation Count	This workflow provides a series of pages that list all the host IP addresses that violate at least one allow list.
Allow List Violations	This workflow includes a table view of allow list violations that lists all violations with the most recently detected violation at the top of the list. Each row in the table contains a single detected violation.
Status	This workflow contains a table view of remediation status, which includes the name of the policy that was violated and the name and status of the remediation that was applied.

Predefined system workflows

The system is delivered with some additional workflows, including system events such as audit events and health events, as well as workflows that list results from rule update imports and active scans.

Table 82: Additional predefined workflows

Workflow name	Description
Audit Log	This workflow contains a table view of the audit log that lists audit events.
Health Events	This workflow displays events triggered by the health monitoring policy.

Workflow name	Description
Rule Update Import Log	This workflow contains a table view listing information about both successful and failed rule update imports.
Scan Results	This workflow contains a table view listing each completed scan.

Custom table workflows

Custom table workflows are workflows that are automatically created when you create custom tables using data from two or more event types. These workflows allow you to view and analyze events associated with the custom table.

Custom table workflow usage

You can use the custom tables feature to create tables that combine data from multiple event types. For example, you can create tables and workflows that correlate intrusion event data with discovery data for simplified searching of critical system events.

When you create a custom table, the Firewall Management Center automatically creates a workflow that you can use to view the events associated with the table. The features in the workflow differ depending on which type of table you use. The specific features of the workflow depend on the type of table you use. For example, custom table workflows based on the intrusion event table always end with the packet view, while those based on discovery events end with the host view.

Unlike workflows based on the predefined event tables, workflows based on custom tables do not have links to other types of workflows.

Use workflows

Workflows provide structured navigation through event data using table views and drill-down pages. You can modify displays, constrain data views, and access additional information through various interactive elements.

Procedure

Step 1 Access the workflow page that you want to use. For instructions, refer to Workflow selection, on page 678.

Step 2 To navigate within the current workflow, use these options:

- To view all of the columns available in your chosen event data type, use table view pages. For more information, refer to Use table view pages, on page 684.
- To view a subset of the columns available in your chosen event data type, use drill-down pages. For more information, refer to Use drill-down pages, on page 684.
- To display the corresponding row in the next page of the workflow, click **Down-Arrow** (▼).

- To move among the pages of a multipage workflow, use the tools at the bottom of each page. For more information, refer to Workflow page traversal tools, on page 681.
- To view the same constraints applied within a workflow for a different type of event, click **Jump to** and choose the event view from the drop-down list.

Step 3 To modify the display of the current workflow, perform these steps:

- Check the check boxes by one or more rows on a page to indicate which row(s) you want to affect, then click one of the buttons at the bottom of the page (for example, **View**) to perform that action for all selected rows.
- Check the check box at the top of the row to select all the rows on the page, then click one of the buttons at the bottom of the page (for example, **View**) to perform that action for all rows on the page.
- Constrain the columns in the display by clicking **Close** () in the column heading that you want to hide. In the pop-up window that appears, click **Apply**

  Note
  To hide or show other columns, check or clear the appropriate check boxes before you click **Apply**. To add a disabled column back to the view, click the expand arrow to expand the search constraints, then click the column name under Disabled Columns.

- Constrain the data view by selected values for selected fields. For information, refer to Event view constraints, on page 698 and Compound event view constraints, on page 700.
- Change the time constraints on the event view. The date range located in the upper right corner of the page sets a time range for events to include in the workflow. For information, refer to Event time constraints, on page 692.

  Note
  Events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.

- To sort data by columns, click the name of a column. To reverse the sort order, click the column name again. The direction indicates which column the data is sorted by, and whether the sort is **Ascending** or **Descending**.
- Click a workflow page link to display that page using any active constraints. Workflow page links appear in the upper left corner of predefined workflow table views and drill-down pages, above events and below the workflow name.

Step 4 To view additional data within the current workflow, perform these steps:

- To view the file's trajectory map in a new window, click network file trajectory in file name and SHA-256 hash value columns. The icon is different depending on the file status. For more information, refer to File trajectory icons, on page 682.
- To display a pop-up window of the host profile associated with an IP address, click host profile in any IP address column. The icon is different depending on the file status. For more information, refer to Host profile icons, on page 682.

- To view the Dynamic Analysis Summary report for the highest threat score associated with a file, click threat score in any threat score column. The icon is different depending on the file's highest threat score. For more information, refer to Threat score icons, on page 683.
- To view user profile information, click **User** or, for users associated with an indication of compromise, **Red User** in any user identity column. The user icon is dimmed if that user cannot be in the database (that is, is an Secure Endpoint Connector user).
- To view vulnerability details for third-party vulnerabilities, click **Vulnerability** in any third-party vulnerability ID column.
- When viewing aggregated data points, hover your pointer over the flag to view the country name.
- When viewing individual data points, click flag to view further geolocation details described in Geolocation, on page 686.

Step 5 To view the same event type using a different workflow, click **(switch workflow)** next to the workflow title, and then choose the workflow you want to use.

Note
You cannot use a different workflow for scan results.

Workflow access by user role

Access to a workflow is determined by the user's role. This table provides information about which workflows each user role can access.

User role	Accessible workflows
Administrator	Can access any workflow, and are the only users who can access the audit log, scan results, and the rule update import log.
Maintenance User	Can access health events.
Security Analyst and **Security Analyst (Read Only)**	Can access intrusion, malware, file, connection, discovery, vulnerability, correlation, and health workflows.

Workflow selection

The system provides predefined workflows for the types of data listed in this table.

Table 83: Features using workflows

Feature	Menu path	Option
Connection events	**Analysis** > **Connections**	Events
Security Intelligence events	**Analysis** > **Connections**	Security Intelligence Events

Feature	Menu path	Option
Correlation events	**Analysis** > **Correlation**	Correlation Events Allow List Events Allow List Violations Status
Malware events	**Analysis** > **Files**	Malware Events
File events	**Analysis** > **Files**	File Events
Captured files	**Analysis** > **Files**	Captured Files
Host events	**Analysis** > **Hosts**	Network Map Hosts Indications of Compromise Applications Application Details Servers Host Attributes Discovery Events
Intrusion events	**Analysis** > **Intrusions**	Events Reviewed Events
User events	**Analysis** > **Users**	Active Sessions User Activity Users Indications of Compromise
Vulnerability events	**Analysis** > **Hosts**	Vulnerabilities Third-Party Vulnerabilities
Scan Results	**Policies** > **Actions** > **Scanners**	—
Health events	**System** (⚙) > **Health** > **Events**	—
Audit events	**System** (⚙) > **Monitoring** > **Audit**	Audit
Rule Update Import Log	Version 7.2.0–7.2.5: System > Updates Version +: System > Content Updates	Rule Updates

When you view any of the kinds of data described in the table above, events appear on the first page of the default workflow for that data. You can specify a different default workflow by configuring your event view settings. Note that workflow access depends on your user role.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Workflow pages

Workflow pages are series of customizable data pages in Firewall Management Center that you can use to evaluate events generated by the system. They share a common set of features across all workflow types despite containing different data and allow you to perform specific actions depending on the page type.

Common workflow page features

Drill-down and table view pages in workflows allow you to quickly narrow your view of the data so you can zero in on events that are significant to your analysis. Table view pages and drill-down pages both support many features you can use to constrain the set of events you want to view or to navigate the workflow. When viewing data on drill-down pages or in the table view in a workflow, you can sort the data in ascending or descending order based on any available column. If the database contains more events than can be displayed on a single workflow page, you can click the links at the bottom of the page to display more events. When you click one of these links, the time window automatically pauses so that you do not see the same events twice; you can unpause the time window when you are ready.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Workflow page types:

- *Table views*: Include a column for each of the fields in the database on which your workflow is based if the page is enabled by default.
- *Drill-down pages*: Generally intermediate pages that you use to narrow your investigation to a few events before moving to a table view page.
- *Graphs*: Workflows based on connection data can include graph pages, also called connection graphs.
- *Final pages*: The final page of a workflow depends on the type of event on which the workflow is based.

Table views

- For best performance, display only the columns you need. The more columns are displayed, the more resources are required to display the data.
- When you disable a column on a table view, the system adds the Count column to the event view if disabling the column could create two or more identical rows, if no more than 6 columns are displayed (excluding the Count column).
- When you click on a value in a table view page, you constrain by that value.
- When you create a custom workflow, you add a table view to it by clicking **Add Table View**.

Drill-down pages

- Contain a subset of columns that are available in the database.
- Allow you to narrow the scope of events you are viewing and to move forward in the workflow.

- If you click on a value in a drill-down page, you constrain by that value and move to the next page in the workflow, focusing more closely on events that match your selected values.
- Clicking a value in a drill-down page does not disable the column where the value is, even if the page you advance to is a table view.
- Drill-down pages for predefined workflows always have a Count column.
- When you create a custom workflow, you add a drill-down page to it by clicking **Add Page**.

For example, a drill-down page for discovery events might include only the IP Address, MAC Address, and Time columns. A drill-down page for intrusion events, on the other hand, might include the Priority, Impact Flag, Inline Result, and Message columns.

Graphs

Workflows based on connection data can include graph pages, also called connection graphs.

For example, a connection graph might display a line graph that shows the number of connections detected by the system over time. Generally, connection graphs are, like drill-down pages, intermediate pages that you use to narrow your investigation.

Final pages

Final page types by workflow event:

- Host view: Final page for workflows based on applications, application details, discovery events, hosts, indications of compromise (IOC), servers, allow list violations, host attributes, or third-party vulnerabilities. Viewing host profiles from this page allows you to easily view data on all IP addresses associated with hosts that have multiple addresses.
- User detail view: Final page for workflows based on users, user activity, and user indications of compromise.
- Vulnerability detail view: Final page for workflows based on Cisco vulnerabilities.
- Packet view: Final page for workflows based on intrusion events.

Workflows based on other kinds of events (for example, audit log events or malware events) do not have final pages.

On the final page of a workflow, you can expand detail sections to view specific information about each object in the set you focused on over the course of the workflow. Although the web interface does not list the constraints on the final page of a workflow, previously set constraints are retained and applied to the set of data.

Workflow page navigation tools

Workflow page navigation tools are visual interface elements that facilitate movement between pages and selection of information to display during event analysis.

Workflow page traversal tools

If a workflow contains multiple pages of data, the bottom of each page displays the number of pages in the workflow, as well as the tools listed in this table which you may use to navigate among the pages:

Table 84: Workflow page traversal tools

Page traversal tool	Action
Page number (To view a different page, enter the number you wish to view, then press **Enter**.)	view a different page
>	view the next page
<	view the previous page
>|	jump to the last page
|<	jump to the first page

File trajectory icons

When a workflow page provides the opportunity to view the trajectory map for a file in a new window, a network trajectory icon appears. This icon differs depending upon the file status.

Table 85: File trajectory icons

File Trajectory Icon	File Status
Clean	Clean
Malware	Malware
Custom detection	Custom detection
Unknown	Unknown
Unavailable	Unavailable

Host profile icons

When a workflow page provides the opportunity to view the host profile associated with an IP address in a pop-up window, a host profile icon appears. If the host profile icon is dimmed, you cannot view the host profile because that host cannot be in the network map (for example, `0.0.0.0`). This icon appears different depending on the status of the host.

Table 86: Host profile icons

Host profile icon	Host status
	Host is not tagged as potentially compromised.
	Host is tagged as potentially compromised by triggered indications of compromise (IOC) rules.

Host profile icon	Host status
	Added to Block List (Appears only if you are performing traffic filtering based on Security Intelligence data.)
	Added to Block List, set to monitor (Appears only if you are performing traffic filtering based on Security Intelligence data.)

Threat score icons

When a workflow page provides the opportunity to view a Dynamic Analysis Summary report for the highest threat score associate with a file, a threat score icon appears. The icon differs depending on the file's highest threat score.

Table 87: Threat score icons

Threat score icon	Threat score level
Low	Low
Medium	Medium
High	High
Very High	Very high

User icons

When a workflow page provides the opportunity to view the user identity associated with a username in a pop-up window, a user icon appears.

Table 88: User icons

User icon	User status
User	User is not associated with any indications of compromise.
Red User	User is associated with one or more indications of compromise.

Workflow toolbar

Each page in a workflow includes a toolbar that offers quick access to related features. This table describes each of the links on the toolbar.

Table 89: Workflow toolbar links

Feature	Description
Bookmark This Page	Bookmarks the current page so you can return to it later. Bookmarking captures the constraints in effect on the page you are viewing so you can return to the same data (assuming the data still exists) at a later time.

Feature	Description
Report Designer	Opens the report designer with the currently constrained workflow as the selection criteria.
Dashboard	Opens a dashboard relevant to your current workflow. For example, Connection Events workflows link to the **Connection Summary** dashboard.
View Bookmarks	Displays a list of saved bookmarks from which you can select.
Search	Displays a Search page where you can perform advanced searches on data in the workflow. You can also click the down arrow icon to select and use a saved search.

Use drill-down pages

Use drill-down pages to navigate between workflow pages while applying specific constraints and filters to focus on relevant data and events. This differs from table views, which only filter the current display without navigating to subsequent pages.

Procedure

Step 1 Access a workflow by choosing the appropriate menu path and option as described in Features Using Workflows.

Step 2 In any workflow, choose from these navigation options:

- To drill down to the next workflow page constraining on a specific value, click a value within a row.
- To drill down to the next workflow page constraining on some events, check the check boxes next to the events you want to view on the next workflow page, then click **View**.
- To drill down to the next workflow page keeping the current constraints, click **View All**.

Note
Clicking a value within a row in a table view only constrains the table view and does not drill down to the next page. Table views always include **Table View** in the page name.

Use table view pages

Table view pages provide some features not available on drill-down, host view, packet view, or vulnerability detail pages.

Procedure

Step 1 Access a workflow by choosing the appropriate menu path. For instructions, refer to Workflow selection, on page 678.

Step 2 Choose a table view from the workflow path displayed beneath the workflow name.

Step 3 If event data is stored remotely, you may see an option to choose whether to display local or remote data. For more information, refer to Work in Firewall Management Center with connection events stored on a Secure Network Analytics appliance, on page 685.

Step 4 Use the features listed below to arrange and navigate within the table view as needed:

- To display the list of disabled columns, click the Search Constraints **Expand Arrow** ( ).
- To hide the list of disabled columns, click the Search Constraints **Collapse Arrow** ( ).
- To add a disabled column back to the event view, click the Search Constraints **Expand Arrow** ( ) to expand the search constraints, then click the column name under Disabled Columns.
- To show or hide (disable) a column, click **Clear** ( ) next to any column name. In the pop-up window that appears, check or clear the appropriate check boxes to indicate which columns you want to display, then click **Apply**.

Work in Firewall Management Center with connection events stored on a Secure Network Analytics appliance

If your devices are sending connection events to a Secure Network Analytics appliance using Security Analytics and Logging (On Premises), you can view and work with these remotely stored events in the Firewall Management Center's event viewer and context explorer, and include them when generating reports. You can also cross-launch from an event in Firewall Management Center to view related data on your Secure Network Analytics appliance.

By default, the system automatically selects the appropriate data source based on the time range you specify. If you want to override the data source, use this procedure.

Note When you change the data source, your selection persists across all of the relevant analytics features that rely on the event data source, including reports, until you change it, even after you sign out. Your selection does not apply to other Firewall Management Center users.

The selected data source is used for low-priority connection events only. All other event types (intrusion, file, and malware events; connection events associated with those events; and Security Intelligence events) are displayed regardless of data source.

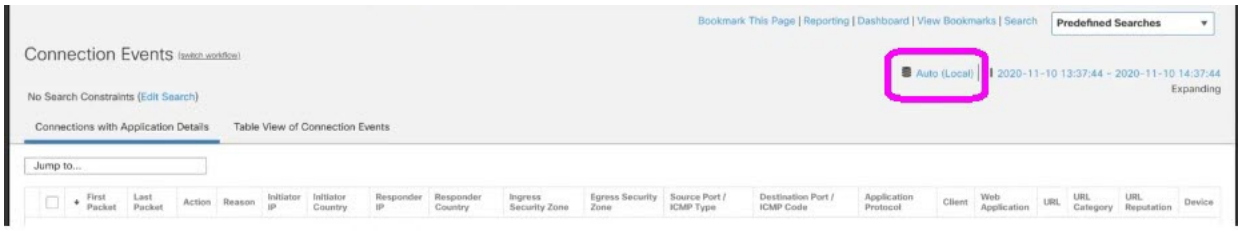
Before you begin

You have used the wizard to send connection events to Security Analytics and Logging (On Premises).

Procedure

Step 1 In the Firewall Management Center web interface, navigate to a page that displays connection event data, such as **Analysis** > **Connections** > **Events**.

Step 2 Click the data source displayed here and select an option:



Caution
If you select **Local**, the system displays only the data available on the Firewall Management Center, even if local data is not available for the entire time range selected. You will not be notified that this situation is occurring.

Step 3 (Optional) To view related data directly in your Secure Network Analytics appliance, right-click (in the unified event viewer, click) a value such as an IP address or domain and choose a cross-launch option.

Geolocation

Geolocation is a database feature that maps IP addresses to geographic locations and enables viewing and filtering traffic based on country and continent by leveraging a geolocation database (GeoDB).

The system may report a continent instead of a specific country for mobile devices and other hosts detected moving from country to country. The system issues periodic updates to the GeoDB. You must regularly update the GeoDB to have accurate geolocation information. For instructions, refer to Geolocation database (GeoDB) updates, on page 258.

Note We no longer provide the geolocation IP package, which contained contextual data associated with routable IP addresses. This saves disk space and does not affect geolocation rules or traffic handling in any way. Any contextual data is now stale, and upgrading to most later versions deletes the IP package. Options to download the IP package or view contextual data have no effect, and are removed in later versions.

Connection event graphs

A connection event graph is a graphical presentation tool that

- displays connection data using data aggregated over five-minute intervals
- presents information that is used for aggregating data, including source and destination IP addresses, destination port, transport protocol, and application protocol, and

- provides an alternative to tabular drill-down pages and final table views of events.

Connection graph types

There are three different types of connection graphs:

- Pie charts display data from one dataset grouped into discrete categories.
- Bar graphs display data from one or more datasets grouped into discrete categories.
- Line graphs plot data from one or more datasets over time, using either a standard or a velocity (rate of change) view.

Like workflow tables, you can drill down and constrain workflow graphs to focus your analysis.

Both bar graphs and line graphs can display multiple datasets; that is, they can display several values on the y-axis for each x-axis data point. For example, you could display the total number of unique initiators and responders. Pie charts can only display one dataset.

You can display different data on a connection graph by changing the x-axis, the y-axis, or both. On a pie chart, changing the x-axis changes the independent variable and changing the y-axis changes the dependent variable.



Note

- You cannot graph Security Intelligence events separately from their associated connection events. For a graphical overview of Security Intelligence event filtering activity, use dashboards and the **Context Explorer**.
- Firewall Management Center displays traffic profiles as line graphs, which you can manipulate in the same way as you would any other connection graph, with some restrictions. To view traffic profiles, you must have **Administrator** access.

Related Topics

Connection summaries (aggregated data for graphs), on page 756

Use connection event graphs

On the Firewall Management Center, you can view connection event graphs and manipulate them depending on the information you are looking for.

The page you see when you access connection graphs differs depending on the workflow you use. You can use a predefined workflow, which terminates in a table view of connection events. You can also create a custom workflow that displays only the information that matches your specific needs.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Connections** > **Events**.

Note

If a connection event table appears instead of a graph, or to view a different graph, click **(switch workflow)** by the workflow title and choose a predefined workflow that includes graphs, or a custom workflow. Note that all predefined connection event workflows—including connection graphs—terminate in a table view of connections.

Step 2 You have these options to work with the graph:

- Time Range — To adjust the time range, which is useful if the graph is blank, see Change the time window, on page 695.
- Field Names — To learn more about the data you can graph, see Connection and Security Intelligence Event Fields, on page 758.
- Host Profile — To view the host profile for an IP address, on a graph displaying connection data by initiator or responder, click either a bar on a bar graph or a wedge on a pie chart and choose **View Host Profile**.
- User Profile — To view user profile information, on a graph displaying connection data by initiator user, click either a bar on a bar graph or a wedge on a pie chart and choose **View User Profile**.
- Other Information — To learn more information about the graphed data, position your cursor over a point on a line graph, a bar in a bar graph, or a wedge in a pie chart.
- Detach — To detach a connection graph so you can perform further analysis without affecting the default time range, click **Detach**.

  Click **New Window** in a detached graph to create a copy. You can then perform different analyses on each of the detached graphs. Note that traffic profiles are detached graphs.
- Export — To export the connection data for a graph as a CSV (comma-separated values) file, **Export Data**. Then, click **Download CSV File** and save the file.

You have these options to customize and filter the data:

- Constrain — To constrain a connection graph by any x-axis (independent variable) criterion without advancing the workflow to the next page, click a point on a line graph, a bar on a bar graph, or a wedge on a pie chart, and choose a **View by...** option.
- Data Selection — To change the data displayed on the graph, click **X-Axis** or **Y-Axis** and choose the new data to graph. Note that changing the x-axis to or from **Time** also changes the graph type; changing the y-axis affects the displayed datasets.
- Datasets — To change the graph's dataset, click **Datasets** and choose a new dataset.
- Drill Down — To drill down to the next page in the workflow, click a point on a line graph, a bar on a bar graph, or a wedge on a pie chart, then choose **Drill-down**. Clicking a point on a line graph changes the time range on the next page to a 10-minute span, centered on the point you clicked. Clicking a bar on a bar graph or a wedge on a pie chart constrains the next page based on the criterion represented by the bar or wedge.

You have these graph types:

- Line — To switch between a standard and velocity (rate of change) line graph, click **Velocity**, then choose **Standard** or **Velocity**.
- Bar and Pie — To switch between a bar graph and pie chart, click **Switch to Bar** or **Switch to Pie**. Because you cannot display multiple datasets on a pie char, if you switch to a pie chart from a bar graph

that has multiple datasets, the pie chart shows only one dataset, which is selected automatically. When choosing which dataset to display, the Firewall Management Center favors total statistics over initiator and responder statistics, and favors initiator statistics over responder statistics.

You have these naviagtion options:

- To navigate between pages in the current workflow, keeping the current constraints, click the appropriate page link at the top left of the workflow page.
- To navigate to other event views to view associated events, click **Jump to** and choose the event view from the drop-down list.
- To recenter a line graph around a point in time without changing the length of the time range, click that point, then choose **Recenter**.
- To recenter a line graph around a point in time while zooming in or out, click that point, choose **Zoom**, then choose a new time span.

Note
Unless you are working with a detached graph, constraining, recentering, and zooming changes the default time range for the Firewall Management Center.

You can now view and manipulate connection event graphs to analyze your network data based on various criteria such as time ranges, ports, IP addresses, and application protocols.

Example: Constraining a Connection Graph

Consider a graph of connections over time. If you constrain a point on the graph by port, a bar graph appears, showing the 10 most active ports based on the number of detected connection events, but constrained by the ten-minute time span that is centered on the point you clicked.

If you further constrain the graph by clicking on one of the bars and choosing **View by Initiator IP**, a new bar graph appears, constrained by not only the same ten-minute time span as before, but also by the port represented by the bar you clicked.

Example: Changing X-Axis and Y-Axis on a Pie Chart

Consider a pie chart that graphs kilobytes per port. In this case, the x-axis is **Responder Port** and the y-axis is **KBytes**. This pie chart represents the total kilobytes of data transmitted over a monitored network during a certain interval. The wedges of the pie represent the percent of the data that was detected on each port.

- If you change the x-axis of the chart to **Application Protocol**, the pie chart still represents the total kilobytes of data transmitted, but the wedges of the pie represent the percentage of the data transmitted for each detected application protocol.
- If you change the y-axis of the chart to **Packets**, the pie chart represents the total number of packets transmitted over the monitored network during a certain interval, and the wedges of the pie represent the percentage of the total number of packets that was detected on each port.

Related Topics

Use workflows, on page 676
Configure event view settings, on page 234

Connection graph data options

You can display different data on a connection graph by changing either the x-axis, the y-axis, or both. On a pie chart, changing the x-axis changes the independent variable and changing the y-axis changes the dependent variable.

Table 90: X-Axis options

X-Axis Option	Graph Type	Graphs This Data
Application Protocol	bar or pie	by the 10 most active application protocols
Device	bar or pie	by the 10 most active managed devices
Initiator IP	bar or pie	by the 10 most active initiator host IP addresses
Initiator User	bar or pie	by the 10 most active initiator users
Responder IP	bar or pie	by the 10 most active responder host IP addresses
Responder Port	bar or pie	by the 10 most active responder ports
Source Device	bar or pie	by the 10 most active NetFlow data exporters, plus a source device named **Firepower** for all connections detected by Secure Firewall System managed devices.
Time	line	over time Changing the y-axis to and from **Time** also changes the graph type and may change the datasets.

Table 91: Y-Axis options

Y-Axis Option	Graphs This Data Using The X-Axis Criterion
Bytes	bytes transmitted
Connections	number of connections
KBytes	kilobytes transmitted
KBytes Per Second	kilobytes per second
Packets	number of packets transmitted
Unique Hosts	number of unique hosts detected
Unique Application Protocols	number of unique application protocols
Unique Users	number of unique users

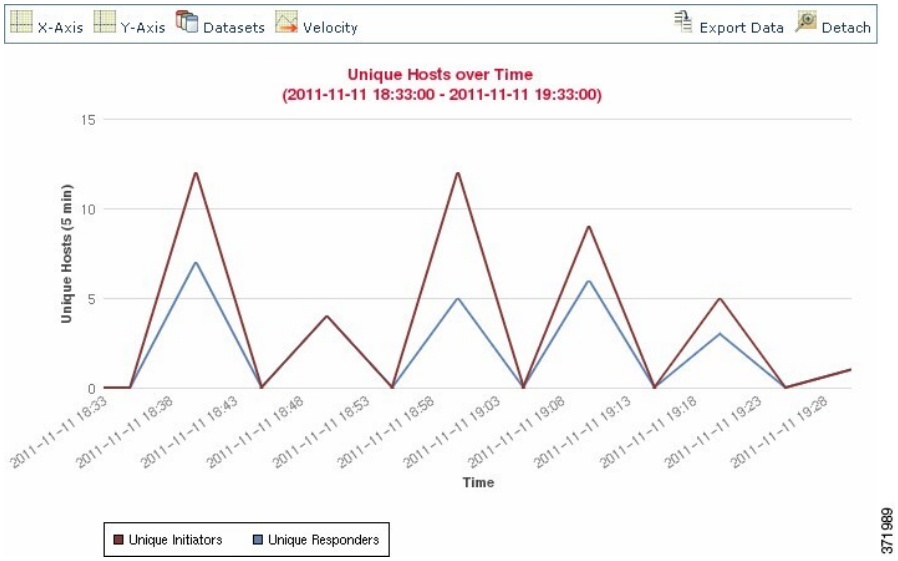
Connection graphs with multiple datasets

Connection graph with multiple datasets is an option that displays several values on the y-axis for each x-axis data point. This option is supported by both bar graphs and line graphs but not pie charts.

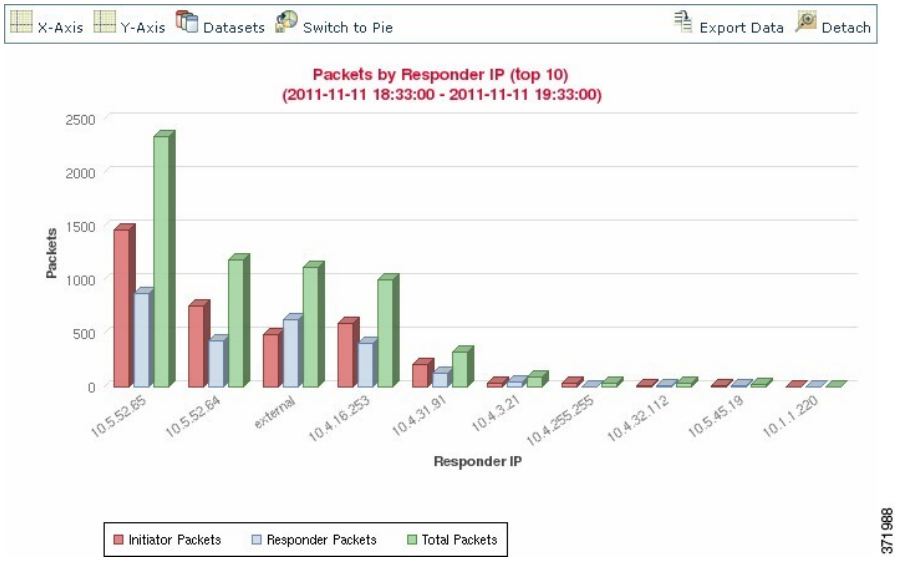
Viewing multiple dataset in connection graphs

Multiple datasets are displayed differently depending on the graph type:

- Line graphs: Multiple datasets appear as separate lines, each shown in a different color. For example, a line graph can display the total number of unique initiators and responders detected on a monitored network over a one-hour interval.



- Bar graphs: Multiple datasets appear as sets of colored bars for each x-axis data point. For example, a bar graph can display total packets transmitted, packets transmitted by initiators, and packets transmitted by responders.



You cannot display multiple datasets on a pie chart. If you switch to a pie chart from a bar graph that has multiple datasets, the pie chart shows only one dataset, which is selected automatically. When selecting which dataset to display, the Firewall Management Center favors total statistics over initiator and responder statistics, and favors initiator statistics over responder statistics.

Connection graph dataset options

This table describes the datasets you can display on the x-axis of a connection graph.

Table 92: Dataset options

If the y-axis displays...	You can select as datasets...
Connections	the default only, which is the number of connections detected on the monitored network (**Connections**). This is the only option for traffic profile graphs.
KBytes	combinations of: • the total kilobytes transmitted on the monitored network (**Total KBytes**) • the number of kilobytes transmitted from host IP addresses on the monitored network (**Initiator KBytes**) • the number of kilobytes received by host IP addresses on the monitored network (**Responder KBytes**)
KBytes Per Second	the default only, which is the total kilobytes per second transmitted on the monitored network (**Total KBytes Per Second**)
Packets	combinations of: • the total packets transmitted on the monitored network (**Total Packets**) • the number of packets transmitted from host IP addresses on the monitored network (**Initiator Packets**) • the number of packets received by host IP addresses on the monitored network (**Responder Packets**)
Unique Hosts	combinations of: • the number of unique session initiators on the monitored network (**Unique Initiators**) • the number of unique session responders on the monitored network (**Unique Responders**)
Unique Application Protocols	the default only, which is the number of unique application protocols on the monitored network (**Unique Application Protocols**)
Unique Users	the default only, which is the number of unique users logged into session initiators on the monitored network (**Unique Initiator Users**)

Event time constraints

Event time constraints are workflow filtering feature that use time stamps to indicate when events occurred and allow you to constrain information by setting time windows or time ranges.

Workflows based on events that can be constrained by time include a time range line at the top of the page.

Time window behavior

By default, workflows use an expanding time window set to the past hour. For example, if you log in at 11:30 AM, you will see events that occurred between 10:30 AM and 11:30 AM. As time moves forward, the time window expands. At 12:30 PM, you will see events that occurred between 10:30 AM and 12:30 PM.

You can change this behavior by setting your own default time window in the event view settings. This governs three properties:

- time window type (static, expanding, or sliding)
- time window length
- the number of time windows (either multiple time windows or a single global time window)

Regardless of the default time window setting, you can manually change the time window during your event analysis by clicking the time range at the top of the page. Depending on the number of time windows you configured and the type of appliance you are using, you can also use the date and time window to change the default time window for the type of event you are viewing.

Finally, you can pause the time window while looking at a sliding or expanding workflow. For more information, refer to Pause the time window to temporarily freeze the data set, on page 696.

Per-session time window customization for events

Per-session time window customization for events is a manual configuration capability that

- allows users to change the time window during event analysis sessions
- applies settings only to the current session, resetting to default when users log out and log back in, and
- affects different workflows based on the number of configured time windows on the appliance.

Time window configuration and workflow effects

Changing the time window for one workflow may affect other workflows on the appliance based on your time window configuration:

- Single, global time window: Changing the time window for one workflow changes it for all other workflows on the appliance.
- Multiple time windows: Changing the audit log or health event workflow time windows has no effect on any other time window, while changing the time window for other kinds of events affects all events that can be constrained by time (with the exception of audit events and health events).

Note Manual time window settings are valid for only the current session. When you log out and then log back in, time windows are reset to the default.

Time window settings have no effect on workflows based on hosts, host attributes, applications, application details, vulnerabilities, users, or allow list violations because not all workflows can be constrained by time.

Use the time window tab on the date and time window to manually configure a time window. Depending on the number of time windows you configured in your default time window settings, the tab's title is one of these:

- **Events Time Window**, if you configured multiple time windows and are setting the time window for a workflow other than the audit log or health events workflow
- **Health Monitoring Time Window**, if you configured multiple time windows and are setting the time window for the health events workflow
- **Audit Log Time Window**, if you configured multiple time windows and are setting the time window for the audit log
- **Global Time Window**, if you configured a single time window

The first decision you must make when configuring a time window is the type of time window you want to use:

- A static time window displays all the events generated from a specific start time to a specific end time.
- An expanding time window displays all the events generated from a specific start time to the present; as time moves forward, the time window expands and new events are added to the event view.
- A sliding time window displays all the events generated from a specific start time (for example, one week ago) to the present; when you refresh the page, the time window slides so that you see only the events in the time range you configured (in this example, for the last week). To temporarily prevent the data set from updating while you are examining it, see Pause the time window to temporarily freeze the data set, on page 696.

Depending on what type you select, the Date/Time window changes to give you different configuration options.

Note The system uses a 24-hour clock based on the time you specified in your time zone preferences.

Time window settings

This table explains the various settings you can configure on the Time Window tab.

Table 93: Time window settings

Setting	Time window type	Description
time window type drop-down list	n/a	Select the type of time window you want to use: static, expanding, or sliding. Note that events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.
Start Time calendar	static and expanding	Specify a start date and time for your time window. The maximum time range for all time windows is from midnight on January 1, 1970 (UTC) to 3:14:07 AM on January 19, 2038 (UTC). Instead of using the calendar, you can use the Presets options, described below.

Setting	Time window type	Description
End Time calendar	static	Specify an end date and time for your time window. The maximum time range for all time windows is from midnight on January 1, 1970 (UTC) to 3:14:07 AM on January 19, 2038 (UTC). Note that if you are using an expanding time window, the End Time calendar is grayed out and specifies that the end time is "Now." Instead of using the calendar, you can use the Presets options, described below.
Show the Last field and drop-down list	sliding	Configure the length of the sliding time window.
Presets: Last	all	Click one of the time ranges in the list to change the time window, based on the local time of the appliance. For example, clicking **1 week** changes the time window to reflect the last week. Clicking a preset changes the calendars to reflect the preset you choose.
Presets: Current	static and expanding	Click one of the time ranges in the list to change the time window, based on the local time and date of the appliance. Clicking a preset changes the calendars to reflect the preset you choose. Note that: • the current day begins at midnight • the current week begins at midnight Sunday • the current month begins at midnight on the first of the month
Presets: Synchronize with	all (not available if you are using a global time window)	Click one of: • **Events Time Window** to synchronize the current time window with the events time window • **Health Monitoring Time Window** to synchronize the current time window with the health monitoring time window • **Audit Log Time Window** to synchronize the current time window with the audit log time window

Change the time window

Change the time window settings for workflows that are constrained by specific time parameters.

Procedure

Step 1 On a workflow constrained by time, click **Time Range** ( ) to go to the date and time window.

Step 2 On **Events Time Window**, set the time window. For instructions, refer to Time window settings, on page 694.

To change the time window back to the default settings, click **Reset**.

Step 3 Click **Apply**.

Pause the time window to temporarily freeze the data set

Pausing the time window allows you to examine a snapshot of the data provided by the workflow without interference from automatic updates.

When an unpaused workflow updates, it may remove events that you want to examine or add events that you are not interested in.

The time window automatically pauses when you click a link at the bottom of the page to display another page of events. When you are finished with your analysis, you can unpause the time window. Unpausing the time window updates it according to your preferences, and also updates the event view to reflect the unpaused time window.

Pausing an event time window has no effect on dashboards, nor does pausing a dashboard have any effect on pausing an event time window.

Procedure

Step 1 Access a workflow that is constrained by time. For instructions, refer to Workflow selection, on page 678.

Step 2 Choose the desired time range control:

- To pause the time window, click time range control **Pause** ( ).
- To unpause the time window, click time range control **Play** ( ).

Default time window for events

A default time window for events is a configurable time range setting that determines the initial time period displayed for event analysis.

During your event analysis, you can use the **Preferences** tab on the date and time window to change the default time window for the type of event you are viewing.

Default time window configuration

The scope of default time window changes depends on your configuration:

- Multiple time windows: Changing the default time window on the **Preferences** tab changes the settings for only the type of event you are viewing (events, health monitoring, or audit log window), specifically whichever time window is indicated by the first tab.
- Single time window: Changing the default time window on the **Preferences** tab changes the default time window for all types of events.

Default time window options for event types

This table explains the various settings you can configure on the **Preferences** tab.

Table 94: Time window preferences

Preference	Description
Refresh Interval	Sets the refresh interval for event views, in minutes. Entering zero disables the refresh option.
Number of Time Windows	Specify how many time windows you want to use: • Select **Multiple** to configure separate default time windows for the audit log, for health events, and for workflows based on events that can be constrained by time. • Select **Single** to use a global time window that applies to all events.
Default Time Window: Show the Last - Sliding	This setting allows you to configure a sliding default time window of the length you specify. The appliance displays all the events generated from a specific start time (for example, 1 hour ago) to the present. As you change event views, the time window "slides" so that you always see events from the last hour.
Default Time Window: Show the Last - Static/Expanding	This setting allows you to configure either a static or expanding default time window of the length you specify. For **static** time windows (enable the **Use End Time** check box), the appliance displays all the events generated from a specific start time (for example, 1 hour ago), to the time when you first viewed the events. As you change event views, the time window stays fixed so that you see only the events that occurred during the static time window. For **expanding** time windows (disable the **Use End Time** check box), the appliance displays all the events generated from a specific start time (for example, 1 hour ago), to the present. As you change event views, the time window expands to the present time.
Default Time Window: Current Day - Static/Expanding	This setting allows you to configure either a static or expanding default time window for the current day. The current day begins at midnight, based on the time zone setting for your current session. For **static** time windows (enable the **Use End Time** check box), the appliance displays all the events generated from midnight to the time when you first viewed the events. As you change event views, the time window stays fixed so that you see only the events that occurred during the static time window. For **expanding** time windows (disable the **Use End Time** check box), the appliance displays all the events generated from midnight to the present. As you change event views, the time window expands to the present time. Note that if your analysis continues for over 24 hours before you log out, this time window can be more than 24 hours.

Preference	Description
Default Time Window: Current Week - Static/Expanding	This setting allows you to configure either a static or expanding default time window for the current week. The current week begins at midnight on the previous Sunday, based on the time zone setting for your current session. For **static** time windows (enable the **Use End Time** check box), the appliance displays all the events generated from midnight to the time when you first viewed the events. As you change event views, the time window stays fixed so that you see only the events that occurred during the static time window. For **expanding** time windows (disable the **Use End Time** check box), the appliance displays all the events generated from midnight Sunday to the present. As you change event views, the time window expands to the present time. Note that if your analysis continues for over 1 week before you log out, this time window can be more than 1 week.

Change the default time window for your event type

Change the default time window settings to customize how time-constrained workflows display events within your preferred time range.

Procedure

Step 1 On a workflow constrained by time, click **Time Range** ( ) to go to the date and time window.

Step 2 Click **Preferences** and change your preferences. For instructions, refer to Default time window options for event types, on page 697.

Step 3 Click **Save Preferences**.

Step 4 Choose one of the following options:

- To apply your new default time window settings to the event view you are using, click **Apply** to close the date and time window and refresh the event view.
- To continue with your analysis without applying the default time window settings, close the date and time window without clicking **Apply**.

Event view constraints

Event view constraints are filtering options in workflow pages that determine the information displayed on workflow pages based on imposed constraints. When you initially open an event workflow, the information is constrained to events that were generated in the previous hour.

- allow users to refine data visibility through multiple constraint methods.

Constraint methods

Use these methods for constraining data in event workflows:

- To advance to the next page in the workflow and constrain the data you are viewing by specific values, select the rows with those values on the page and click **View**.

Note If you select a row with multiple non-count values and click **View**, you create a compound constraint.

- To advance to the next page in the workflow retaining the current constraints and carrying forward all events, select **View All**.
- To constrain the page to the rows with values that you selected and also add the selected value to the list of constraints at the top of the page, click a value within a row on the page.

If you are viewing a list of logged connections and want to constrain the list to only those you allowed using access control, click **Allow** in the **Action** column.

As another example, if you are viewing intrusion events and want to constrain the list to only events where the destination port is 80, click **80 (http)/tcp** in the **Destination Port/ICMP Code** column.

Note The procedure for constraining connection events based on Monitor rule criteria is slightly different and you may need to take some extra steps. Additionally, you cannot constrain connection events by associated file or intrusion information.

Use searches to contrain data

You can also use searches to constrain the information in a workflow. Use this feature when you want to constrain against multiple values in a single column. For example, if you want to view the events related to two IP addresses, click **Edit Search**, then modify the appropriate IP address field on the Search page to include both addresses, and then click **Search**.

The search criteria you enter on the search page are listed as the constraints at the top of the page, with the resulting events constrained accordingly. On the Firewall Management Center, the current constraints are also applied when navigating to other workflows, unless they are compound constraints.

When searching, you must pay careful attention to whether your search constraints apply to the table you are searching. For example, client data is not available in connection summaries. If you search for connection events based on the detected client in the connection and then view the results in a connection summary event view, the Firewall Management Center displays connection data as if you had not constrained it at all. Invalid constraints are labeled as not applicable (N/A) and are marked with a strikethrough.

Constrain events

Constraining events help you narrow down large datasets in workflows by applying filters based on specific values or combinations of values. You can save these constraints for reuse and apply them across different event views.

Procedure

Step 1 Access a workflow by choosing the appropriate menu path. For instructions, refer to Workflow selection, on page 678.

Step 2 Apply event constraints using these options:

- To constrain the view to events that match a single value, click the desired value within a row on the page.
- To constrain the view to events that match multiple values, check the check boxes for events with those values, and click **View**.

  Note
  A compound constraint is added if the row contains multiple non-count values.

- To remove a constraint, click the Search Constraints **Expand Arrow** ( ) and click the name of the constraint in the expanded Search Constraints list.
- To edit constraints using the Search page, click **Edit Search.**
- To save constraints as a saved search, click **Save Search** and give the query a name.

  Note
  You cannot save queries containing compound constraints.

- To use the same constraints with another event view, click **Jump to** and choose the event view.

  Note
  You do not retain compound constraints when you switch to another workflow.

- To toggle the display of constraints click the Search Constraints **Expand Arrow** ( ) or the Search Constraints **Collapse Arrow** ( ). This is useful when the list of constraints is large and takes up most of the screen.

Compound event view constraints

A compound event view constraint is a filtering option that retrieves events based on all non-count values for a specific event. When you select a row with multiple non-count values, this option allows you to set constraints that match all the non-count values in that row on that page.

How compound constraints work

When you select a row that has a source IP address of `10.10.31.17` and a destination IP address of `10.10.31.15` and a row that has a source IP address of `172.10.10.17` and a destination IP address of `172.10.10.15`, you retrieve all of this:

- Events that have a source IP address of 10.10.31.17 and a destination IP address of 10.10.31.15

OR

- Events that have a source IP address of 172.10.31.17 and a destination IP address of 172.10.31.15

When you combine compound constraints with simple constraints, the simple constraints are distributed across each set of compound constraints. If, for example, you added a simple constraint for a protocol value of `tcp` to the compound constraints listed above, you retrieve all of this:

- Events that have a source IP address of 10.10.31.17 and a destination IP address of 10.10.31.15 and a protocol of `tcp`, or
- events that have a source IP address of 172.10.31.17 and a destination IP address of 172.10.31.15 and a protocol of `tcp`.

You cannot perform a search or save a search on a compound constraint. You also cannot retain compound constraints when you use the event view links or click **(switch workflow)** to switch to another workflow. If you bookmark an event view with compound constraints applied, the constraints are not saved with the bookmark.

Use compound event view constraints

Compound constraints allow you to filter workflow data by selecting multiple rows with non-count values, providing more refined data analysis capabilities.

Procedure

Step 1 Access a workflow by choosing the appropriate menu path. For instructions, refer to Workflow selection, on page 678.

Step 2 To manage compound constraints, you have these options:

- To create a compound constraint, choose one or more rows with multiple non-count values and click **View**.
- To clear compound constraints, click the Search Constraints **Expand Arrow** ( ) and click **Compound Constraints**.

Navigate between workflows

The **Jump to...** drop-down list on a workflow page is a feature that enables you to navigate between different workflows. It transfers applicable shared properties and constraints from selected rows to the new workflow, and enhances investigation capabilities by allowing correlation of data across different workflows.

Navigation behavior and constraints

When you select a new workflow, properties shared by the rows you select and the constraints you set are used in the new workflow, if they are applicable. If configured constraints or event properties do not map to fields in the new workflow, they are dropped. In addition, compound constraints are not retained when you

switch from one workflow to another. In addition, constraints from the captured files workflow only transfer to file and malware event workflows.

Note

- When you view event counts over a time range, the total number of events may not reflect the number of events for which more detailed data is available. This occurs because the system sometimes prunes older event details to manage disk space usage. To minimize the occurrence of event detail pruning, you can fine-tune event logging to log only those events most important to your deployment.
- Unless you have either paused the time window or have configured a static time window, the time window changes when you change workflows.

Security investigation workflow

This feature enhances your ability to investigate suspicious activity. For example, if you are viewing connection data and notice that an internal host is transmitting an abnormally large amount of data to an external site, you can select the responder IP address and the port as constraints and then jump to the **Applications** workflow. The applications workflow will use the responder IP address and port as IP Address and Port constraints and display additional information about the application, such as what kind of application it is. You can also click **Hosts** at the top of the page to view the host profile for the remote host.

After finding more information about the application, you can select **Correlation Events** to return to the connection data workflow, remove the Responder IP from the constraints, add the Initiator IP to constraints, and select **Application Details** to see what client the user on the initiating host used when transferring data to the remote host. Note that the Port constraint is not transferred to the Application Details page. While keeping the local host as a constraint, you can also use other navigation buttons to find additional information:

- To discover if any policies have been violated by the local host, keep the IP address as a constraint and select **Correlation Events** from the **Jump to** drop-down list.
- To find out if an intrusion rule triggered against the host, indicating a compromise, select **Intrusion Events** from the **Jump to** drop-down list.
- To view the host profile for the local host and determine if the host is susceptible to any vulnerabilities that may have been exploited, select **Hosts** from the **Jump to** drop-down list.

Work with the unified event viewer

Unified Events provide you a single-screen view of multiple types (connection, intrusion, file, malware, and some security-related connection events) of firewall events. The Unified Events table is highly customizable. You can create and apply custom filters to fine-tune the information displayed on the event viewer. The **Live View** option in the unified events table lets you see the firewall events in real time and monitor the activity on your network.

Procedure

Step 1 Choose **Analysis** > **Unified Events**.

Step 2 Select a time range (fixed or sliding) to view the firewall events from a specific period. By default, unified event viewer table displays events from the previous hour. You can filter the table to get more granular context of the security event, customize the table columns, or enable live view and see the event updates in real time.

Bookmarks

The bookmarks option is a workflow navigation feature that

- enables you to save a specific location and time in an event analysis, and
- allows quick return to the saved point for further investigation.

You can share a bookmark with other users who have bookmark access. This is useful when you identify events that require detailed analysis and need to turn over the investigation to another user with the appropriate privileges.

Bookmark information retention

Bookmarks retain information about:

- the workflow you are using
- the part of the workflow you are viewing
- the page number within the workflow
- any search constraints
- any disabled columns
- the time range you are using

Note If the events saved under a bookmark are deleted, either by a user or due to automatic database cleanup, the bookmark does not display the original set of events.

Create bookmarks

Create bookmarks to save specific event analysis views and configurations for quick access in future sessions.

In a multidomain deployment, you can only view bookmarks created in the current domain.

Procedure

Step 1 During an event analysis, with the events of interest displayed, click **Bookmark This Page**.

Step 2 In the **Bookmark Name** field, enter a name.

Step 3 Click **Save Bookmark**.

View bookmarks

View bookmarks to quickly access previously saved event views and filter configurations.

In a multidomain deployment, you can only view bookmarks created in the current domain.

Procedure

Step 1 Access a workflow page. For instructions, refer to Workflow selection, on page 678.

Step 2 From any event view, you have two options:

- Hover your pointer over **View Bookmarks**, and click on the desired bookmark in the drop-down menu.
- Click on click **View Bookmarks** and on the View Bookmarks page, click on the desired bookmark name or **View** ( ) next to it.

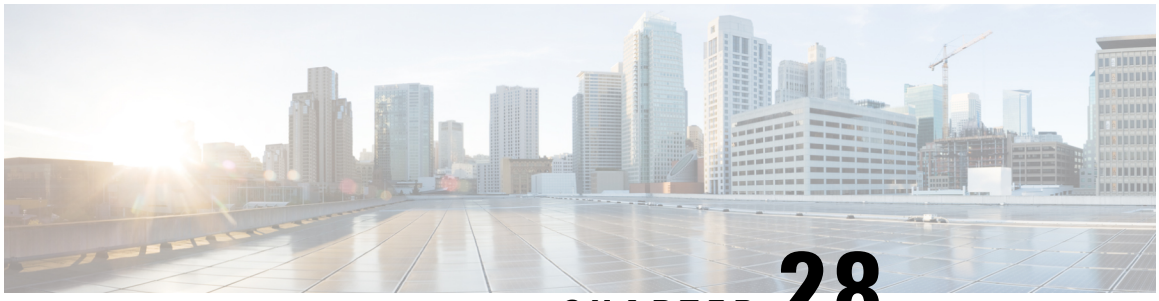
Note
If the events that originally appeared in a bookmark are deleted (either directly by a user or by automatic database cleanup), the bookmark no longer displays the original set of events.

History for workflows

This table provides a comprehensive history of workflow-related features, including deprecated functionality, new capabilities, and improvements introduced across Firewall Management Center software versions.

Table 95: Workflow feature history

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Deprecated: intrusion incidents and event clipboard	7.1	Any	Intrusion incidents and event clipboard are deprecated. Deprecated screens: • **Analysis** > **Intrusions** > **Clipboard** • **Analysis** > **Intrusions** > **Incidents**
Unified event viewer.	7.0	Any	View and work in a single table with multiple event types: connection (including security-rekated connection), intrusion, file, and malware. New/modified screens: **Analysis** > **Unified Events**
Work with events stored remotely.	7.0	Any	Use Firewall Management Center to work with connection events stored on a Secure Network Analytics appliance. The system automatically uses the most appropriate data source, or you can explicitly choose the source. This option appears only if you have completed the Security Analytics and Logging (On Premises) wizard. New/modified screens: pages that display connection events, for example, the event viewer, dashboard, context explorer, and reports.
Improved loading speed of workflow tables in certain cases.	6.6	Any	Tables on workflow pages now show a **Count** column for rows that are identical only when no more than six columns are displayed. This minimizes the amount of calculation required and thus improves table loading speed. New/modified screens: Any events table.



CHAPTER 28

Event Search

The following topics describe how to search for events within a workflow:

- Event searches, on page 707
- Query overrides via the shell, on page 715
- History for search for events, on page 716

Event searches

An event search is a customizable query that

- retrieves specific information stored as events in database tables based on defined search criteria,
- can be saved with custom names and shared with other users or kept private for individual use, and
- returns only records that match the specified criteria for all fields in the search parameters.

Event search functionality

The system generates information that is stored as events in database tables. Events contain multiple fields that describe the activity that caused the appliance to generate the event. You can create and save searches customized for your environment for any of the different event types and save them to reuse later.

When you save a search you give it a name and specify whether the search will be available to you alone or to all users of the appliance. If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search. If you previously saved a search, you can load it, make any necessary modifications, and then start the search. Custom analysis dashboard widgets, report templates, and custom roles can also use saved searches. If you have saved searches, you can delete them from the Search page.

For some event types, the system provides predefined searches that serve as examples and can provide quick access to important information about your network. You can modify fields within the predefined searches for your network environment, then save the searches to reuse later.

The search criteria you can use depends on the type of search, but the mechanics are the same. Searches return only records that match the search criteria specified for all fields.

Note Searching a custom table requires a slightly different procedure.

Related Topics

Search custom tables, on page 731

Search constraints

A search constraint is a event data search feature that enables you to enter search criteria values to apply to fields defined for event tables and constrain the data displayed to your specific needs.

Each database table has its own search page where you can enter search constraint values to apply to fields defined for the table. Depending on the type of field, special syntax may be used to specify criteria such as wildcard characters or a range of numeric values.

Search results appear on workflow pages displaying each table field in columnar layout. Some database tables can additionally be searched using fields that are not displayed as columns on workflow pages. To determine whether such a constraint applies to your search results when viewing the results on a workflow page, click

Expand Arrow ( ) to view the active search constraints.

Best practice for general search constraints

When searching for events, consider these general guidelines:

- Many fields require wildcards for partial-match searches. All fields accept wildcards for these searches. For more information, refer to Wildcards and symbols in searches, on page 709.
- All fields accept negation (`!`).
- All fields accept comma-separated lists of search values. Records that contain any of the listed values in the specified field match that search criteria.
- All fields accept comma-separated lists enclosed in quotation marks as search values.
  - For fields that may contain only a single value, records with the specified field containing the exact string specified within the quotation marks match the search criteria. For instance, a search for `A, B, "C, D, E"` will match records where the specified field contains `"A"` or `"B"` or `"C, D, E"`. This permits matching on fields that include the comma in possible values.
  - For fields that may contain multiple values at the same time, records with the specified fields containing all of the values in the quote-enclosed comma-separated list match that search criteria.
  - For fields that may contain multiple values at the same time, search criteria may include single values as well as quote-enclosed comma-separated lists. For instance, a search for `A, B, "C, D, E"` on a field that may contain one of more of these letters matches records where the specified field contains `A` or `B`, or all of `C`, `D`, and `E`.
- Specify `n/a` in any field to identify events where information is not available for that field; use `!n/a` to identify the events where that field is populated.
- You can precede many numeric fields with greater than (`>`), greater than or equal to (`>=`), less than (**<**), less than or equal to (`<=`), equal to (`=`), or not equal to (`<>`) operators.

Note When you search a field with long complicated values (such as SHA-256 hash values), copy the search criteria value from source material and paste it into the appropriate field on the search page.

Wildcards and symbols in searches

When searching in all text fields in connection and Security Intelligence events and in most text fields in other event types, searches for partial matches in text fields require an asterisk (*) to represent unspecified characters in a string. Searches without an asterisk are exact-match searches in these fields. Even in fields that do not require wildcards, we recommend always using wildcards for partial-match searches.

For example, to find `example.com`, `www.example.com`, or `department.example.com`, search for `*.example.com`. Searching for `example.com` in most cases returns only `example.com`.

Search for non-alphanumeric characters

If you want to search for non-alphanumeric characters (including the asterisk character), enclose the search string in quotation marks. For example, to search for the string:

```
Find an asterisk (*)
```

enter:

```
"Find an asterisk (*)"
```

Objects and application filters in searches

The system allows you to create named objects, object groups, and application filters that can be used as part of your network configuration. You can use these objects, groups, and filters as search criteria when performing or saving searches.

When you perform a search, objects, object groups, and application filters appear in the format, `${object_name}`. For example, a network object with the object name `ten_ten_network` appears as `${ten_ten_network}` in a search.

Click the **Object** (+) icon that appears next to a search field to use an object as a search criterion.

Related Topics

Use the object manager

Time constraints in searches

The formats accepted by search criteria fields that take a time value are shown in the tables that follow.

Table 96: Time specification in search fields

Time formats	Example
`today [at HH:MMam|PM]`	`today` `today at 12:45pm`
`YYYY-DDMM-HH:MM:SS`	`2006-03-22 14:22:59`

You can precede a time value with one of the operators shown in this table:

Table 97: Time specification operators

Operator	Example	Explanation
<	`< 2006-03-22 14:22:59`	Returns events with a timestamp before 2:23 PM, March 22, 2006.
>	`> today at 2:45pm`	Returns events with a timestamp later than today at 2:45 PM.

IP addresses in searches

When specifying IP addresses in searches, you can enter an individual IP address, a comma-separated list of addresses, an address block, or a range of IP addresses separated with a hyphen (-). You can also use negation.

For searches that support IPv6 (such as intrusion event, connection data, and correlation event searches) you can enter IPv4 and IPv6 addresses and CIDR/prefix length address blocks in any combination. When you search for hosts by IP address, the results include all hosts for which at least one IP address matches your search conditions, that is, a search for an IPv6 address may return hosts whose primary address in IPv4.

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

Because IP addresses can be represented by network objects, you can also click the add network **Object** (+) that appears next to an IP address search field to use a network object as an IP address search criterion.

Table 98: Acceptable IP address syntax

To specify...	Type...	For example...
a single IP address	the IP address.	`192.168.1.1` `2001:db8::abcd`
multiple IP addresses using a list	a comma-separated list of IP addresses. Do not add a space before or after the commas.	`192.168.1.1,192.168.1.2` `2001:db8::b3ff,2001:db8::0202`
a range of IP addresses that can be specified with a CIDR block or prefix length	the IP address block in IPv4 CIDR or IPv6 prefix length notation.	`192.168.1.0/24` This specifies any IP in the 192.168.1.0 network with a subnet mask of 255.255.255.0, that is, 192.168.1.0 through 192.168.1.255.
a range of IP addresses that cannot be specified with a CIDR block or prefix	the IP address range using a hyphen. Do not add a space before or after the hyphen.	`192.168.1.1-192.168.1.5` `2001:db8::0202-2001:db8::8329`

negation of any of the other ways to specify IP addresses or ranges of IP addresses	an exclamation point in front of the IP address, block, or range.	`192.168.0.0/32,!192.168.1.10`
		`!2001:db8::/32`
		`!192.168.1.10,!2001:db8::/32`

To specify...	Type...	For example...
hosts that are blocked or monitored (but would have been blocked) Refer to Host profile icons, on page 682.	In connection and Security Intelligence events, in Initiator IP and Responder IP fields: • block • monitor	--

Related Topics

IP address conventions, on page 27

URLs in searches

When searching for URLs, include wildcards. For example, use **`*example.com*`** to find all variations of the domain, such as **`https://example.com`** and **`division.example.com`** and **`example.com/division/`**.

Managed devices in searches

When you group devices—whether just on the Firewall Management Center, or as actual high availability or scalability configurations—searching for the name for the group correctly returns results for all devices in the group.

If the system finds a match for a group, it replaces the group name with the appropriate member device names for the purpose of performing the search. When you save a search that uses a device group in the device field the system saves the name specified in the device field and performs the device name replacement again each time the search is executed.

Ports in searches

The Firewall Management Center accepts specific syntax for port numbers in searches.

You can enter port numbers in these formats:

- A single port number.
- A comma-separated list of port numbers.
- Two port numbers separated by a dash to represent a range of port numbers.
- A port number followed by a protocol abbreviation, separated by a forward slash (only when searching for intrusion events).
- A port number or range of port numbers preceded by an exclamation mark to indicate a negation of the specified ports.

Note Avoid using spaces when specifying port numbers or ranges.

Table 99: Port syntax examples

Example	Description
`21`	Returns all events on port 21, including TCP and UDP events.
`!23`	Returns all events except those on port 23.
`25/tcp`	Returns all TCP-related intrusion events on port 25.
`21/tcp,25/tcp`	Returns all TCP-related intrusion events on ports 21 and 25.
`21-25`	Returns all events on ports 21 through 25.

Event fields in searches

These sections identify the available event fields that you can use as search criteria when performing event searches.

System, audit, and operational events:

- Audit log workflow fields, on page 427
- Health event fields, on page 414
- Intrusion rule update logs, on page 265
- Remediation status table fields, on page 968

Network traffic and security events:

- Connection and Security Intelligence Event Fields, on page 758
- Intrusion event fields, on page 790
- Correlation event fields, on page 960
- File and malware event fields, on page 841
- Captured file fields, on page 859
- Allow list event fields, on page 964
- Allow list violation fields, on page 966

Application events:

- Application data fields, on page 931
- Application detail data fields, on page 933

Discovery, host, and asset data:

- Discovery event fields, on page 913
- Host data fields, on page 915
- Host attribute data fields, on page 921

- Server data fields, on page 927
- User-related fields, on page 942

Vulnerability and scan data:

- Vulnerability data fields, on page 935
- Third-party vulnerability data fields, on page 940
- See *Nmap Scan Results Fields* in the *Cisco Secure Firewall Management Center Device Configuration Guide*

Perform a search

Search for specific events or data records within the system to analyze security events, audit logs, file events, health events, and other types of monitored data.

Before you begin

You must have **Admin**, **Security Analyst**, or **Security Analyst (Read Only)** user role to perform a search.

Procedure

Step 1 Choose **Analysis** > **Search**.

Alternatively, you may also click **Search** from any page on a workflow.

Step 2 From the table drop-down list, select the type of event or data to search.

Step 3 Enter your search criteria in the appropriate fields. For more information about applicable search fields, refer to Event Fields in Searches.

Step 4 If you want to use the search again in the future, save the search as described in Save a search, on page 713.

Step 5 Click **Search** to start the search.

Your search results appear in the default workflow for the table you are searching, constrained by time (if applicable).

What to do next

- To analyze the search results using workflows, see Use workflows, on page 676.

Related Topics

Configure event view settings, on page 234

Save a search

Save search criteria for future use to streamline repeated searches and maintain consistent filtering parameters.

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

Before you begin

- Establish search criteria as described in Perform a search, on page 713, or load a saved search as described in Load a saved search, on page 714.
- You must have **Admin**, **Security Analyst**, or **Security Analyst (Read Only)** user role to save a search.

Procedure

Step 1 From the Search page, if you want to save the search as private so only you can access it, check the **Private** checkbox.

Note
To use the search as a data restriction for a custom user role, you must save it as a private search.

Step 2 You have these options:

- To save a new version of a loaded search, click **Save As New**.
- To save a new search, or overwrite a custom search using the same name, click **Save**. If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Load a saved search

Load a saved search to quickly access previously configured search criteria without having to recreate the search constraints.

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

Before you begin

You must have **Admin**, **Security Analyst**, or **Security Analyst (Read Only)** privileges to load a saved search.

Procedure

Step 1 Choose **Analysis** > **Search**.

Alternatively, you may also click **Search** from any page on a workflow.

Step 2 From the table drop-down list, choose the type of event or data to search.

Step 3 Choose the search you want to load from the **Custom Searches** list or the **Predefined Searches** list.

Step 4 If you want to use different search criteria, change the search constraints.

Step 5 If you want to use a changed search again in the future, save the search as described in Save a search, on page 713.

Step 6 Click **Search**.

Query overrides via the shell

The query override tool in Firewall Management Center is a Linux shell-based query management option that allows you to locate queries running longer than a specified number of minutes and stop those queries. The tool logs an event to the audit log and to syslog when you stop a query.

The internal **Admin** user can access the Firewall Management Center CLI. If you use an external authentication object that grants CLI access, users matching the shell access filter can also log into the CLI.

Note If you leave the search page in the Firewall Management Center web interface, the query continues to run. Queries that take a long time to return results may impact system performance while the query is running.

Shell-Based query management syntax

Use this syntax to manage long-running queries:

```
query_manager [-v] [-l [minutes]] [-k query_id [...]] [--kill-all minutes]
```

Table 100: query_manager options

Option	Description
-h, --help	Prints a brief help message.
-l, --list [minutes]	Lists all queries taking longer than passed-in minutes. By default it will show all queries taking longer than 1 minute.
-k, --kill query_id [...]	Kills the query with the passed-in id. The option can take multiple ids.
--kill-all minutes	Kills all queries taking longer than passed-in minutes.
-v, --verbose	Verbose output including full SQL queries.

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

Stop long-running queries

Stopping long-running queries helps maintain system performance and prevents resource exhaustion.

Before you begin

You must be the **admin** user or externally authenticated user with CLI access to perform this task.

Procedure

Step 1 Connect to the Secure Firewall Management Center via `ssh`.

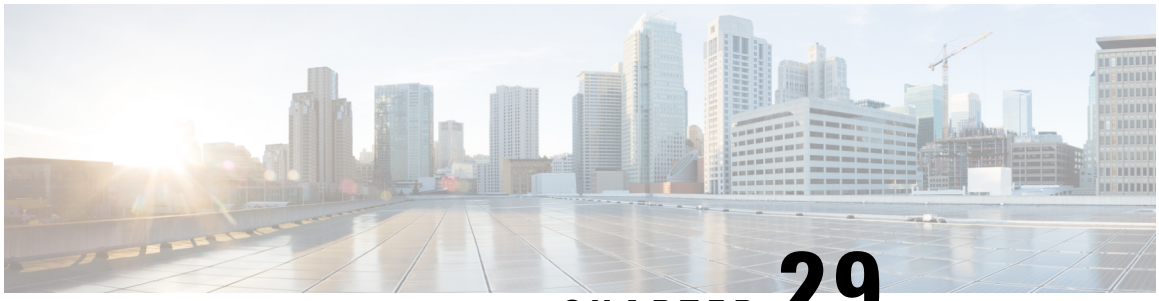
Step 2 Use the CLI `expert` command to access the Linux shell.

Step 3 Run `query_manager` under `sudo` using the syntax described in Shell-Based query management syntax, on page 715.

History for search for events

This history table provides version history information for the searching for events feature, including the changes and improvements across different Firewall Management Center software versions.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Partial-match searches in many fields now require wildcards	6.6	Any	For example, when searching for URLs, use **`*example.com*`** to find all variations of **`example.com`**. This behavior change applies to searches on the **Analysis** > **Search** page, when searching for connection or security-related connection events. This search page can also be accessed via links on other pages. In fields that do not require wildcards for partial-match searches, they can optionally be used. Affected platforms: Firewall Management Center



CHAPTER 29

Custom Workflows

The following topics describe how to use custom workflows:

- Custom workflows, on page 717
- Saved custom workflows, on page 717
- How custom workflows work, on page 718
- How to access custom workflows, on page 721

Custom workflows

Custom workflows are workflows you create to meet your specific needs when the predefined or Cisco-provided workflows do not meet your requirements.

Custom workflow features

- When creating a custom workflow, you select the event type or database table that the workflow will use as its basis.

  The Firewall Management Center allows you to base a custom workflow on a custom table.

- You can include drill-down pages, table views, and host or packet views in custom workflows to visualize and analyze data.
- If your event evaluation process changes, you can edit custom workflows to accommodate new requirements.

Note You can set a custom workflow as the default workflow for any event type.

Saved custom workflows

In addition to predefined workflows, which cannot be modified, the Firewall Management Center includes several saved custom workflows. Each of these workflows is based on a custom table and can be modified.

In a multidomain deployment, these saved workflows belong to the global domain and cannot be modified in lower domains.

Table 101: Saved custom workflows

Workflow name	Description
Events by Priority and Classification	This workflow lists events and their type in order of event priority, along with a count showing how many times each event has occurred. This workflow is based on the Intrusion Events custom table.
Hosts with Servers Default Workflow	You can use this workflow to quickly view the basic information in the Hosts with Servers custom table. This workflow is based on the Hosts with Servers custom table.
Server and Host Details	You can use this workflow to determine what servers are most frequently used on your network and which hosts are running those servers. This workflow is based on the Hosts with Servers custom table.

How custom workflows work

If the predefined and Cisco-provided custom workflows do not meet your needs, you can create custom workflows.



Note Instead of creating a new custom workflow, you can export a custom workflow from another appliance and then import it onto your appliance. You can then edit the imported workflow to suit your needs.

Summary

The key components in the process of creating a custom workflow are:

- Source table: Provides data foundation for the workflow.
- Drill-down pages: Display detailed data with customizable columns and sorting.
- Table view pages: Provide a standard table display without customizable properties.
- Final page: Automatically added based on selected table type.

Workflow

The custom workflow creation process involves these stages:

1. Initiate workflow creation and configure basic settings:
   - Select a table to be the source of the workflow.
   - Enter a workflow name.
   - Add drill-down pages and table view pages to the workflow.
2. Configure the page properties for each drill-down page in these ways:

- Provide a name that appears at the top of the page in the web interface.
- Specify the default sort order as ascending or descending.
- You can include up to five columns on each page.

Note

- Add at least one drill-down page or a table view of events to each custom workflow.
- If you selected **Vulnerabilities** as the table type, then add **IP Address** as a table column, the **IP Address** column does not appear when you view vulnerabilities using your custom workflow. It appears only when you use the search feature to limit the workflow to a specific IP address or block of addresses.

3. Add table view pages. You can add table view pages anywhere in the sequence of workflow pages. They do not support editable properties such as page name, sort order, or user-definable columns.
4. The Firewall Management Center automatically adds a final page based on the selected table type, as described in this table. This final page is added by default when you create the workflow.

Table 102: Custom workflow final pages

Event/Asset Type	Final Page
Discovery events	Hosts
Vulnerabilities	Vulnerability detail
Third-party vulnerabilities	Hosts
Users	Users
Indications of compromise	Hosts or users
Intrusion events	Packets

The system does not add a final page for custom workflows based on other kinds of events (for example, audit log or malware events).

Custom workflows based on connection data are similar to other workflows, except you can include drill-down pages containing connection summary data, connection data graph pages, drill-down pages with data for individual connections, and table view pages.

Create custom workflows based on non-connection data

Custom workflows based on non-connection data allow you to organize and analyze security information in a structured format that suits your specific monitoring and analysis needs.

Before you begin

You must have **Admin** or **Security Analyst** privileges to create a custom workflow based on non-connection data.

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Custom Workflows**.

Step 2 Click **Create Custom Workflow**.

Step 3 Enter a name for the workflow in the **Name** field.

Step 4 Enter a **Description**. This is optional.

Step 5 Choose the table you want to include from the **Table** drop-down list.

Step 6 To add one or more drill-down pages to the workflow, complete these steps:

a) Click **Add Page**.

b) Enter a name for the page in the **Page Name** field.

c) Under **Column 1**, set the preferred sort priority and select a table column to display as the left-most column.

Example:

To create a page showing the destination ports that are targeted, and to sort the page by count, choose **2** from the **Sort Priority** drop-down list and **Destination Port/ICMP Code** from the **Field** drop-down list.

d) Continue choosing fields and set their sort priority until you have specified all the fields to appear on the page.

Step 7 To add a table view page to the workflow, click **Add Table View**.

You cannot customize or configure table views.

Step 8 Click **Save**.

Create custom connection data workflows

Create custom workflows based on connection data to visualize and analyze network connection information through various graph types and data views.

Custom connection data workflows let you add connection data graph pages, drill-down pages, and table view pages. You can add any number of each page type to the workflow in any order. Each connection data graph page includes one graph: a line, bar, or pie chart. You can include more than one dataset on line and bar graphs.

Before you begin

You must have **Admin** privileges to create a custom workflow based on connection data.

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Custom Workflows**.

Step 2 Click **Create Custom Workflow**.

Step 3 Enter a name for the workflow in the **Name** field.

Step 4 Enter a **Description**. This is optional.

Step 5 From the **Table** drop-down list, choose **Connection Events**.

Step 6 To add one or more drill-down pages to the workflow, complete these steps:

a) Click **Add Page** to add a drill-down page that contains data on individual connections.

Alternatively, click **Add Summary Page** to add a drill-down page that contains connection summary data.

b) Enter a name for the page in the **Page Name** field.

c) Under **Column 1**, set the preferred sort priority and select a table column to display as the left-most column.

d) Continue choosing fields and set their sort priority until you have specified all the fields to appear on the page.

Example:

To create a page showing the amount of traffic transmitted over your monitored network and to sort the page by the responders that transmitted the most traffic, choose **1** from the **Sort Priority** drop-down list and **Responder Bytes** from the **Field** drop-down list.

Step 7 To add one or more graph pages to the workflow, complete these steps:

a) Click **Add Graph**.

b) Enter a name for the page in the **Graph Name** field.

c) Choose the type of graph you want to include on the page.

d) Specify what kind of data you want to graph by choosing the x-axis and y-axis of the graph.

On a pie chart, the x-axis represents the independent variable and the y-axis represents the dependent variable.

e) Choose the datasets you want to include on the graph.

Note that pie charts can include only one data set.

Step 8 To add a table view of connection data, click **Add Table View**.

You cannot customize or configure table views.

Step 9 Click **Save**.

How to access custom workflows

The method you use to view a workflow depends on whether the workflow is based on one of the predefined event tables or on a custom table.

- If your custom workflow is based on a predefined event table, access it in the same way that you would access a workflow that ships with the appliance. For example, to access a custom workflow based on the **Hosts** table, choose **Analysis** > **Hosts heading** > **Hosts**.
- If your workflow is based on a custom table, choose **Analysis** > **Advanced** > **Custom Tables** to access it.

If your event evaluation process changes, you can edit custom workflows to meet your new needs.

Note

- You cannot edit any of the predefined workflows.
- You can set a custom workflow as the default workflow for any event type.

View custom workflows based on predefined tables

View and manage custom workflows that utilize predefined tables in Firewall Management Center.

Before you begin

You must have **Admin**, **Maintenance User**, or **Security Analyst** user privileges to view a custom workflow.

Procedure

Step 1 Choose the appropriate menu path and option for the table on which you based your custom workflow, as described in the Workflow selection, on page 678.

Step 2 To switch workflows, click **(switch workflow)** next to the current workflow title and choose a different workflow.

Step 3 If no events appear and the workflow can be constrained by time, adjust the time range as needed to display relevant events. For more information, refer to Event time constraints, on page 692.

View custom workflows based on custom tables

View custom workflows that are based on custom tables. These workflows provide insights into custom data analysis and security event monitoring.

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

Before you begin

You must have **Admin** or **Security Analyst** privileges to view a custom workflow that is based on custom tables.

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Custom Tables**.

Step 2 Click **View** ( ) next to the custom table that you want to view.

Step 3 To switch workflows, click **(switch workflow)** next to the current workflow title and choose a different workflow, including custom workflows.

Step 4 If no events appear and the workflow can be constrained by time, adjust the time range as needed to display relevant events. For more information, refer to Event time constraints, on page 692.

Edit custom workflows

Edit custom workflows to modify existing workflow configurations and adapt them to changing security requirements.

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

Before you begin

You must have **Admin** or **Security Analyst** privileges to edit a custom workflow.

Procedure

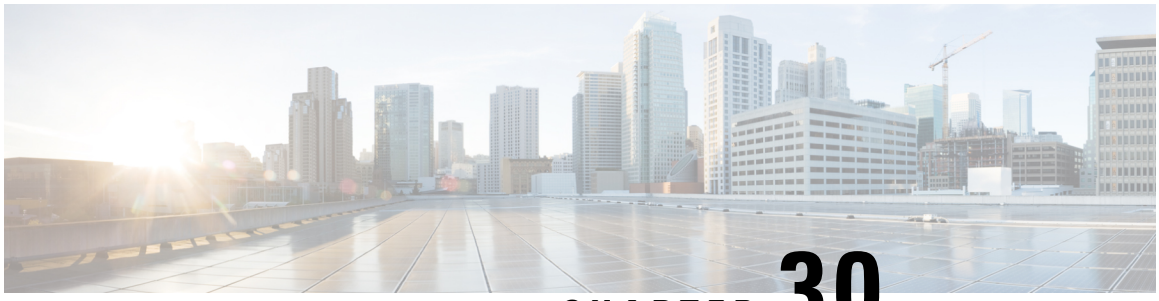
Step 1 Choose **Analysis** > **Advanced** > **Custom Workflows**.

Step 2 Click **Edit** ( ) next to the workflow that you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Edit the workflow as needed.

Step 4 Click **Save**.



CHAPTER 30

Custom Tables

The following topics describe how to use custom tables:

- Custom tables, on page 725
- Predefined custom tables, on page 725
- User-defined custom tables, on page 729
- Search custom tables, on page 731
- History for custom tables, on page 732

Custom tables

A custom table is a user-defined database table that combines fields from different predefined tables to enhance analysis of network activity.

The Firewall Management Center collects data about your network and stores it in a series of database tables. When you use a workflow to view specific information, the Firewall Management Center retrieves the data from one of these tables.For example, the **Network Applications by Count** workflow displays columns that are sourced from fields in the **Applications** table.

If you determine that your analysis of the activity on your network would be enhanced by combining fields from different tables, you can create a custom table.

Predefined custom tables

A predefined custom table is a system-delivered custom table that

- contains fields from two or more predefined tables
- correlates data across different table types for specific workflow needs, and
- provides a foundation for creating additional user-defined custom tables.

For example, Firewall Management Center includes predefined custom tables that correlate intrusion event data with host data. This enables you to search for events that impact critical systems and view the results in one workflow.

This table describes the custom tables provided in Firewall Management Center.

Table 103: Predefined custom tables

Table	Description
Hosts with Servers	Includes fields from the Hosts and Servers tables, providing you with information about the detected applications running on your network, as well as basic operating system information about the hosts running those applications.

You can create additional custom tables that contain only information that matches your specific needs. In a multidomain deployment, the predefined custom tables belong to the global domain. You cannot modify these tables in lower domains.

Possible table combinations

When you create a custom table, you can combine fields from predefined tables that have related data. This table lists the predefined tables you can combine to create a new custom table. Keep in mind that you can create a custom table that combines fields from more than two predefined custom tables.

Table 104: Custom table combinations

You can combine fields from...	With fields from...
Applications	• Correlation Events • Intrusion Events • Connection Summary Data • Host Attributes • Application Details • Discovery Events • Hosts • Servers • Allow List Events
Correlation Events	• Applications • Host Attributes • Hosts
Intrusion Events	• Applications • Host Attributes • Hosts • Servers

You can combine fields from...	With fields from...
Connection Summary Data	• Applications • Host Attributes • Hosts • Servers
HostIndications of Compromise	• Applications • Application Details • Captured Files • Connection Summary Data • Correlation Events • Discovery Events • Host Attributes • Hosts • Intrusion Events • Security Intelligence Events • Servers • Allow List Events
Host Attributes	• Applications • Correlation Events • Intrusion Events • Connection Summary Data • Application Details • Discovery Events • Hosts • Servers • Allow List Events
Application Details	• Applications • Host Attributes • Hosts

You can combine fields from...	With fields from...
Discovery Events	• Applications • Host Attributes • Hosts
Security Intelligence Events	• Applications • Host Attributes • Hosts • Servers
Hosts	• Applications • Correlation Events • Intrusion Events • Connection Summary Data • Host Attributes • Application Details • Discovery Events • Servers • Allow List Events
Servers	• Applications • Intrusion Events • Connection Summary Data • Host Attributes • Hosts
Allow List Events	• Applications • Host Attributes • Hosts

Sometimes a field in one table maps to more than one field in another table.

When you create a new custom table, a default workflow that displays all the columns in the table is automatically created. Also, just as with predefined tables, you can search custom tables for data that you want to use in your network analysis. You can also generate reports based on custom tables, as you can with predefined tables.

User-defined custom tables

A user-defined custom table is a table that you create or import to combine selected fields from multiple predefined tables to create customized data views. It lets you view related data from different tables in a single context, such as host details associated with specific events.

To create a custom table, first decide which predefined tables contain the fields you want to include. You can then select the desired fields and, if necessary, configure field mappings for any common fields.

Instead of creating a new custom table, you can export a custom table from another Firewall Management Center, then import it onto your Firewall Management Center.

Correlation events and hosts table combination

Consider a custom table that combines fields from the **Correlation Events** table and the **Hosts** table. This custom table enables you to get detailed information about the hosts involved in violations of your correlation policies. Note that you must decide whether to display data from the **Hosts** table that matches the source IP address or the destination IP address in the **Correlation Events** table.

If you view the table for this custom table, it displays correlation events (one per row) and may include these information:

- the date and time the event was generated
- the name of the correlation policy that was violated
- the name of the rule that triggered the violation
- the IP address associated with the source, or initiating, host involved in the correlation event
- the source host's NetBIOS name
- the operating system and version the source host is running, and
- the source host criticality.

Create a custom table

Create custom tables that contain only information that matches your specific needs. Custom tables allow you to combine information from multiple data sources to create comprehensive views of security events and associated host details.

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Custom Tables**.

Step 2 Click **Create Custom Table**.

Step 3 In the **Name** field, enter a name for the custom table.

Step 4 From the **Tables** drop-down list, choose a predefined table that has related data.

Step 5 From **Fields** drop-down list, choose the fields that contain data for your specific needs.

You can combine fields from multiple predefined tables that have related data. After completing Step 5, select another predefined table that contains the required data and repeat the selection process as needed.

Note
To select multiple fields at once, use **Ctrl**, **command**, or **Shift** key while selecting fields. You can also click and drag to select multiple adjacent fields. If you want to specify the order in which fields appear in the table view, add them one at a time.

Step 6 After selecting all required fields, click **Add** to add them to the table.

Step 7 Click **Save**.

Modify a custom table

Modify a custom table to update its structure, remove unnecessary fields, or change table properties to better meet your reporting needs.

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

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Custom Tables**.

Step 2 Click the **Edit** ( ) icon next to the custom table you want to modify.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Optionally, remove fields from the table by clicking **Delete** ( ) next to the fields you want to remove.

Note
If you delete a field that is currently in use in a report, the Firewall Management Center will prompt you to confirm whether you want to remove any sections from the report that are using this field.

Step 4 Make other changes as needed.

Step 5 Click **Save**.

Delete a custom table

Delete custom tables that are no longer needed from the current domain.

In a multidomain deployment, the system displays custom tables created in the current domain, which you can delete. It also displays custom tables created in ancestor domains, which you cannot delete. To delete custom tables in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Custom Tables**.

Step 2 Click the **Delete** ( ) icon next to the custom table you want to delete.

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

View a workflow based on a custom table

View workflows to examine events in custom tables and access table views, packet views, or hosts pages depending on your custom table configuration.

When you create a custom table, Firewall Management Center automatically creates a default workflow for it. The first page of this workflow displays a table view of events. If you include intrusion events in your custom table, the second page of the workflow is the packet view. Otherwise, the second page of the workflow is a hosts page. You can also create your own custom workflows based on your custom table.

Note If you create a custom workflow based on a custom table, you can specify it as the default workflow for that table.

You can use the same method to view events in your custom table that you use for event views based on predefined tables.

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

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Custom Tables**.

Step 2 Click the **View** ( ) icon next to the custom table related to the workflow that you want to see.

Search custom tables

Search for specific data within custom tables to locate relevant records based on your search criteria.

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

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Custom Tables**.

Step 2 Click the **View** ( ) icon next to the custom table you want to search.

Note
To use a different workflow, including a custom workflow, click **(switch workflow)** next the workflow title.

Step 3 Click **Search**.

Note
To search the database for a different kind of event or data, choose it from the table drop-down list.

Step 4 Enter your search criteria in the appropriate fields.

If you enter criteria for multiple fields, the search returns only the records that match search criteria specified for all fields.

Note
Click **Object** (+) next to a search field to use an object as a search criterion.

Step 5 Optionally, if you plan to save the search, you can check the **Private** check box to save the search as private so only you can access it. Otherwise, leave the check box clear to save the search for all users.

Note
If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search.

Step 6 Optionally, you can save the search to be used again in the future. You have the following options:

- Click **Save** to save the search criteria. The search is visible only to your account if you checked the **Private** check box.
- Click **Save As New** to save a new search or assign a name to a search you created by altering a previously-saved search. The search is saved and visible only to your account if you checked the **Private** check box.

Step 7 Click **Search** to start the search.

Your search results appear in the default workflow for the custom table, constrained by the current time range (if applicable).

History for custom tables

This table provides a chronological history of Custom Tables feature updates and improvements, enabling you to track feature availability and requirements across different Firewall Management Center versions.

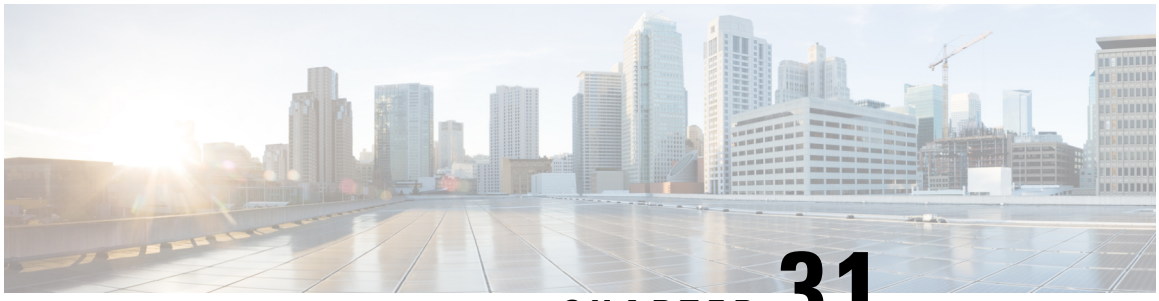
Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Support for connection events in custom tables was removed	6.6	Any	You can no longer create custom tables that include connection events. If you upgraded to version 6.6: Existing tables with connection events will be listed as deprecated and will show no data, and you cannot export or edit them. Existing reports, custom workflows, and dashboards may include deprecated tables; you may want to review these. Modified screens: **Analysis** > **Advanced** > **Custom Tables** and the page for adding or editing custom tables.



PART VIII

Events and Assets

- Connection Logging, on page 737
- Connection and Security-Related Connection Events, on page 755
- Intrusion Events, on page 787
- File/Malware Events and Network File Trajectory, on page 835
- Host Profiles, on page 871
- Discovery Events, on page 899
- Correlation and Compliance Events, on page 959



CHAPTER 31

Connection Logging

The following topics describe how to configure the system to log connections made by hosts on your monitored network:

- Connection logging, on page 737
- Connection logging limitations, on page 746
- Best practices for connection logging, on page 747
- Requirements and prerequisites for connection logging, on page 750
- Connection logging, on page 750

Connection logging

Connection logging is a system feature that

- generates logs of connections detected by managed devices known as *connection events*,
- provides granular control through rules and policies over which connections to log, when to log them, and where to store the data, and
- creates special connection events called *security intelligence events* that represent connections blocked by the reputation-based Security Intelligence feature.

Connection event data

Connection events contain data about the detected sessions. The specific data available for any event can vary, but generally includes:

- Basic connection properties: timestamp, source and destination IP address, ingress and egress zones, the device that handled the connection, and so on
- Additional connection properties discovered or inferred by the system: applications, requested URLs, or users associated with the connection, and so on
- Metadata about why the connection was logged: which configuration handled the traffic, whether the connection was allowed or blocked, details about encrypted and decrypted connections, and so on

Log connections according to the security and compliance needs of your organization. When setting up connection logging, remember that the system can log a connection for multiple reasons. Disabling logging in one place does not guarantee that matching connections will not get logged elsewhere.

The information contained in a connection event depends on factors such as traffic characteristics and relevant configuration.



Note You can supplement the connection logs gathered by your managed devices with connection data generated from exported NetFlow records. This is particularly useful if you have NetFlow-enabled routers or other devices monitoring networks that managed devices cannot access.

Logged connections

Logged connections are network connections that

- are automatically saved to the Firewall Management Center database as end-of-connection events,
- generate logs regardless of other logging configurations unless connection event storage is disabled, and
- include connections associated with intrusions, file and malware events, intelligent application bypass, and monitored traffic.

Connection types and logging behavior

The system automatically logs specific connection types based on their association with security events and traffic monitoring.

Connection types that generate automatic logs:

- Connections associated with intrusions: The system automatically logs connections associated with intrusion events, unless the connection is handled by the access control policy's default action.
- Connections associated with file and malware events: The system automatically logs connections associated with file and malware events.
- Connections associated with intelligent application bypass: The system automatically logs bypassed and would-have-bypassed connections associated with IAB.
- Monitored connections: The system always logs the ends of connections for monitored traffic, even if the traffic matches no other rules and you do not enable default action logging.



Note File events generated by inspecting NetBIOS-SSN (SMB) traffic do not immediately generate connection events because the client and server establish a persistent connection. The system generates connection events after the client or server ends the session.

Special behavior for default action connections:

- When an intrusion policy associated with the access control default action generates an intrusion event, the system does *not* automatically log the end of the associated connection. Instead, you must explicitly enable default action connection logging.

- However, if you enable beginning-of-connection logging for the default action, the system *does* log the end of the connection when an associated intrusion policy triggers, in addition to logging the beginning of the connection.

For more information about monitored connections, refer to Logging for monitored connections, on page 740.

Other Connections You Can Log

So that you log only critical connections, enable connection logging on a per-rule basis. If you enable connection logging for a rule, the system logs all connections handled by that rule.

You can also log connections handled by policy default actions. Depending on the rule or default action (and for access control, a rule's inspection configuration), your logging options differ.

Prefilter Policy: Rules and Default Action

You can log connections (including entire plaintext, passthrough tunnels) that you fastpath or block with a prefilter policy.

Prefiltering uses outer-header criteria to handle traffic. For tunnels that you log, the resulting connection events contain information from the outer, encapsulation headers.

For traffic subject to further analysis, logging in the prefilter policy is disabled, although matching connections may still be logged by other configurations. The system performs all further analysis using inner headers, that is, the system independently handles and logs each connection within an allowed tunnel.

SSL Policy: Rules and Default Action

You can log connections that match an SSL rule or SSL policy default action.

For blocked connections, the system immediately ends the session and generates an event. For monitored connections and connections that you pass to access control rules, the system generates an event when the session ends.

Access Control Policy: Security Intelligence Decisions

You can log a connection whenever it is blocked by the reputation-based Security Intelligence feature.

Optionally, and recommended in passive deployments, you can use a monitor-only setting for Security Intelligence filtering. This allows the system to further analyze connections that would have been blocked by Security Intelligence, but still log the match. Security Intelligence monitoring also allows you to create traffic profiles using Security Intelligence information.

When the system logs a connection event as the result of Security Intelligence filtering, it also logs a matching Security Intelligence event, which is a special kind of connection event that you can view and analyze separately, and that is also stored and pruned separately. So that you can identify the matching IP address in the connection, host icons beside blocked and monitored IP addresses look slightly different in the tables on the pages under the **Analysis** > **Connections** > **Events** menus.

Access Control Policy: Rules and Default Action

You can log connections that match an access control rule or access control policy default action.

Related Topics

Rules and policy actions effect on logging, on page 740

Rules and policy actions effect on logging

Rules and policy actions effect on logging is a traffic handling mechanism that

- determines how the system inspects and handles matching traffic,
- controls when and how details about matching traffic can be logged, and
- provides metadata about why the connection was logged, including which configurations handled the traffic.

Connection event metadata

Connection events contain metadata about why the connection was logged, including which configurations handled the traffic. Where you can configure connection logging, rule actions, and policy default actions determine not only how the system inspects and handles matching traffic, but also when and how you can log details about matching traffic.

Related Topics

Connection and Security Intelligence Event Fields, on page 758

Logging for fastpathed connections

Logging for fastpathed connections is a monitoring capability that

- records connection events for traffic that matches tunnel rules and prefilter rules with **Fastpath** actions,
- captures information for non-encrypted tunnels including outer session details for tunnel rules, and
- provides limited information in connection events since fastpathed traffic bypasses access control and QoS processing.

Supported traffic types

You can log fastpathed connections and non-encrypted tunnels, which include traffic matching these rules and actions in the prefilter policy:

- Tunnel rules—**Fastpath** action (logs the outer session)
- Prefilter rules—**Fastpath** action

Fastpathed traffic bypasses the rest of access control and QoS, so connection events for fastpathed connections contain limited information.

Logging for monitored connections

Logging for monitored connections is a system function that

- automatically logs connection ends for traffic matching specific monitoring configurations,
- operates independently of default action logging settings, and
- consolidates multiple Monitor rule matches into single connection events.

Monitored traffic configurations

The system always logs the ends of connections for traffic matching these configurations, even if the traffic matches no other rules and you do not enable default action logging:

- Security Intelligence: Block lists set to monitor (also generates a Security Intelligence event)
- SSL rules: **Monitor** action
- Access control rules: **Monitor** action

The system does not generate a separate event each time a single connection matches a Monitor rule. Because a single connection can match multiple Monitor rules, each connection event can include and display information on the first eight Monitor access control rules that the connection matches, as well as the first matching SSL Monitor rule.

Similarly, if you send connection events to an external syslog or SNMP trap server, the system does not send a separate alert each time a single connection matches a Monitor rule. Rather, the alert that the system sends at the end of the connection contains information on the Monitor rules the connection matched.

Logging for trusted connections

Logging for trusted connections is a system capability that

- records the beginnings and ends of trusted connections for traffic matching access control rules with **Trust** action
- captures events for traffic matching the access control default action **Trust All Traffic**, and
- generates limited information because trusted connections are not subject to deep inspection or discovery.

Trusted connection logging behavior

You can log the beginnings and ends of trusted connections, which include traffic matching these rules and actions:

- Access control rules: **Trust** action
- Access control default action: **Trust All Traffic**



Note Although you *can* log trusted connections, we recommend you do not do so because trusted connections are not subject to deep inspection or discovery, so connection events for trusted connections contain limited information.

TCP connections detected by a Trust rule on the first packet generate only an end-of-connection event. The system generates the event one hour after the final session packet.

Logging for blocked connections

Logging for blocked connections is a security feature that

- records traffic matching blocking rules and actions across multiple policy types,
- generates connection events with different timing based on the blocking mechanism, and

- provides visibility into denied traffic for security monitoring and analysis.

Blocked connection types

You can log blocked connections, which include traffic matching these rules and actions:

- Tunnel rules—**Block**
- Prefilter rules—**Block**
- Prefilter default action—**Block all tunnel traffic**
- Security Intelligence—Block lists not set to Monitor (also generates a Security Intelligence event)
- SSL rules—**Block** and **Block with reset**
- SSL default action—**Block** and **Block with reset**
- Access control rules—**Block**, **Block with reset**, and **Interactive Block**
- Access control default action—**Block All Traffic**

Only devices deployed inline (that is, using routed, switched, or transparent interfaces, or inline interface pairs) can block traffic. Because blocked connections are not actually blocked in passive deployments, the system may report multiple beginning-of-connection events for each blocked connection.



Caution Firewall Threat Defense device logs end-of-connection events for TCP connections where the three-way handshake does not complete, including SYN-only, SYN-SYN/ACK (no ACK), and SYN-RST/ACK scenarios (except when TLS Identity is enabled). These logs provide visibility into incomplete or suspicious connection attempts. However, 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.

Beginning vs. end-of-connection logging for blocked connections

When you log a blocked connection, how the system logs it depends on why the connection was blocked. This is important to keep in mind when configuring correlation rules based on connection logs:

- For SSL rules SSL policy default actions that block encrypted traffic, the system logs **end**-of-connection events. This is because the system cannot determine if a connection is encrypted using the first packet in the session.
- For other blocking actions, the system logs **beginning**-of-connection events. Matching traffic is denied without further inspection.

Logging bypassed interactive blocks

Interactive blocking access control rules, which cause the system to display a warning page when a user browses to a prohibited website, allow you to configure end-of-connection logging. This is because if the user clicks through the warning page, the connection is considered a new, allowed connection which the system can monitor and log.

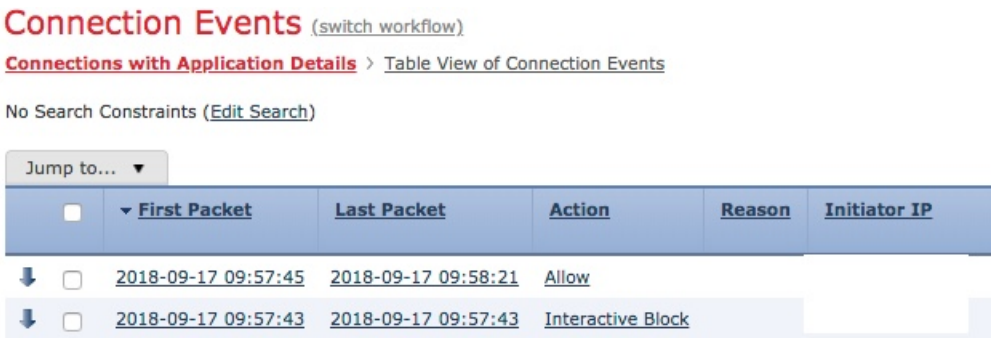
Therefore, for packets that match an Interactive Block or Interactive Block with Reset rule, the system can generate these connection events:

- A beginning-of-connection event when a user's request is initially blocked and the warning page is displayed; this event has an associated action of **Interactive Block** or **Interactive Block with Reset**
- Multiple beginning- or end-of-connection events if the user clicks through the warning page and loads the originally requested page; these events have an associated action of **Allow** and a reason of **User Bypass**

Interactive block example

The following figure shows an example of an interactive block followed by allow.

Figure 28: Interactive block followed by allow



Logging for allowed connections

Logging for allowed connections is a security feature that

- generates file and malware events when file policies detect or block files during traffic inspection, and
- provides configurable logging options on a per-access-control-rule basis.

Connection logging details

You can log allowed connections, which includes traffic matching the following rules and actions:

- SSL rules—**Decrypt** action
- SSL rules—**Do not decrypt** action
- SSL default action—**Do not decrypt**
- Access control rules—**Allow** action
- Access control default action—**Network Discovery Only** and any intrusion prevention option

Enabling logging for these configurations ensures the connection is logged, while also permitting (or specifying) the next phase of inspection and traffic handling. SSL logging is always end-of-connection; access control configurations also allow beginning-of-connection logging.

Although the **Analyze** action in tunnel and prefilter rules also allows connections to continue with access control, logging is disabled for rules with this action. Matching connections may still be logged by other configurations. Allowed tunnels might have their encapsulated sessions evaluated and logged individually.

When you allow traffic with an access control rule or default action, you can use an associated intrusion policy to further inspect traffic and block intrusions. For access control rules, you can also use a file policy to detect and block prohibited files, including malware. Unless you disable connection event storage, the system automatically logs most allowed connections associated with intrusion, file, and malware events. For detailed information, refer to Logged connections, on page 738.

Connections with encrypted payloads are not subject to deep inspection, so connection events for encrypted connections contain limited information.

When a file policy detects or blocks a file, it logs one of these events to the Firewall Management Center database:

- *File events*, which represent detected or blocked files, including malware files.
- *Malware events*, which represent detected or blocked malware files only.
- *Retrospective malware events*, which are generated when the malware disposition for a previously detected file changes.

You can disable this logging on a per-access-control-rule basis. You can also disable file and malware event storage entirely.

Note We recommend you leave file and malware event logging enabled.

Connection logging modes

A connection logging mode is a connection monitoring approach that

- provides two distinct timing options for capturing connection events,
- handles blocked traffic differently than allowed traffic, and
- requires optimization choices to balance performance with information completeness.

Connection logging timing and exceptions

You can log a connection at its beginning or its end, with specific exceptions for blocked traffic:

- Blocked traffic: Because blocked traffic is immediately denied without further inspection, only beginning-of-connection events are typically available for logging; there is no unique end-of-connection to log.
- Blocked encrypted traffic: When you enable connection logging in an SSL policy, the system logs end-of-connection rather than beginning-of-connection events. This occurs because the system cannot determine if a connection is encrypted with the first packet and therefore cannot immediately block encrypted sessions.

To optimize performance, log either the beginning or the end of any connection, but not both. Monitoring a connection for any reason forces end-of-connection logging. For a single non-blocked connection, the

end-of-connection event contains all information present in the beginning-of-connection event, as well as data gathered over the duration of the session.

The table presents differences between beginning and end-of-connection events, along with the advantages of logging each.

Table 105: Comparing beginning and end-of-connection events

Attributes	Beginning-of-Connection Events	End-of-Connection Events
Can be generated...	When the system detects the beginning of a connection (or, after the first few packets if event generation depends on application or URL identification).	When the system: • Detects the close of a connection. • Does not detect the end of a connection after a period of time. • Can no longer track the session due to memory constraints.
Can be logged for...	All connections except those blocked by the SSL policy.	Most connections.
Contain...	Only information that can be determined in the first packet (or the first few packets, if event generation depends on application or URL identification).	All information in the beginning-of-connection event, plus information determined by examining traffic over the duration of the session; for example, the total amount of data transmitted or the timestamp of the last packet in the connection. **Note** The connection event does not count the amount of data transmitted after the threat defense returns a snort verdict for the connection or if you fastpath the connection.
Are useful...	If you want to log: • Blocked connections. • Only the beginning of a connection because the end-of-connection information does not matter to you.	If you want to: • Log encrypted connections handled by an SSL policy. • Perform any kind of detailed analysis on, or trigger correlation rules using, information collected over the duration of the session. • View connection summaries (aggregated connection data) in custom workflows, view connection data in graphical format, or create and use traffic profiles.

Secure Firewall Management Center vs external logging

A log storage option in the Firewall Management Center enables analysis of connection and Security Intelligence event logs through dashboards, event views, traffic profiling, and correlation policies. The number of stored

events depends on the model. Logging usually occurs at connection end, with critical connections logged automatically. Events can also be sent externally via syslog or SNMP traps using alert responses or specific configurations, influencing log storage and processing methods.

Secure Firewall Management Center logging features

If you store connection and Security Intelligence event logs on the Firewall Management Center, you can use the system's reporting, analysis, and data correlation features. For example:

- Dashboards and the Context Explorer provide you with graphical, at-a-glance views of the connections logged by the system.
- Event views (most of the options available under the Analysis menu) present detailed information on the connections logged by the system, which you can display in a graphical or tabular format or summarize in a report.
- Traffic profiling uses connection data to create a profile of your normal network traffic that you can then use as a baseline against which to detect and track anomalous behavior.
- Correlation policies allow you to generate events and trigger responses (such as alerts or external remediations) to specific types of connections or traffic profile changes.

The number of events the Firewall Management Center can store depends on its model.

Note To use these features, you **must** log connections (and in most cases, the end of those connections rather than the beginning). This is why the system automatically logs critical connections—those associated with logged intrusions, prohibited files, and malware.

You can also log events to an external syslog or SNMP trap server, or to other external tools, using these options:

- For external logging on any device:

  A connection you configure called an *alert response*.

- For external logging on Firewall Threat Defense devices:

  Refer to *About Configuring Syslog* and *Configure SNMP Traps* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

- For additional options related to external logging:

  Refer to Event Analysis Using External Tools, on page 635.

Related Topics

Configuring external alerts with alert responses, on page 561

Connection logging limitations

Connection logging limitations are restrictions that prevent certain types of network connections from being logged by firewall systems.

Connection types that cannot be logged

Connection logging cannot capture information for these connection types:

- The outer session of a plaintext, passthrough tunnel whose encapsulated connections are inspected by access control.
- TCP connections if the three-way handshake is not completed, to avoid denial-of-service attacks against your firewalls. To monitor or debug failed connections, you can use the **show asp drops** CLI command or the packet capture feature (Packet capture, on page 460).

If a connection event does not contain the information you think it should, refer to Requirements for populating connection event fields, on page 776 and Information available in connection event fields, on page 778.

Event viewer event appearances

Event viewer event appearances are the display characteristics that

- require manual page refreshes for events under the Analysis menu,
- typically display events within seconds of detection, but may experience delays under heavy traffic or system operations, and
- show all connection events logged per defined rules without filtering options for unified logging.

Event display characteristics

These points are applicable to all types of events:

- If you are looking at a page under the Analysis menu, you must refresh the page to display new events.
- Events generally are available for viewing within a few seconds of the time the traffic was detected. However, there can be an arbitrary delay under situations such as: Exceptionally heavy traffic conditions; the Firewall Management Center is managing a lot of devices on a low-bandwidth network; or during operations such as event backup which pause event processing.
- All the connection events logged as per the defined rules appear in the event viewer. An option to filter the events is not available for the unified logging of the connection events.

Best practices for connection logging

Best practices for connection logging are guidelines that

- ensure you log only critical connections by enabling connection logging on a per-access-control-rule basis,
- help avoid logging unnecessary traffic such as trusted connections or infrastructure traffic not related to your service offering, and
- optimize storage and performance by preventing duplicate logging and configuring appropriate event limits.

Connections that are always logged

The system automatically logs these connections:

- Some connections associated with detected files, malware, intrusions, and Intelligent Application Bypass (IAB).

  For more information, refer to Logged connections, on page 738.

- Monitored connections.

  For more information, refer to Logging for monitored connections, on page 740.

Connections to never log

Do not enable logging for the following:

- Access control rules with a Trust action.

  Trusted connections are not subject to deep inspection or discovery, so connection events for trusted connections contain limited information.

- Do not enable logging for Block rules in passive deployments. To log connections that the system would have blocked if your devices were deployed inline, use a Monitor rule instead of a Block rule.

  Only devices deployed inline (that is, using routed, switched, or transparent interfaces, or inline interface pairs) can block traffic. Because blocked connections are not actually blocked in passive deployments, the system may report multiple beginning-of-connection events for each blocked connection.

- Traffic you're not interested in. Examples follow:

  - Specific allowed traffic, such as DNS requests to a trusted DNS host.
  - Infrastructure traffic that is not related to your service offering.

  (As previously mentioned, you can still monitor this traffic for threats.)

As discussed in Logged connections, on page 738, even if you disable logging for the preceding, intrusion events, malware, and IAB are still logged.

Avoid logging what is being logged elsewhere

If another device or service is logging connection data for a network segment, disable logging for that segment's data in the Firewall Management Center. Examples:

- If a router logs connection events on the same network segment as the Firewall Management Center, avoid logging the same connections on the Firewall Management Center unless you need those connection events for something else, such as correlation policies or traffic profiles.

  For more information about correlation policies, refer to Correlation policies and rules, on page 989. For more information about traffic profiles, refer to Introduction to traffic profiles, on page 1027.

- If you use Secure Network Analytics to leverage NetFlow records reported from switches and routers to identify potential behavioral anomalies and suspicious traffic patterns, you can disable connection logging for rules monitoring those segments and instead rely on Secure Network Analytics for behavioral analytics for those parts of your network.

  For more information, consult the Secure Network Analytics documentation.

Log either the beginning or end of the connection (not both)

If you have a choice between beginning and end-of-connection logging, enable end-of-connection logging. This is because end-of-connection logs information from beginning-of-connection events, as well as information gathered over the duration of the session.

Log the beginning of connections only if you want to log blocked connections, or if end-of-connection information does not matter to you.

For more information, refer to Connection logging modes, on page 744.

Logging for blocked traffic

Because blocked traffic is immediately denied without further inspection, usually you can log only beginning-of-connection events.

For more information, refer to Logging for blocked connections, on page 741.

Log events to an external location

If your company's security policy permits it, you can save disk space on your Firewall Management Center by streaming logs to an external source using any of the following:

- eStreamer, which enables you to stream logs from a Firewall Management Center to a custom-developed client application. For more information, refer to the *Secure Firewall Management Center Event Streamer Integration Guide* .
- Syslog or SNMP trap are referred to as alert responses. For more information, refer to Configuring external alerts with alert responses, on page 561.

Specify the maximum number of event records

Consider the minimum and maximum number of records that can be stored in the database for each event type. For example, a virtual Firewall Management Center can store a maximum of 50 million connection events. Additionally, other event types have their own configurable limits. Go to **System** (⚙) > **Configuration** > **Database** to adjust these limits to meet your storage and performance needs.

For a list of all Firewall Management Center models and their event database sizes, refer to Database event limits, on page 68.

Control what is displayed in connection events

To specify the number of rows displayed in connection events, click your username in the upper right of the Firewall Management Center, click **User Preferences**, and then click the **Event View Settings**tab. The maximum you can set is 1000 events per page.

Set up connection event reports

To make sure you do not miss connection events, you can set up automated reports in .csv format and optionally schedule them to occur at a regular interval. For more information, refer to the following:

- To create a report template, click **Analysis** > **Connections** > **Events** and then from the event view page, click **Create Report**. For more information, refer to Create a report template from an event view, on page 536.
- Schedule tasks (**System** (⚙) > **Tools** > **Scheduling**): Scheduling, on page 505.

Requirements and prerequisites for connection logging

This reference provides the technical requirements and prerequisites that must be met to implement and use connection logging features.

Model support

Any.

Supported domains

Any

User roles

- Admin
- Access Admin
- Network Admin

Connection logging

Connection logging is a network monitoring feature that records network connection activities to match various rules and conditions.

Logging connections with tunnel and prefilter rules

This task enables logging for prefilter policy rules to monitor and analyze connection events for troubleshooting and security analysis purposes.

The prefilter policy applies to Secure Firewall Threat Defense devices only.

Before you begin

- Set the rule action to **Block** or **Fastpath**. Logging is disabled for the **Analyze** action, which allows connections to continue with access control, where other configurations determine their handling and logging.
- Logging is performed on inner flows, not on the encapsulating flow.

Follow these steps to configure logging for tunnel and prefilter rules:

Procedure

Step 1 In the prefilter policy editor, click **Edit** ( ) next to the rule where you want to configure logging.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 2 Click **Logging**.

Step 3 Specify whether you want to **Log at Beginning of Connection** or **Log at End of Connection**.

To optimize performance, log either the beginning or the end of any connection, but not both. Because blocked traffic is immediately denied without further inspection, you cannot log end-of-connection events for Block rules.

Step 4 Specify where to send connection events:

Step 5 Click **Save** to save the rule.

Step 6 Click **Save** to save the prefilter policy.

What to do next

Deploy configuration changes; see the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Configure logging for decryptable connections with rules

Logging decryptable connections allows you to track and analyze traffic that is being monitored by decryption rules. This is particularly important for security analysis and compliance requirements.

Procedure

Step 1 In the SSL policy editor, click **Edit** ( ) next to the rule where you want to configure logging.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 2 Click **Logging**.

Step 3 Check the **Log at End of Connection** check box.

For monitored traffic, end-of-connection logging is required.

Step 4 Specify where to send connection events.

Send events to the event viewer if you want to perform Firewall Management Center-based analysis on these connection events. For monitored traffic, this is required.

Step 5 Click **Save** to save the rule.

Step 6 Click **Save** to save the policy.

What to do next

Deploy configuration changes; see the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Logging connections with security intelligence

Enable Security Intelligence logging to monitor and analyze network connections, IP addresses, URLs, and domain names for security threats and policy compliance.

The Security Intelligence policy requires the Threat Smart License or Protection Classic License.

Procedure

Step 1 In the access control policy editor, click **Security Intelligence**.

Step 2 Click the **Logging** ( ) icon to enable Security Intelligence logging using the following criteria:

- By IP address: Click the logging icon next to **Networks**.
- By URL: Click the logging icon next to **URLs**.
- By Domain Name: Click the logging icon next to the **DNS Policy** drop-down list.

If the logging icon is disabled, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Check the **Log Connections** check box.

Step 4 Specify where to send connection and Security Intelligence events.

Send events to the event viewer if you want to perform Firewall Management Center-based analysis, or if you set a Block list to monitor-only.

Step 5 Click **OK** to set logging options.

Step 6 Click **Save** to save the policy.

What to do next

Deploy configuration changes; see the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Logging connections with access control rules

Depending on your choices for the rule action and deep inspection options, your logging options differ; refer to Rules and policy actions effect on logging, on page 740.

Procedure

Step 1 In the access control policy editor, click **Edit** ( ) next to the rule where you want to configure logging.

If **View** ( ) appears instead, the configuration is inherited from an ancestor policy, belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 2 Click the **Logging** tab.

Step 3 Specify whether you want to **Log at Beginning of Connection** or **Log at End of Connection**.

To optimize performance, log either the beginning or the end of any connection, but not both.

Step 4 (Optional) Check the **Log Files** check box to log file and malware events associated with the connection.

It is recommended to leave this option enabled.

Step 5 Specify where to send the connection events:

- **Event Viewer** (or a product name): Send connection events to Firewall Management Center (or other device manager) if you want to perform Firewall Management Center-based analysis on these connection events, or if the rule action is **Monitor**.
- **Syslog Server**: Send connection events to the syslog server configured in the Logging tab in Access Control Policy, unless overridden.

  Show Overrides: Displays the options to override the settings configured in the access control policy.

  - **Override Severity**: When you choose this option and select a severity for the rule, connection events for this rule will have the selected severity regardless of the severity configured in the Logging tab in Access Control Policy.
  - **Override Default Syslog Destination**: Send the syslog generated for the connection event for this rule to destination specified in this alert.

- **SNMP Trap**: Connection events are sent to the selected SNMP trap.

Step 6 Click **Save** to save the rule.

What to do next

Deploy configuration changes; see the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Logging connections with a policy default action

A policy's default action determines how the system handles traffic that matches none of the rules in the policy (except Monitor rules in access control and SSL policies, which match and log—but do not handle or inspect—traffic).

Logging settings for the SSL policies default action also govern how the system logs undecryptable sessions.

Before you begin

For prefilter default action logging, set the default action to **Block all tunnel traffic**. Logging is disabled for the **Allow all tunnel traffic** action, which allows connections to continue with access control, where other configurations determine their handling and logging.

Procedure

Step 1 In the policy editor, click the **Logging** ( ) next to the **Default Action** drop-down list.

Step 2 Specify when you want to log matching connections:

- Log at Beginning of Connection: Not supported for SSL default actions.
- Log at End of Connection: Not supported if you choose the access control **Block All Traffic** default action or the prefilter **Block all tunnel traffic** default action.

To optimize performance, log either the beginning or the end of any connection, but not both.

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration. In an access control policy, the configuration may also be inherited from an ancestor policy.

Step 3 Specify where to send connection events.

Send events to the event viewer if you want to perform Firewall Management Center-based analysis on these connection events.

Step 4 Click **OK**.

Step 5 Click **Save** to save the policy.

What to do next

Deploy configuration changes; see the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Limiting logging of long URLs

End-of-connection events for HTTP traffic record the URL requested by monitored hosts. Disabling or limiting the number of stored URL characters may improve system performance. Disabling URL logging (storing zero characters) does not affect URL filtering. The system filters traffic based on requested URLs even though the system does not record them.

Procedure

Step 1 In the access control policy editor, click **Advanced**, then click **Edit** ( ) next to **General Settings**.

If **View** ( ) appears instead, the configuration is inherited from an ancestor policy, belongs to an ancestor domain, or you do not have permission to modify the configuration. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

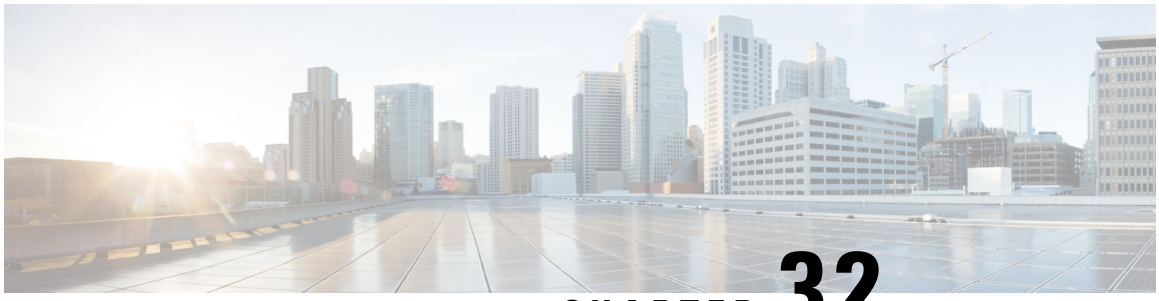
Step 2 Enter the **Maximum URL characters to store in connection events**.

Step 3 Click **OK**.

Step 4 Click **Save** to save the policy.

What to do next

Deploy configuration changes; see the *Cisco Secure Firewall Management Center Device Configuration Guide*.



CHAPTER 32

Connection and Security-Related Connection Events

The following topics describe how to use connection and security events tables.

- Connection events, on page 755
- Connection and Security Intelligence Event Fields, on page 758
- View Connection and Security Intelligence Event Tables , on page 779
- View the connection summary page, on page 783
- History for Connection and Security Intelligence Events, on page 784

Connection events

A connection event is a log that records connections that your managed devices monitor. It includes *Security Intelligence events*. These are connections that the reputation-based Security Intelligence feature blocks. It also provides granular control over which connections to log, when to log them, and where to store the data.

Connection event detection sources

Connection events generally include transactions detected by:

- Access control policies
- SSL policies
- Prefilter policies (captured by prefilter or tunnel rules)
- DNS Block lists
- URL Block lists
- Network (IP address) Block lists

Settings in rules and policies give you granular control over which connections you log, when you log them, and where you store the data.

For more information, refer to Connection Logging, on page 737.

Related Topics

Security intelligence

Connection vs. Security Intelligence events

A *Security Intelligence event* is a connection event that

- is generated whenever a session is blocked or monitored by the reputation-based Security Intelligence feature, and
- has an identical corresponding connection event that can be viewed and analyzed independently.

Event processing and storage

The system creates an identical connection event for every Security Intelligence event. You can view and analyze Security Intelligence events independently. The system stores and prunes these events separately.

The system enforces Security Intelligence before more resource-intensive evaluations. When Security Intelligence blocks a connection, the resulting event lacks information from subsequent evaluations, such as user identity.

Note In this guide, information about connection events also pertains to Security Intelligence events, unless otherwise noted.

NetFlow connections

NetFlow connections are unidirectional end-of-connection events that

- supplement connection data gathered by managed devices,
- are generated from records broadcast by NetFlow exporters, and
- are logged in the Secure Firewall Management Center database.

NetFlow connection characteristics

The system logs NetFlow records as unidirectional end-of-connection events in the Secure Firewall Management Center database. The available information for these connections differs somewhat from connections detected by your access control policy. For more information, refer to Differences between NetFlow and managed device data.

Related Topics

NetFlow data

Connection summaries (aggregated data for graphs)

Connection summaries are collections of aggregated data that

- aggregate connection data over 5-minute intervals,
- generate connection graphs and traffic profiles, and
- support custom workflows based on connection summary data.

The system aggregates connection data collected over five-minute intervals into connection summaries, which the system uses to generate connection graphs and traffic profiles. Optionally, you can create custom workflows based on connection summary data, which you use in the same way as you use workflows based on individual connection events.

To be aggregated, multiple connections must:

- represent the end of connections
- have the same source and destination IP addresses, and use the same port on the responder (destination) host
- use the same protocol (TCP or UDP)
- use the same application protocol
- either be detected by the same managed device or by the same NetFlow exporter

Each connection summary includes total traffic statistics, as well as the number of connections in the summary. NetFlow exporters generate unidirectional connections. The system increments a summary's connection count by 2 for every NetFlow-based connection.

There are no connection summaries specifically for Security Intelligence events, although the system can aggregate corresponding end-of-connection events into connection summary data.



Note Connection summaries do not contain all information associated with their aggregated connections. For example, because client information is not used to aggregate connections into connection summaries, summaries do not contain client information.

Long-running connections

A long-running connection is a monitored session that

- spans two or more five-minute intervals during which the system aggregates connection data,
- increases the connection count only when the session starts, and
- uses estimated calculations based on constant rates rather than actual data for each interval.

Connection counting and traffic calculation behavior

When the system calculates the number of packets and bytes for long-running connections, it does not report actual transmission data for each five-minute interval. Instead, it uses a constant transmission rate and calculates estimates based on the total number of packets and bytes transmitted, connection length, and the connection portion in each interval.

Combined connection summaries from external responders

A combined connection summary is a data optimization technique that

- reduces connection data storage space by aggregating summaries when one host is not on the monitored network,

- improves connection graph rendering performance by combining connections that meet summary aggregation criteria, and
- displays `external` instead of an IP address for non-monitored hosts in connection summaries and graphs.

Aggregation criteria and display behavior

To reduce the space required to store connection data and speed up the rendering of connection graphs, the system combines connection summaries when:

- one of the hosts involved in the connection is not on your monitored network
- except for the IP address of the external host, the connections in the summaries meet the summary aggregation criteria

When viewing connection summaries in the **Analysis** > **Connections** submenu pages, and when working with connection graphs, the system displays `external` instead of an IP address for the non-monitored hosts.



Note As a consequence of this aggregation, if you attempt to drill down to the table view of connection data (that is, access data on individual connections) from a connection summary or graph that involves an external responder, the table view contains no information.

Connection and Security Intelligence Event Fields



Note You cannot use the connection/Security Intelligence events Search page to search for events associated with a connection.

Access Control Policy (Syslog: ACPolicy)

The access control policy that monitored the connection.

Access Control Rule (Syslog: AccessControlRuleName)

The access control rule or default action that handled the connection, as well as up to eight Monitor rules matched by that connection.

If the connection matched one Monitor rule, the Firewall Management Center displays the name of the rule that handled the connection, followed by the Monitor rule name. If the connection matched more than one Monitor rule, the number of matching Monitor rules is displayed, for example, `Default Action + 2 Monitor Rules`.

To display a pop-up window with a list of the first eight Monitor rules matched by the connection, click ***N*** **Monitor Rules**.

Action (Syslog: AccessControlRuleAction)

The action associated with the configuration that logged the connection.

For Security Intelligence-monitored connections, the action is that of the first non-Monitor access control rule triggered by the connection, or the default action. Similarly, because traffic matching a Monitor rule

is always handled by a subsequent rule or by the default action, the action associated with a connection logged due to a Monitor rule is never Monitor. However, you can still trigger correlation policy violations on connections that match Monitor rules.

Action	Description
Allow	Connections either allowed by access control explicitly, or allowed because a user bypassed an interactive block.
Block, Block with reset	Blocked connections, including: • tunnels and other connections blocked by the prefilter policy • connections blocked by Security Intelligence. • encrypted connections blocked by an SSL policy. • connections where an exploit was blocked by an intrusion policy. • connections where a file (including malware) was blocked by a file policy. For connections where the system blocks an intrusion or file, system displays `Block`, even though you use access control Allow rules to invoke deep inspection.
Fastpath	Non-encrypted tunnels and other connections fastpathed by the prefilter policy.
Interactive Block, Interactive Block with reset	Connections logged when the system initially blocks a user's HTTP request using an Interactive Block rule. If the user clicks through the warning page that the system displays, additional connections logged for the session have an action of Allow.
Trust	Connections trusted by access control. The system logs trusted TCP connections differently depending on the device model.
Default Action	Connections handled by the access control policy's default action.
(Blank/empty)	The connection closed before enough packets had passed to match a rule. This can happen only if a facility other than access control, such as intrusion prevention, causes the connection to be logged.

Application Protocol (Syslog: ApplicationProtocol)

In the Firewall Management Center web interface, this value constrains summaries and graphs.

The application protocol, which represents communications between hosts, detected in the connection.

Application Protocol Category and Tag

Criteria that characterize the application to help you understand the application's function.

Application Risk

The risk associated with the application traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of application detected in the connection has an associated risk; this field displays the highest of those.

Business Relevance

The business relevance associated with the application traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of application detected in the connection has an associated business relevance; this field displays the lowest (least relevant) of those.

Client and Client Version (Syslog: Client, ClientVersion)

The client application and version of that client detected in the connection.

If the system cannot identify the specific client used in the connection, the field displays the word "client" appended to the application protocol name to provide a generic name, for example, FTP client.

Client Category and Tag

Criteria that characterize the application to help you understand the application's function.

Connection Counter (Syslog Only)

A counter that distinguishes one connection from another simultaneous connection. This field has no significance on its own.

The following fields collectively uniquely identify a connection event: DeviceUUID, First Packet Time, Connection Instance ID, and Connection Counter.

Connection Instance ID (Syslog Only)

The Snort instance that processed the connection event. This field has no significance on its own.

The following fields collectively uniquely identify a connection event: DeviceUUID, First Packet Time, Connection Instance ID, and Connection Counter.

ConnectionDuration (Syslog Only)

This field exists only as a syslog field; it does not exist in the Firewall Management Center web interface. (The web interface conveys this information using the First Packet and Last Packet columns.)

This field has a value only when logging occurs at the end of the connection. For a start-of-connection syslog message, this field is not output, as it is not known at that time.

For an end-of-connection syslog message, this field indicates the number of seconds between the first packet and the last packet, which may be zero for a short connection. For example, if the timestamp of the syslog is 12:34:56 and the ConnectionDuration is 5, then the first packet was seen at 12:34:51.

Connections

The number of connections in a connection summary. For long-running connections, that is, connections that span multiple connection summary intervals, only the first connection summary interval is incremented. To view meaningful results for searches using the **Connections** criterion, use a custom workflow that has a connection summary page.

Count

The number of connections that match the information that appears in each row. Note that the **Count** field appears only after you apply a constraint that creates two or more identical rows. If you create a custom workflow and do not add the **Count** column to a drill-down page, each connection is listed individually and packets and bytes are not summed.

Detection Type (Syslog: DetectionType)

This field shows the source of detection of a client application. It displays **AppID** or **Encrypted Visibility** when:

- The Encrypted Visibility Engine identifies the client application from TLS traffic, this field displays **Encrypted Visibility Engine**.
- Applications are not to process names, this field displays **AppID**, indicating that the engine that identified the client application was AppID.

Destination Port/ICMP Code (Syslog: Separate fields - DstPort, ICMPCode)

In the Firewall Management Center web interface, these values constrain summaries and graphs.

The port or ICMP code used by the session responder.

DestinationSecurityGroup (Syslog Only)

This field holds the text value associated with the numeric value in **DestinationSecurityGroupTag**, if available. If the group name is not available as a text value, then this field contains the same integer value as the DestinationSecurityGroupTag field.

DestinationSecurityGroupType (Syslog Only)

This field displays the source from which a security group tag was obtained.

Value	Description
Inline	Destination SGT value is from packet
Session Directory	Destination SGT value is from ISE via session directory topic
SXP	Destination SGT value is from ISE via SXP topic

Destination SGT (Syslog: DestinationSecurityGroupTag)

The numeric Security Group Tag (SGT) attribute of the destination involved in the connection.

The Destination SGT value is obtained from the source specified in the **DestinationSecurityGroupType** field.

Device

In the Firewall Management Center web interface, this value constrains summaries and graphs.

The managed device that detected the connection or, for connections generated from NetFlow data, the managed device that processed the data.

DeviceUUID (Syslog Only)

The unique identifier of the firewall device that generated an event.

The following fields collectively uniquely identify a connection event: DeviceUUID, First Packet Time, Connection Instance ID, and Connection Counter.

DNS Query (Syslog: DNSQuery)

The DNS query submitted in a connection to the name server to look up a domain name.

This field can also hold the domain name for URL filtering matches when DNS filtering is enabled. In this case, the URL field will be blank and the URL Category and URL Reputation fields contain the values associated with the domain.

For more information about DNS filtering, refer to DNS filtering.

DNS Record Type (Syslog: DNSRecordType)

The type of the DNS resource record used to resolve a DNS query submitted in a connection.

DNS Response (Syslog: DNSResponseType)

The DNS response returned in a connection to the name server when queried.

DNS Sinkhole Name (Syslog: DNS_Sinkhole)

The name of the sinkhole server where the system redirected a connection.

DNS TTL (Syslog: DNS_TTL)

The number of seconds a DNS server caches the DNS resource record.

Domain

The domain of the managed device that detected the connection or, for connections generated from NetFlow data, the domain of the managed device that processed the data. This field is only present if you have ever configured the Firewall Management Center for multitenancy.

Encrypted Visibility Process Name (Syslog: EncryptedVisibilityProcessName)

Process or client in the TLS client hello packet that was analyzed by the Encrypted Visibility Engine (EVE).

Encrypted Visibility Confidence Score (Syslog: EncryptedVisibilityConfidenceScore)

The confidence value in the range 0-100% that the encrypted visibility engine has detected the right process. For example, if the process name is Firefox and if the confidence score is 80%, it means that the engine is 80% confident that the process it has detected is Firefox.

Encrypted Visibility Threat Confidence (Syslog: EncryptedVisibilityThreatConfidence)

dd_EncryptedVisibilityThreatConfidenceThe probability level that the process detected by the encrypted visibility engine contains threat. This field indicates the bands (Very High, High, Medium, Low, or Very Low) based on the value in the threat confidence score.

Encrypted Visibility Threat Confidence Score (Syslog: EncryptedVisibilityThreatConfidenceScore)

The confidence value in the range 0-100% that the process detected by the encrypted visibility engine contains threat. If the threat confidence score is very high, say 90%, then the Encrypted Visibility Process Name field displays "Malware."

Endpoint Location

The IP address of the network device that used ISE to authenticate the user, as identified by ISE.

Endpoint Profile (Syslog: Endpoint Profile)

The user's endpoint device type, as identified by ISE.

Event Priority (Syslog Only)

Whether or not the connection event is a high priority event. `High` priority events are connection events that are associated with an intrusion, Security Intelligence, file, or malware event. All other events are `Low` priority.

Files (Syslog: FileCount)

The number of files (including malware files) detected or blocked in a connection associated with one or more file events.

In the Firewall Management Center web interface, the **View Files icon** links to a list of files. The number on the icon indicates the number of files (including malware files) detected or blocked in that connection.

First Packet or Last Packet (Syslog: See the ConnectionDuration field)

The date and time the first or last packet of the session was seen.

First Packet Time (Syslog Only)

The time the system encountered the first packet.

The following fields collectively uniquely identify a connection event: DeviceUUID, First Packet Time, Connection Instance ID, and Connection Counter.

HTTP Referrer (Syslog: HTTPReferer)

The HTTP referrer, which represents the referrer of a requested URL for HTTP traffic detected in the connection (such as a website that provided a link to, or imported a link from, another URL).

HTTP Response Code (Syslog: HTTPResponse)

The HTTP status code sent in response to a client's HTTP request over a connection.

Ingress/Egress Interface (Syslog: IngressInterface, EgressInterface)

The ingress or egress interface associated with the connection. If your deployment includes an asymmetric routing configuration, the ingress and egress interface may not belong to the same inline pair.

Ingress/Egress Security Zone (Syslog: IngressZone, EgressZone)

The ingress or egress security zone associated with the connection.

For rezoned encapsulated connections, the ingress field displays the tunnel zone you assigned, instead of the original ingress security zone. The egress field is blank.

Ingress Virtual Router/Egress Virtual Router (Syslog: IngressVRF, EgressVRF)

In networks using virtual routing, the names of the virtual routers through which traffic entered and exited the network.

Initiator/Responder Bytes (Syslog: InitiatorBytes, ResponderBytes)

The total number of bytes transmitted by the session initiator or received by the session responder.

Initiator/Responder Continent

When a routable IP is detected, the continent associated with the IP address for the session initiator or responder.

Initiator/Responder Country

When a routable IP is detected, the country associated with the IP address of the session initiator or responder. The system displays an icon of the country's flag, and the country's ISO 3166-1 alpha-3 country code. Hover your pointer over the flag icon to view the country's full name.

Initiator/Responder IP (Syslog: SrcIP, DstIP)

In the Firewall Management Center web interface, these values constrain summaries and graphs.

The IP address (and host name, if DNS resolution is enabled) of the session initiator or responder.

See also Initiator and responder, source and destination, and sender and receiver fields, on page 773.

In the Firewall Management Center web interface, the host icon identifies the IP address that caused the connection to be blocked.

For plaintext, passthrough tunnels either blocked or fastpathed by the prefilter policy, initiator and responder IP addresses represent the tunnel endpoints—the routed interfaces of the network devices on either side of the tunnel.

Initiator/Responder Packets (Syslog: InitiatorPackets, ResponderPackets)

The total number of packets transmitted by the session initiator or received by the session responder.

Initiator User (Syslog: User)

In the Firewall Management Center web interface, this value constrains summaries and graphs.

The user logged into the session initiator. If this field is populated with **No Authentication**, the user traffic:

- matched an access control policy without an associated identity policy
- did not match any rules in the identity policy

If applicable, the username is preceded by `<realm>\`.

See also Initiator and responder, source and destination, and sender and receiver fields, on page 773.

Intrusion Events (Syslog: IPSCount)

The number of intrusion events, if any, associated with the connection.

In the Firewall Management Center web interface, the **View Intrusion Events icon** links to a list of events.

IOC

Whether the event triggered an indication of compromise (IOC) against a host involved in the connection.

NAT Source/Destination IP (Syslog: NAT_InitiatorIP, NAT_ResponderIP)

The NAT translated IP address of the session initiator or responder.

NAT Source/Destination Port (Syslog: NAT_InitiatorPort, NAT_ResponderPort)

The NAT translated port of the session initiator or responder.

NetBIOS Domain (Syslog: NetBIOSDomain)

The NetBIOS domain used in the session.

NetFlow SNMP Input/Output

For connections generated from NetFlow data, the interface index for the interface where connection traffic entered or exited the NetFlow exporter.

NetFlow Source/Destination Autonomous System

For connections generated from NetFlow data, the border gateway protocol autonomous system number for the source or destination of traffic in the connection.

NetFlow Source/Destination Prefix

For connections generated from NetFlow data, the source or destination IP address ANDed with the source or destination prefix mask.

NetFlow Source/Destination TOS

For connections generated from NetFlow data, the setting for the type-of-service (TOS) byte when connection traffic entered or exited the NetFlow exporter.

Network Analysis Policy (Syslog: NAPPolicy)

The network analysis policy (NAP), if any, associated with the generation of the event.

Original Client Country

The country where the original client IP address belongs. To obtain this value, the system extracts the original client IP address from an X-Forwarded-For (XFF), True-Client-IP, or custom-defined HTTP header, then maps it to the country using the geolocation database (GeoDB). To populate this field, you must enable an access control rule that handles proxied traffic based on its original client.

Original Client IP (Syslog: originalClientSrcIP)

The original client IP address from an X-Forwarded-For (XFF), True-Client-IP, or custom-defined HTTP header. To populate this field, you must enable an access control rule that handles proxied traffic based on its original client.

Prefilter Policy (Syslog: Prefilter Policy)

The prefilter policy that handled the connection.

Protocol (Syslog: Protocol)

In the Firewall Management Center web interface:

- This value constrains summaries and graphs.
- This field is available only as a search field.

The transport protocol used in the connection. To search for a specific protocol, use the name or number protocol as listed in http://www.iana.org/assignments/protocol-numbers.

QoS-Applied Interface

For rate-limited connections, the name of the interface where you applied rate limiting.

QoS-Dropped Initiator/Responder Bytes

The number of bytes dropped from the session initiator or session responder due to rate limiting.

QoS-Dropped Initiator/Responder Packets

The number of packets dropped from the session initiator or session responder due to rate limiting.

QoS Policy

The QoS policy that rate limited the connection.

QoS Rule

The QoS rule that rate limited the connection.

Reason (Syslog: AccessControlRuleReason)

The reason or reasons the connection was logged, in many situations. For a full list, refer to Connection event reasons, on page 774.

Connections with a Reason of IP Block, DNS Block, and URL Block have a threshold of 15 seconds per unique initiator-responder pair. After the system blocks one of those connections, it does not generate connection events for additional blocked connections between those two hosts for the next 15 seconds, regardless of port or protocol.

Referenced Host (Syslog: ReferencedHost)

If the protocol in the connection is HTTP or HTTPS, this field displays the host name that the respective protocol was using.

SecIntMatchingIP (Syslog Only)

Which IP address matched.

Possible values: **`None`**, **`Destination`**, or **`Source`**.

Security Context (Syslog: Context)

For connections handled by ASA FirePOWER in multiple context mode, the metadata identifying the virtual firewall group through which the traffic passed.

Security Intelligence Category (Syslog: URLSICategory, DNSSICategory , IPReputationSICategory)

The name of the object that represents or contains the blocked URL, domain, or IP address in the connection. The Security Intelligence category can be the name of a network object or group, a Block list, a custom Security Intelligence list or feed, a TID category related to an observation, or one of the categories in the Intelligence Feed.

In the Firewall Management Center web interface, DNS, Network (IP address), and URL Security Intelligence connection events are combined into a single category field. In syslog messages, those events are specific by type.

Security-related connection events include security intelligence events and other connection events such as the ones that triggered intrusion or malware events. The **Security Intelligence Summary** workflow displays all the security intelligence events by their category and count. The events without a security intelligence category are grouped and displayed with the count only.

For more information about the categories in the Intelligence Feed, refer to Security intelligence categories.

Source Device

In the Firewall Management Center web interface, this value constrains summaries and graphs.

The IP address of the NetFlow exporter that broadcast the data used to generate for the connection. If the connection was detected by a managed device, this field displays `Firepower`.

Source Port/ICMP Type (Syslog: SrcPort, ICMPType)

In the Firewall Management Center web interface, these values constrain summaries and graphs.

The port or ICMP type used by the session initiator.

SourceSecurityGroup (Syslog Only)

This field holds the text value associated with the numeric value in **SourceSecurityGroupTag**, if available. If the group name is not available as a text value, then this field contains the same integer value as the SourceSecurityGroupTag field. Tags can be obtained from inline devices (no source SGT name specified) or from ISE (which specifies a source).

SourceSecurityGroupType (Syslog Only)

This field displays the source from which a security group tag was obtained.

Value	Description
Inline	Source SGT value is from packet

Value	Description
Session Directory	Source SGT value is from ISE via session directory topic
SXP	Source SGT value is from ISE via SXP topic

Source SGT (Syslog: SourceSecurityGroupTag)

The numeric representation of the Security Group Tag (SGT) attribute of the packet involved in the connection. The SGT specifies the privileges of a traffic source within a trusted network. Security Group Access (a feature of both Cisco TrustSec and Cisco ISE) applies the attribute as packets enter the network.

SSL Actual Action (Syslog: SSLActualAction)

In the Firewall Management Center web interface, this field is a search field only.

The system displays field values in the **SSL Status** field on search workflow pages.

The action the system applied to encrypted traffic in the SSL policy.

Action	Description
Block/Block with reset	Represents blocked encrypted connections.
Decrypt (Resign)	Represents an outgoing connection decrypted using a re-signed server certificate.
Decrypt (Replace Key)	Represents an outgoing connection decrypted using a self-signed server certificate with a substituted public key.
Decrypt (Known Key)	Represents an incoming connection decrypted using a known private key.
Default Action	Indicates the connection was handled by the default action.
Do not Decrypt	Represents a connection the system did not decrypt.

SSL Certificate Information (Syslog: SSLCertificate)

In the Firewall Management Center web interface, this field is a search field only.

The information stored on the public key certificate used to encrypt traffic, including:

- Subject/Issuer Common Name
- Subject/Issuer Organization
- Subject/Issuer Organization Unit
- Not Valid Before/After
- Serial Number
- Certificate Fingerprint

- Public Key Fingerprint

SSL Certificate Status (Syslog: SSLServerCertStatus)

This applies only if you configured a Certificate Status SSL rule condition. If encrypted traffic matches an SSL rule, this field displays one or more of the following server certificate status values:

- Self Signed
- Valid
- Invalid Signature
- Invalid Issuer
- Expired
- Unknown
- Not Valid Yet
- Revoked

If undecryptable traffic matches an SSL rule, this field displays `Not Checked`.

SSL Cipher Suite (Syslog: SSSLCipherSuite)

A macro value representing a cipher suite used to encrypt the connection. See https://www.iana.org/assignments/tls-parameters/tls-parameters.xhtml for cipher suite value designations.

SSL Encryption applied to the connection

This field is available only as a search field in the Secure Firewall Management Center web interface.

Enter **`yes`** or **`no`** in the **SSL** search field to view TLS/SSL-encrypted or non-encrypted connections.

SSL Expected Action (Syslog: SSLExpectedAction)

In the Firewall Management Center web interface, this field is a search field only.

The action the system expected to apply to encrypted traffic, given the SSL rule in effect.

Enter any of the values listed for SSL **Actual Action**.

SSL Failure Reason (Syslog: SSLFlowStatus)

The reason the system failed to decrypt encrypted traffic:

- Unknown
- No Match
- Success
- Uncached Session
- Unknown Cipher Suite
- Unsupported Cipher Suite
- Unsupported SSL Version
- SSL Compression Used

- Session Undecryptable in Passive Mode
- Handshake Error
- Decryption Error
- Pending Server Name Category Lookup
- Pending Common Name Category Lookup
- Internal Error
- Network Parameters Unavailable
- Invalid Server Certificate Handle
- Server Certificate Fingerprint Unavailable
- Cannot Cache Subject DN
- Cannot Cache Issuer DN
- Unknown SSL Version
- External Certificate List Unavailable
- External Certificate Fingerprint Unavailable
- Internal Certificate List Invalid
- Internal Certificate List Unavailable
- Internal Certificate Unavailable
- Internal Certificate Fingerprint Unavailable
- Server Certificate Validation Unavailable
- Server Certificate Validation Failure
- Invalid Action

Field values are displayed in the **SSL Status** field on the search workflow pages.

SSL Flow Error

The error name and hexadecimal code if an error occurred during the TLS/SSL session; `Success` if no error occurred.

SSL Flow Flags

The first ten debugging level flags for an encrypted connection. On a workflow page, to view all flags, click the ellipsis (**...**).

The message OVER_SUBSCRIBED is displayed if your managed device is overloaded. For more information, refer to Troubleshoot TLS/SSL oversubscription.

SSL Flow Messages

The keywords below indicate encrypted traffic is associated with the specified message type exchanged between client and server during the TLS/SSL handshake. See http://tools.ietf.org/html/rfc5246 for more information.

- HELLO_REQUEST
- CLIENT_ALERT
- SERVER_ALERT
- CLIENT_HELLO
- SERVER_HELLO
- SERVER_CERTIFICATE
- SERVER_KEY_EXCHANGE
- CERTIFICATE_REQUEST
- SERVER_HELLO_DONE
- CLIENT_CERTIFICATE
- CLIENT_KEY_EXCHANGE
- CERTIFICATE_VERIFY
- CLIENT_CHANGE_CIPHER_SPEC
- CLIENT_FINISHED
- SERVER_CHANGE_CIPHER_SPEC
- SERVER_FINISHED
- NEW_SESSION_TICKET
- HANDSHAKE_OTHER
- APP_DATA_FROM_CLIENT
- APP_DATA_FROM_SERVER
- SERVER_NAME_MISMATCH

  The server certificate seen in the session has a Common Name or SAN values not corresponding to the destined domain name.

- CERTIFICATE_CACHE_HIT

  A certificate matching the destined domain name was found in the cache.

- CERTIFICATE_CACHE_MISS

  A certificate matching the destined domain name was not found in the cache.

The message HEARTBEAT is displayed if applications are using the TLS/SSL heartbeat extension. For more information, refer to TLS heartbeat.

SSL Policy (Syslog: SSLPolicy)

The SSL policy that handled the connection.

If TLS server identity discovery is enabled in the access control policy advanced settings, and there is no SSL policy associated with the access control policy, this field holds `none` for all SSL events.

SSL Rule (Syslog: SSLRuleName)

The SSL rule or default action that handled the connection, as well as the first Monitor rule matched by that connection. If the connection matched a Monitor rule, the field displays the name of the rule that handled the connection, followed by the Monitor rule name.

SSLServerName (Syslog Only)

This field exists only as a syslog field; it does not exist in the Firewall Management Center web interface.

Hostname of the server with which the client established an encrypted connection.

SSL Session ID (Syslog: SSLSessionID)

The hexadecimal Session ID negotiated between the client and server during the TLS/SSL handshake.

SSL Status

The action associated with the **SSL Actual Action** (SSL rule, default action, or undecryptable traffic action) that logged the encrypted connection. The **Lock icon** links to SSL certificate details. If the certificate is unavailable (for example, for connections blocked due to TLS/SSL handshake error), the lock icon is dimmed.

If the system fails to decrypt an encrypted connection, it displays the **SSL Actual Action** (undecryptable traffic action) taken, as well as the **SSL Failure Reason**. For example, if the system detects traffic encrypted with an unknown cipher suite and allows it without further inspection, this field displays `Do Not Decrypt (Unknown Cipher Suite)`.

When searching this field, enter one or more of the **SSL Actual Action** and **SSL Failure Reason** values to view encrypted traffic the system handled or failed to decrypt.

SSL Subject/Issuer Country

This field is available only in the Firewall Management Center web interface, and only as a search field.

A two-character ISO 3166-1 alpha-2 country code for the subject or issuer country associated with the encryption certificate.

SSL Ticket ID (Syslog: SSLTicketID)

A hexadecimal hash value of the session ticket information sent during the TLS/SSL handshake.

SSLURLCategory (Syslog Only)

URL categories for the URL visited in the encrypted connection.

This field exists only as a syslog field; in the Firewall Management Center web interface, values in this field are included in the URL Category column.

See also **URL**.

SSL Version (Syslog: SSLVersion)

The TLS/SSL protocol version used to encrypt the connection:

- Unknown
- SSLv2.0
- SSLv3.0
- TLSv1.0
- TLSv1.1

- TLSv1.2
- TLSv1.3

TCP Flags (Syslog: TCPFlags)

For connections generated from NetFlow data, the TCP flags detected in the connection.

When searching this field, enter a list of comma-separated TCP flags to view all connections that have *at least* one of those flags.

Time

The ending time of the five-minute interval that the system used to aggregate connections in a connection summary. This field is not searchable.

Total Packets

This field is available only as a search field.

The total number of packets transmitted in the connection.

Traffic (KB)

This field is available only as a search field.

The total amount of data transmitted in the connection, in kilobytes.

Tunnel/Prefilter Rule (Syslog: Tunnel or Prefilter Rule)

The tunnel rule, prefilter rule, or prefilter policy default action that handled the connection.

URL, URL Category, and URL Reputation (Syslog: URL, URLCategory and SSLURLCategory, URLReputation)

The URL requested by the monitored host during the session and its associated category and reputation, if available.

For an event to display URL category and reputation, you must include the applicable URL rules in an access control policy and configure the rule with URL category and URL reputation under the **URLs** tab.

URL category and reputation do not appear in an event if the connection is processed before it matches a URL rule.

If the URL column is empty and DNS filtering is enabled, the DNS Query field shows the domain, and the URL Category and URL Reputation values apply to the domain.

If the system identifies or blocks a TLS/SSL application, the requested URL is in encrypted traffic, so the system identifies the traffic based on an SSL certificate. For TLS/SSL applications, therefore, this field indicates the common name contained in the certificate.

See also **SSLURLCategory**, above.

User Agent (Syslog: UserAgent)

The user-agent string application information extracted from HTTP traffic detected in the connection.

VLAN ID (Syslog: VLAN_ID)

The innermost VLAN ID associated with the packet that triggered the connection.

Web Application (Syslog: WebApplication)

The web application, which represents the content or requested URL for HTTP traffic detected in the connection.

If the web application does not match the URL for the event, the traffic is probably referred traffic, such as advertisement traffic. If the system detects referred traffic, it stores the referring application (if available) and lists that application as the web application.

If the system cannot identify the specific web application in HTTP traffic, this field displays `Web Browsing`.

Web Application Category and Tag

Criteria that characterize the application to help you understand the application's function.

Connection and Security Intelligence Event data overview

In the Secure Firewall Management Center web interface, you can view and search connection and Security Intelligence events using tabular and graphical workflows under the **Analysis** > **Connections** > **Security Intelligence Events**.



Note For each Security Intelligence event, there is an identical, separately stored connection event. Each Security Intelligence event has a populated **Security Intelligence Category** field.

The information available for any individual event can vary depending on how, why, and when the system logged the connection.

Search constraints

Fields marked with an asterisk (*****) on search pages constrain connection graphs and connection summaries. Because connection graphs are based on connection summaries, the same criteria that constrain connection summaries also constrain connection graphs.

If invalid search constraints are used for connection summaries, and the results are viewed in a connection summary page within a custom workflow, they are labeled as not applicable (N/A) and marked with a strikethrough.

Syslog fields

Most fields appear both in the Secure Firewall Management Center web interface and as syslog messages. Fields without a listed syslog equivalent are not available in syslog messages. A few fields are syslog-only, as noted, and few others are separate fields in syslog messages but are consolidated fields in the web interface or vice-versa.

Initiator and responder, source and destination, and sender and receiver fields

A field category is a classification of network event terminology that

- distinguishes between different roles in network connections, security events, and file transfers,
- applies to specific event types such as connection, intrusion, and file/malware events, and

- represents different aspects of network communication flow that may not always align with each other.

Field terminology comparison

These field types have different relationships based on the event context.

Table 106: Comparison of terms

Fields	Event Type	Description
Initiator/Responder	Connection	The initiator and responder of the connection. The connection initiator is not necessarily the same as the intrusion source or the malware file sender.
Source/Destination	Intrusion	The source and destination of the attack. The source of an intrusion event can be the initiator or the responder of the connection.
Sender/Receiver **(Sending..., Receiving...)**	File, Malware	The sender and receiver of a file or malware. The file sender is not necessarily the connection initiator. Files may be uploaded or downloaded.

Connection event reasons

Lists the reasons for connection events logged by the system, including descriptions for each reason.

The Reason field in a connection event displays why the connection was logged. Each reason corresponds to specific actions or circumstances in the connection event.

The table lists each connection event reason logged by the system, with a description of what triggers each reason and the resulting system action.

Reason	Description
Content Restriction	The system modified the packet to enforce content restrictions related to the Safe Search feature.
DNS Block	The system denied the connection without inspection, based on the domain name and Security Intelligence data. A reason of DNS Block is paired with an action of Block, Domain not found, or Sinkhole, depending on the DNS rule action.
DNS Monitor	The system would have denied the connection based on the domain name and Security Intelligence data, but you configured the system to monitor, rather than deny, the connection.

Reason	Description
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 larger than 1GB/10 seconds. You can adjust the byte and time thresholds for identifying elephant flows in the Firewall Threat Defense CLI using the **system support elephant-flow-detection** command. For more information, refer to the *Cisco Secure Firewall Threat Defense Command Reference*. **Note** A flow is considered as elephant flow only when both the byte and time thresholds are surpassed. You can create a custom dashboard to correlate elephant flows and other interrelated metrics, for example, CPU metrics such as Snort, System, and Physical Cores. For more information, refer to System Monitoring and Troubleshooting chapter.
File Block	The connection contained a file or malware file that the system prevented from being transmitted. File Block reason is always paired with an action of Block.
File Custom Detection	The connection contained a file on the custom detection list that the system prevented from being transmitted.
File Monitor	The system detected a particular type of file in the connection.
File Resume Allow	File transmission was originally blocked by a Block Files or Block Malware file rule. After a new access control policy allowing the file was deployed, the HTTP session automatically resumed. This reason only appears in inline deployments.
File Resume Block	File transmission was originally allowed by a Detect Files or Malware Cloud Lookup file rule. After a new access control policy blocking the file was deployed, the HTTP session automatically stopped. This reason only appears in inline deployments.
Intelligent App Bypass	The Intelligent Application Bypass (IAB) mode: • If the action is Trust, IAB was in bypass mode. Matching traffic passed without further inspection. • If the action is Allow, IAB was in test mode. Matching traffic was available for further inspection.
Intrusion Block	Snort2 Engine: 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. Snort3 Engine: When there is a "would have dropped" result, the connection event reason is blank, instead of "Intrusion block." The "would have dropped" event is treated the same as "Allow" in regards to the connection event reason being populated.
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.

Reason	Description
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.
IP Monitor	The system would have denied the connection based on the IP address and Security Intelligence data, but you configured the system to monitor, rather than deny, the connection.
SSL Block	The system blocked an encrypted connection based on the TLS/SSL inspection configuration. A reason of SSL Block is always paired with an action of Block.
URL Block	The system denied the connection without inspection, based on the URL and Security Intelligence data. A reason of URL Block is always paired with an action of Block.
URL Monitor	The system would have denied the connection based on the URL and Security Intelligence data, but you configured the system to monitor, rather than deny, the connection.
User Bypass	The system initially blocked a user's HTTP request, but the user clicked through a warning page to view the site. A reason of User Bypass is always paired with an action of Allow.

Requirements for populating connection event fields

A connection event field population requirement is a system condition that

- determines what information appears in connection events based on multiple system factors,
- depends on appliance models, licensing, traffic characteristics, and detection methods, and
- varies according to evaluation stage, logging method, and specific system configurations.

Factors affecting connection event field population

Several factors determine which fields are populated in connection events:

Appliance Model and License

Many features require that you enable specific licensed capabilities on target devices, and many features are only available on some models.

Traffic Characteristics

The system reports only information that is present and detectable in network traffic. For example, there could be no user associated with an initiator host, or no referenced host detected in a connection where the protocol is not DNS, HTTP, or HTTPS.

Origin/Detection Method: Traffic-Based Detection vs NetFlow

With the exception of NetFlow-only fields, the information available in NetFlow records is more limited than the information generated by traffic-based detection; refer to Differences between NetFlow and managed device data.

Evaluation Stage

Each type of traffic inspection and control occurs where it provides maximum flexibility and performance.

For example, the system enforces Security Intelligence before more resource-intensive evaluations. When a connection is blocked by Security Intelligence, the resulting event does not contain the information that the system would have gathered from subsequent evaluation, for example, user identity.

Logging Method: Beginning or End of Connection

When the system detects a connection, whether you can log it at its beginning or its end (or both) depends on how you configure the system to detect and handle it.

Beginning-of-connection events do not have information that must be determined by examining traffic over the duration of the session; for example, the total amount of data transmitted or the timestamp of the last packet in the connection. Beginning-of-connection events are also not guaranteed to have information about application or URL traffic in the session, and do not contain any details about the session's encryption. Beginning-of-connection logging is usually the only option for blocked connections.

Connection Event Type: Individual vs Summary

Connection summaries do not contain all of the information associated with their aggregated connections. For example, because client information is not used to aggregate connections into connection summaries, summaries do not contain client information.

Keep in mind that connection graphs are based on connection summary data, which use only end-of-connection logs. Configure your system to log end-of-connection data for connection graphs and summaries to display data.



Note Security-related connection events include security intelligence events and other connection events, such as the ones that triggered intrusion or malware events. The **Security Intelligence Summary** workflow groups the security-related connection events that do not have a security intelligence category and displays the count without a **Security Intelligence Category** value.

Other Configurations

Other configurations that affect connection logging include, but are not limited to:

- ISE-related fields are populated only if you configure ISE, in connections associated with users who authenticate via an Active Directory domain controller. Connection events do not contain ISE data for users who authenticate via LDAP, RADIUS, or RSA domain controllers.
- The Security Group Tag (SGT) fields are populated only if you configure ISE as an identity source or add custom SGT rule conditions.
- Prefilter-related fields (including tunnel zone information in security zone fields) are populated only in connections handled by a prefilter policy.
- TLS/SSL-related fields are populated only in encrypted connections handled by an SSL policy. You can view the values of the fields using a Do Not Decrypt rule action if you do not need to decrypt the traffic.
- File information fields are populated only in connections logged by access control rules associated with file policies.
- Intrusion information fields are populated only in connections logged by access control rules either associated with intrusion policies or using the default action.
- QoS-related fields are populated only in connections subject to rate limiting.

- The Reason field is populated only in specific situations, such as when a user bypasses an Interactive Block configuration.
- The Domain field is only present if you have ever configured the Secure Firewall Management Center for multitenancy.

An advanced setting in the access control policy controls the number of characters the system stores in the connection log for each URL requested by monitored hosts in HTTP sessions. If you use this setting to disable URL logging, the system does not display individual URLs in the connection log, although you can still view category and reputation data, if it exists.

For the connection event to display URL category and reputation, you must include the applicable URL rules in an access control policy and configure the rule with URL category and URL reputation under the **URLs** tab. URL category and reputation do not appear in an event if the connection is processed before it matches a URL rule.

Related Topics

Differences between NetFlow and managed device data

Information available in connection event fields

The table lists the conditions in which the system populates connection and security intelligence fields. Columns represent different event types:

- Origin: Direct (events representing connections detected and handled by a system-managed device)
- Origin: NetFlow (events representing connections exported by a NetFlow exporter)
- Logging: Start (events representing connections logged at their beginning)
- Logging: End (events representing connections logged at their end)

A "yes" in the table means the system can populate a field but only reports detectable information in network traffic. For example, TLS/SSL-related fields populate only for encrypted connections handled by an SSL policy.

Connection Event Field	Origin: Direct	Origin: NetFlow	Logging: Start	Logging: End
Access Control Policy	yes	no	yes	yes
Access Control Rule	yes	no	yes	yes
Action	yes	no	yes	yes
Application Protocol	yes	yes	if available	yes
Application Protocol Category & Tag	yes	yes	if available	yes
Client Application	yes	no	if available	yes
Client Application Category & Tag	yes	no	if available	yes
Payload Application	yes	no	if available	yes

Connection Event Field	Origin: Direct	Origin: NetFlow	Logging: Start	Logging: End
Payload Application Category & Tag	yes	no	if available	yes
Web Application	yes	no	if available	yes
Web Application Category & Tag	yes	no	if available	yes

View Connection and Security Intelligence Event Tables

You can use the Firewall Management Center to view a table of connection or Security Intelligence events. Then, you can manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access connection graphs differs depending on the workflow you use. You can use a predefined workflow, which terminates in a table view of events. You can also create a custom workflow that displays only the information that matches your specific needs.

When you are using a connection or Security Intelligence workflow table, you can perform many common actions.

Note that when you constrain connection events on a drill-down page, the packets and bytes from identical events are summed. However, if you are using a custom workflow and did not add a **Count** column to a drill-down page, the events are listed individually and packets and bytes are not summed.

Note that **Connection Events** table view displays **1 of Many** instead of how many pages of events are available if your system generates more than 25 connection events.

Before you begin

You must be an Admin or Security Analyst user to perform this task.

Follow these steps to view and interact with connection and security intelligence event tables.

Procedure

Step 1 Choose either of the following:

- **Analysis** > **Connections** > **Events** (for connection events)
- **Analysis** > **Connections** > **Security Intelligence Events**

Note
If a connection graph appears instead of a table, click **(switch workflow)** by the workflow title, and choose the predefined **Connection Events** workflow, or a custom workflow. Note that all predefined connection event workflows, including connection graphs, terminate in a table view of connections.

Step 2 You have the following choices:

- Time Range: To adjust the time range, which is useful if no events appear, refer to Change the time window, on page 695.

- Data Source: If data is stored remotely using Security Analytics and Logging (On Premises), and you have good reason to change the data source, choose a data source. For important information about this option, refer to Work in Firewall Management Center with connection events stored on a Secure Network Analytics appliance, on page 685.
- Field Names: To learn more about the contents of the columns in the table, refer to Connection and Security Intelligence Event Fields, on page 758.

Note
In the table view of events, multiple fields are hidden by default. To change the fields that appear, click the Disable Column in any column name to display a field chooser.

- Additional information: To view data in available sources external to your system, right-click an event value. The options you see depend on the data type and include public sources; other sources depend on the resources you have configured. For more information, refer to Event investigation using web-based resources, on page 638.
- External intelligence: To gather intelligence about an event, right-click an event value in the table and choose from a Cisco or third-party intelligence source. For example, you can get details about a suspicious IP address from Cisco Talos. The options you see depend on the data type and the integrations that are configured on your system. For more information, refer to Event investigation using web-based resources, on page 638.
- Host Profile: To view the host profile for an IP address, click **Host Profile** or, for hosts with active indications of compromise (IOC) tags, **Compromised Host** that appears next to the IP address.
- User Profile: To view user identity information, click the user icon that appears next to the **User Identity**, or for users associated with IOCs, **Red User**.
- Files and Malware: To view the files, including malware, detected or blocked in a connection, click **View Files** and proceed as described in View files and malware detected in a connection, on page 781.
- Intrusion Events: To view the intrusion events associated with a connection, as well as their priority and impact, click **Intrusion Events** in the **Intrusion Events** column and proceed as described in View intrusion events for a connection, on page 782.

Note
To quickly view intrusion, file, or malware events associated with one or more connections, check the connections using the check boxes in the table, then choose the appropriate option from the **Jump to** drop-down list. Note that because they are blocked before access control rule evaluation, there can be no files or intrusions associated with connections blocked by Security Intelligence. You can only see this information for a Security Intelligence event if you configured Security Intelligence to monitor, rather than block, connections.

- Certificate: To view details about an available certificate used to encrypt a connection, click **Enabled Lock** in the **SSL Status** column.
- Constrain: To constrain the columns that appear, click **Close** ( ) in the column heading that you want to hide. In the pop-up window that appears, click **Apply**.

Note
To hide or show other columns, check or clear the appropriate check boxes before you click **Apply**. To add a disabled column back to the view, expand the search constraints, then click the column name under Disabled Columns.

- Delete Events: (Security Intelligence event tables only). To delete some or all items in the current constrained view, check the check boxes next to items you want to delete and click **Delete** or click **Delete All**.

- Drill Down: For more information, refer to Use drill-down pages, on page 684.

  Note
  To drill down using one of several Monitor rules that matched a logged connection, click an *N* **Monitor Rules** value. In the pop-up window that appears, click the Monitor rule you want to use to constrain connection events.

- Navigate This Page: For more information, refer to Workflow page traversal tools, on page 681.
- Navigate Between Pages: To navigate between pages in the current workflow, keeping the current constraints, click the appropriate page link at the top left of the workflow page.
- Navigate Between Event Views: To navigate to other event views to view associated events, click **Jump to** and choose the event view from the drop-down list.
- Sort: To sort data in a workflow, click the column title. Click the column title again to reverse the sort order.

Related Topics

Workflows, on page 667
Configure event view settings, on page 234

View files and malware detected in a connection

View file events associated with connections that match access control rules with file policies to identify detected files and malware.

If you associate a file policy with one or more access control rules, the system can detect files (including malware) in matching traffic. Use the **Analysis** > **Connections** > **Events** menu options to see the file events, if any, associated with the connections logged by those rules. Instead of a list of files, the Secure Firewall Management Center displays view files () in the **Files** column. The number on the view files indicates the number of files (including malware files) detected or blocked in that connection.

Not all file and malware events are associated with connections. Specifically:

- Malware events detected by Secure Endpoint ("endpoint-based malware events") are not associated with connections. Those events are imported from your Secure Endpoint deployment.
- Many IMAP-capable email clients use a single IMAP session, which ends only when the user exits the application. Although long-running connections are logged by the system, files downloaded in the session are not associated with the connection until the session ends.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

You must be an Admin or Security Analyst user to perform this task.

Follow these steps to view files and malware detected in a connection:

Procedure

Step 1 Go to **Analysis** > **Connections** > **Events** and choose the relevant option.

Step 2 While using a connection event table, click **View Files**.
A pop-up window appears with a list of the files detected in the connection as well as their types, and if applicable, their malware dispositions.

Step 3 You have the following choices:

- View: To view a table view of file events, click a **File's View**.
- View: To view details in a table view of malware events, click a **Malware File's View**.
- Track: To track the file's transmission through your network, click a **File's Trajectory**.
- View: To view details on all of the connection's detected file or malware events detected by Malware Defense ("network-based malware events"), click **View File Events** or **View Malware Events**.

Related Topics

Workflows, on page 667
Configure event view settings, on page 234

View intrusion events for a connection

Enable detailed analysis of intrusion events linked to active or logged connections, including their priority and the potential impact to your network.

If you associate an intrusion policy with an access control rule or default action, the system can detect exploits in matching traffic. Use the **Analysis** > **Connections** > **Events** menu options to see the intrusion events, if any, associated with logged connections, as well as their priority and impact.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

You must be an Admin or Security Analyst user to perform this task.

Follow these steps to view intrusion events for a connection.

Procedure

Step 1 Go to **Analysis** > **Connections** > **Events** and choose the relevant option.

Step 2 While using a connection event table, click **Intrusion Events** in the **Intrusion Events** column.

Step 3 In the pop-up window that appears:

- Click a **Listed Event's View** to view details in the packet view.

- Click **View Intrusion Events** to view details on all of the connection's associated intrusion events.

Related Topics

Workflows, on page 667
Configure event view settings, on page 234

Encrypted connection certificate details

You can use options under the **Analysis** > **Connections** menu to display the public key certificate (if available) used to encrypt a connection handled by the system. The certificate contains the information listed in this reference.

The table lists and describes the certificate attributes you can view when displaying the public key certificate for encrypted connections handled by the system.

Table 107: Encrypted connection certificate details

Attribute	Description
Subject/Issuer Common Name	The host and domain name of the certificate subject or certificate issuer.
Subject/Issuer Organization	The organization of the certificate subject or certificate issuer.
Subject/Issuer Organization Unit	The organizational unit of the certificate subject or certificate issuer.
Not Valid Before/After	The dates when the certificate is valid.
Serial Number	The serial number assigned by the issuing CA.
Certificate Fingerprint	The SHA hash value used to authenticate the certificate.
Public Key Fingerprint	The SHA hash value used to authenticate the public key contained within the certificate.

Related Topics

Workflows, on page 667
Configure event view settings, on page 234

View the connection summary page

Display summary graphs of network connections to analyze activity by time interval, device, or other criteria.

The connection summary page provides aggregated graphs that help monitor network activity. Only users with specific role permissions can access this page. Graphs are organized by criteria such as connection count over chosen intervals; however, you cannot drill down into individual connection events from these summaries.

You can perform almost all the same actions on connection summary graphs that you can perform on connection graphs. However, because the graphs on the connection summary page are based on aggregated data, you cannot examine the individual connection events on which the graphs are based. In other words, you cannot drill down to a connection data table view from a connection summary graph.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Overview** > **Dashboards heading** > **Dashboard**.

Step 2 Click **(switch dashbaord)**, and then choose **Connection Summary**.

Step 3 From the **Select Device** list, choose the device whose summary you want to view, or choose **All** to view a summary of all devices.

Step 4 To manipulate and analyze the connection graphs, proceed as described in Use connection event graphs, on page 687.

Tip
To detach a connection graph so you can perform further analysis without affecting the default time range, click **View**.

History for Connection and Security Intelligence Events

Documents the history and details of connection and security intelligence events, including feature updates and changes across various versions.

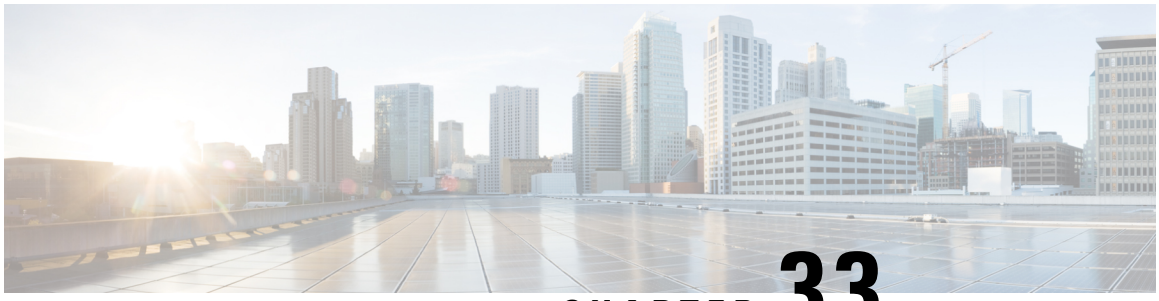
Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
New Connection Event Reason - Elephant Flow.	7.1	Any	Refer to Connection event reasons, on page 774.
NAT Translated IP Address and Port	7.1	Any	Four new fields are added to the connection and security intelligence event table: • NAT Source IP • NAT Destination IP • NAT Source Port • NAT Destination Port
Ability to choose a data source when working with certain events stored remotely	7.0	Any	Refer to History for workflows, on page 704.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
DNS filtering	7.0 6.7 (Beta feature)	Any	When DNS filtering is enabled: • The DNS Query field may hold the domain associated with DNS filtering matches. • If the URL field is empty but DNS Query, URL Category, and URL Reputation have values, the event was generated by the DNS filtering feature, and the category and reputation apply to the domain specified in DNS Query. • Refer to *DNS Filtering and Events* in the *Cisco Secure Firewall Management Center Device Configuration Guide* .
Removal of support for custom tables for connection events	6.6	Any	You can no longer create custom tables for connection events. If you upgrade, any pre-existing custom tables for connection events are still available but always return no results. There is no change to other types of custom tables. New/Modified screens: The Tables option on **Analysis** > **Advanced Custom Tables** Platform: Firewall Management Center
Removal of ability to Delete and Delete All connection events	6.6	Any	The Delete and Delete All buttons have been removed from connection events table pages. To purge all connection events, refer to Data Purge and Storage, on page 521. New/Modified screens: **Analysis** > **Connections** > **Events** Platform: Firewall Management Center
New fields for VRF and SGT	6.6	Any	• Ingress Virtual Router (Syslog: IngressVRF) • Egress Virtual Router (Syslog: EgressVRF) • DestinationSecurityGroupType (Syslog only) • SourceSecurityGroupType (Syslog only)

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details

New and changed Security Group Tag fields	6.5	Any	Changes to fields in the Firewall Management Center web interface:
			• Changed fields: Security Group Tag is now Source SGT
			• New fields: Destination SGT
			
			Changes to syslog fields:
			• Changed fields:
			SecurityGroup is now SourceSecurityGroupTag
			• New fields:
			  • SourceSecurityGroup
			  • DestinationSecurityGroup
			  • DestinationSecurityGroupTag
			
			Supported Platforms: Firewall Management Center, managed devices

New syslog field: Event Priority	6.5	Any	This field identifies connection events as High priority when they are associated with intrusion, file, malware, or Security Intelligence events.
Unique identifier for connection event in syslogs	6.4.0.4	Any	The following syslog fields collectively uniquely identify a connection event: DeviceUUID, First Packet Time, Connection Instance ID, and Connection Counter.



CHAPTER 33

Intrusion Events

The following topics describe how to work with intrusion events.

- Intrusion events, on page 787
- Tools for reviewing and evaluating intrusion events, on page 788
- License requirements for intrusion events, on page 788
- Requirements and prerequisites for intrusion events, on page 789
- View intrusion events, on page 789
- Intrusion event workflow pages, on page 807
- View intrusion event statistics, on page 826
- View intrusion event performance graphs, on page 828
- View intrusion event graphs, on page 832
- History for intrusion events, on page 832

Intrusion events

An intrusion event is a security record that

- documents the date, time, and type of exploit when the system identifies a possible intrusion,
- includes contextual information about the source of the attack and its target,
- records a copy of the packet or packets that triggered the event for packet-based events, and
- is transmitted from managed devices to the Secure Firewall Management Center for aggregated analysis.

Network monitoring capabilities

The system can help you monitor your network for traffic that could affect the availability, integrity, and confidentiality of a host and its data. By placing managed devices on key network segments, you can examine the packets that traverse your network for malicious activity. The system has several mechanisms it uses to look for the broad range of exploits that attackers have developed.

When the system identifies a possible intrusion, it generates an *intrusion event* (sometimes called by a legacy term, "IPS event"), which is a record of the date, time, the type of exploit, and contextual information about the source of the attack and its target. For packet-based events, a copy of the packet or packets that triggered the event is also recorded. Managed devices transmit their events to the Secure Firewall Management Center

where you can view the aggregated data and gain a greater understanding of the attacks against your network assets.

You can also deploy a managed device as an inline, switched, or routed intrusion system, which allows you to configure the device to drop or replace packets that you know to be harmful.

Tools for reviewing and evaluating intrusion events

Tools for reviewing and evaluating intrusion events are a collection of features that

- provide overview capabilities for current activity on managed devices,
- enable automated alerting and correlation responses to specific intrusion events, and
- support both text-based and graphical reporting with customizable scheduling options.

Available intrusion event tools

These tools help you review intrusion events and evaluate their importance:

- An event summary page that gives you an overview of the current activity on your managed devices
- Text-based and graphical reports that you can generate for any time period you choose; you can also design your own reports and configure them to run at scheduled intervals
- An incident-handling tool that you can use to gather event data related to an attack; you can also add notes to help you track your investigation and response
- Automated alerting that you can configure for SNMP, email, and syslog
- Automated correlation policies that you can use to respond to and remediate specific intrusion events
- Predefined and custom workflows that you can use to drill down through the data to identify the events that you want to investigate further
- External tools for managing and analyzing data. You can send data to those tools using syslog or eStreamer. For more information, refer to Event Analysis Using External Tools, on page 635

Additionally, you can use publicly-available information such as the predefined resources on the **Analysis** > **Advanced** > **Contextual Cross-Launch** page to learn more about malicious entities.

To search for a particular message string and retrieve documentation for the rule that generated an event, refer to https://www.snort.org/rule_docs/.

License requirements for intrusion events

This reference provides license requirement information for configuring intrusion events in your security deployment.

Threat Defense License

IPS

Classic License

Protection

Requirements and prerequisites for intrusion events

Model support

Any

Supported domains

Any

User roles

- Admin
- Intrusion Admin

View intrusion events

The initial intrusion events view differs depending on the workflow you use to access the page. You can use one of the predefined workflows, which includes one or more drill-down pages, a table view of intrusion events, and a terminating packet view, or you can create your own workflow. You can also view workflows based on custom tables, which may include intrusion events.

An event view may be slow to display if it contains a large number of IP addresses and you have enabled the **Resolve IP Addresses** event view setting.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Intrusions** > **Events**.

Step 2 You have these choices:

- Adjust time range: Adjust the time range for the event view as described in Change the time window, on page 695.
- Change workflows: If you are using a custom workflow that does not include the table view of intrusion events, choose any of the system-provided workflows by clicking **(switch workflow)** next to the workflow title.
- Constrain: To narrow your view to the intrusion events that are important to your analysis, refer to Use intrusion event workflows, on page 808.
- Delete event: To delete an event from the database, click **Delete** to delete the event whose packet you are viewing or click **Delete All** to delete all the events whose packets you previously selected.

- Mark reviewed: To mark intrusion events reviewed, refer to Mark intrusion events reviewed, on page 803.
- View connection data: To view connection data associated with intrusion events, refer to View connection data associated with intrusion events, on page 803.
- View contents: To view the contents of the columns in the table as described in Intrusion event fields, on page 790.

Intrusion event data overview

An intrusion event captures the date, time, and type of exploit when the system identifies a possible intrusion and provides contextual information about the source of the attack and its target. It also includes a copy of the packet or packets that triggered the event for packet-based events.

You can view intrusion event data in the Secure Firewall Management Center web interface at **Analysis** > **Intrusions** > **Events** or emit data from certain fields as syslog messages for consumption by an external tool. Syslog fields are indicated in the list below; fields without a listed syslog equivalent are not available in syslog messages.

When searching intrusion events, keep in mind that the information available for any individual event can vary depending on how, why, and when system logged the event. For example, only intrusion events triggered on decrypted traffic contain TLS/SSL information.



Note In the Secure Firewall Management Center web interface, some fields in the table view of intrusion events are disabled by default. To enable a field for the duration of your session, expand the search constraints, then click the column name under **Disabled Columns**.

Intrusion event fields

This reference provides field definitions for intrusion events generated by security systems. These fields contain critical information for monitoring, analyzing, and responding to security threats.

Access Control Policy (Syslog: ACPolicy)

The access control policy associated with the intrusion policy where the intrusion, preprocessor, or decoder rule that generated the event is enabled.

Access Control Rule (Syslog: AccessControlRuleName)

The access control rule that invoked the intrusion policy that generated the event. `Default Action` indicates that the intrusion policy where the rule is enabled is not associated with a specific access control rule but, instead, is configured as the default action of the access control policy.

This field is empty (or, for syslog messages, omitted)if there is:

- No associated rule/default action: Intrusion inspection was associated with neither an access control rule nor the default action, for example, if the packet was examined by the intrusion policy specified to handle packets that must pass before the system can determine which rule to apply. (This policy is specified in the Advanced tab of the access control policy.)

- No associated connection event: The connection event logged for the session has been purged from the database, for example, if connection events have higher turnover than intrusion events.

Application Protocol (Syslog: ApplicationProtocol)

The application protocol, if available, which represents communications between hosts detected in the traffic that triggered the intrusion event.

Application Protocol Category and Tag

Criteria that characterize the application to help you understand the application's function.

Application Risk

The risk associated with detected applications in the traffic that triggered the intrusion event: Very High, High, Medium, Low, and Very Low. Each type of application detected in a connection has an associated risk; this field displays the highest risk of those.

Business Relevance

The business relevance associated with detected applications in the traffic that triggered the intrusion event: Very High, High, Medium, Low, and Very Low. Each type of application detected in a connection has an associated business relevance; this field displays the lowest (least relevant) of those.

Classification (Syslog: Classification)

The classification where the rule that generated the event belongs.

See a list of possible classification values in Intrusion Event Details.

When searching this field, enter the classification number, or all or part of the classification name or description for the rule that generated the events you want to view. You can also enter a comma-separated list of numbers, names, or descriptions. Finally, if you add a custom classification, you can also search using all or part of its name or description.

Client (Syslog: Client)

The client application, if available, which represents software running on the monitored host detected in the traffic that triggered the intrusion event.

Client Category and Tag

Criteria that characterize the application to help you understand the application's function.

Connection Counter (Syslog Only)

A counter that distinguishes one connection from another simultaneous connection. This field has no significance on its own.

The following fields collectively uniquely identify the connection event associated with a particular intrusion event: DeviceUUID, First Packet Time, Connection Instance ID, and Connection Counter.

Connection Instance ID (Syslog Only)

The Snort instance that processed the connection event. This field has no significance on its own.

The following fields collectively uniquely identify the connection event associated with a particular intrusion event: DeviceUUID, First Packet Time, Connection Instance ID, and Connection Counter.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows. This field is not searchable.

CVE ID

This field is a search field only.

Search by the identification number associated with the vulnerability in MITRE's Common Vulnerabilities and Exposures (CVE) database (https://cve.mitre.org/).

Destination Continent

The continent of the receiving host involved in the intrusion event.

Destination Country

The country of the receiving host involved in the intrusion event.

Destination Host Criticality

The destination host criticality (value of the Host Criticality attribute for that corresponding host) when the event is generated.

Keep in mind that this field is not updated when the criticality of the host changes. However, new events will have the new criticality value.

Destination IP (Syslog: DstIP)

The IP address used by the receiving host involved in the intrusion event.

See also Initiator and responder, source and destination, and sender and receiver fields, on page 773.

Destination Port / ICMP Code (Syslog: DstPort, ICMPCode)

The port number for the host receiving the traffic. For ICMP traffic, where there is no port number, this field displays the ICMP code.

Destination User

The username associated with the Responder IP of the connection event. This host may or may not be the host receiving the exploit. This value is typically known only for users on your network.

.

See also Initiator and responder, source and destination, and sender and receiver fields, on page 773.

Device

The managed device where the access control policy was deployed.

DeviceUUID (Syslog Only)

The unique identifier of the firewall device that generated an event.

The following fields collectively uniquely identify the connection event associated with a particular intrusion event: DeviceUUID, First Packet Time, Connection Instance ID, and Connection Counter.

Domain

The domain of the device that detected the intrusion. This field is only present if you have ever configured the Firewall Management Center for multitenancy.

Egress Interface (Syslog: EgressInterface)

The egress interface of the packet that triggered the event. This interface column is not populated for a passive interface.

Egress Security Zone (Syslog: EgressZone)

The egress security zone of the packet that triggered the event. This security zone field is not populated in a passive deployment.

Egress Virtual Router

In networks using virtual routing, the name of the virtual router through which traffic exited the network.

Email Attachments

The MIME attachment file name that was extracted from the MIME Content-Disposition header. To display attachment file names, you must enable the SMTP preprocessor **Log MIME Attachment Names** option. Multiple attachment file names are supported.

Email Headers

This field is a search field only.

The data that was extracted from the email header.

To associate email headers with intrusion events for SMTP traffic, you must enable the SMTP preprocessor **Log Headers** option.

Email Recipient

The address of the email recipient that was extracted from the SMTP RCPT TO command. To display a value for this field, you must enable the SMTP preprocessor **Log To Addresses** option. Multiple recipient addresses are supported.

Email Sender

The address of the email sender that was extracted from the SMTP MAIL FROM command. To display a value for this field, you must enable the SMTP preprocessor **Log From Address** option. Multiple sender addresses are supported.

First Packet Time (Syslog Only)

The time the system encountered the first packet.

The following fields collectively uniquely identify the connection event associated with a particular intrusion event: DeviceUUID, First Packet Time, Connection Instance ID, and Connection Counter.

Generator

The component that generated the event.

See also information about the following intrusion event fields: GID, Message, and Snort ID.

GID (Syslog Only)

Generator ID; the ID of the component that generated the event.

See also information about the following intrusion event fields: Generator, Message, and Snort ID.

HTTP Hostname

The host name, if present, that was extracted from the HTTP request Host header. Note that request packets do not always include the host name.

To associate host names with intrusion events for HTTP client traffic, you must enable the HTTP Inspect preprocessor **Log Hostname** option.

In table views, this column displays the first fifty characters of the extracted host name. You can hover your pointer over the displayed portion of an abbreviated host name to display the complete name, up to 256 bytes. You can also display the complete host name, up to 256 bytes, in the packet view.

HTTP Response Code (Syslog: HTTPResponse)

The HTTP status code sent in response to a client's HTTP request over the connection that triggered the event.

HTTP URI

The normalized URI, if present, associated with the HTTP request packet that triggered the intrusion event. Note that request packets do not always include a URI.

To associate URIs with intrusion events for HTTP traffic, you must enable the HTTP Inspect preprocessor **Log URI** option.

To see the associated HTTP URI in intrusion events triggered by HTTP responses, you should configure HTTP server ports in the **Perform Stream Reassembly on Both Ports** option; note, however, that this increases resource demands for traffic reassembly.

This column displays the first fifty characters of the extracted URI. You can hover your pointer over the displayed portion of an abbreviated URI to display the complete URI, up to 2048 bytes. You can also display the complete URI, up to 2048 bytes, in the packet view.

Impact

The impact level in this field indicates the correlation between intrusion data, network discovery data, and vulnerability information.

When searching this field, do not specify impact icon colors or partial strings. For example, do not use `blue`, `level 1`, or `0`. Valid case-insensitive values are:

- Impact 0, Impact Level 0
- Impact 1, Impact Level 1

- Impact 2, Impact Level 2
- Impact 3, Impact Level 3
- Impact 4, Impact Level 4
- Impact 5, Impact Level 5

Because no operating system information is available for hosts added to the network map from NetFlow data, the system cannot assign Vulnerable (impact level 1: red) impact levels for intrusion events involving those hosts. In such cases, use the host input feature to manually set the operating system identity for the hosts.

Ingress Interface (Syslog: IngressInterface)

The ingress interface of the packet that triggered the event. Only this interface column is populated for a passive interface.

Ingress Security Zone (Syslog: IngressZone)

The ingress security zone or tunnel zone of the packet that triggered the event. Only this security zone field is populated in a passive deployment.

Ingress Virtual Router

In networks using virtual routing, the name of the virtual router through which traffic entered the network.

Inline Result (Syslog: InlineResult)

In workflow and table views, this field displays one of the following:

Table 108: Inline Result Field Contents in Workflow and Table Views

This Icon	Indicates
	The system dropped the packet that triggered the rule.
	IPS would have dropped the packet if you enabled the **Drop when Inline** intrusion policy option (in an inline deployment), or if a Drop and Generate rule generated the event while the system was pruning.
	IPS may have transmitted or delivered the packet to the destination, but the connection that contained this packet is now blocked.
No icon (blank)	The triggered rule was not set to Drop and Generate Events

The following table lists the possible reasons for the inline results — Would have dropped and Partially dropped.

Table 109:

Inline Result	Reason	Detailed Reason
Would Have Dropped	Interface in Passive or Tap mode	You have configured the interfaces in inline tap or passive mode.
	Intrusion Policy in "Detection" Inspection Mode	You have set the inspection mode in the intrusion policy to Detection.
	Connection Timed Out	The Snort inspection engine has suspended the inspection as the TCP/IP connection timed out.
Partially Dropped	Connection Closed (0x01)	While creating a new flow, if the allocated flows are more than the allowed number of flows, the Snort inspection engine prunes the least recently used flows.
	Connection Closed (0x02)	When reloading the Snort inspection engine causes a memory adjustment, the engine prunes the least recently used flows.
	Connection Closed (0x04)	When the Snort inspection engine is gracefully shutting down, the engine purges all the active flows.

In a passive deployment, the system does not drop packets, including when an inline interface is in tap mode, regardless of the rule state or the inline drop behavior of the intrusion policy.

When searching this field, enter either of the following:

- **`dropped`** to specify whether the packet is dropped in an inline deployment.
- **`would have dropped`** to specify whether the packet would have dropped if the intrusion policy had been set to drop packets in an inline deployment.
- **`partially dropped`** to specify whether the packet is transmitted or delivered to the destination, but the connection that contained this packet is now blocked.

Intrusion Policy (Syslog: IntrusionPolicy)

The intrusion policy where the intrusion, preprocessor, or decoder rule that generated the event was enabled. You can choose an intrusion policy as the default action for an access control policy, or you can associate an intrusion policy with an access control rule.

IOC (Syslog: NumIOC)

Whether the traffic that triggered the intrusion event also triggered an indication of compromise (IOC) for a host involved in the connection.

When searching this field, specify **`triggered`** or **`n/a`**.

Message (Syslog: Message)

The explanatory text for the event. For rule-based intrusion events, the event message is pulled from the rule. For decoder- and preprocessor-based events, the event message is hard coded.

The Generator and Snort IDs (GID and SID) and the SID version (Revision) are appended in parentheses to the end of each message in the format of numbers separated by colons (GID:SID:version). For example `(1:36330:2)`.

MITRE

A count of techniques that you can click to bring up a modal, which shows the full list of MITRE tactics and techniques within that hierarchy.

MPLS Label (Syslog: MPLS_Label)

The Multiprotocol Label Switching label associated with the packet that triggered the intrusion event.

Network Analysis Policy (Syslog: NAPPolicy)

The network analysis policy, if any, associated with the generation of the event.

This field displays the first fifty characters of the extracted URI. You can hover your pointer over the displayed portion of an abbreviated URI to display the complete URI, up to 2048 bytes. You can also display the complete URI, up to 2048 bytes, in the packet view.

Original Client IP

The original client IP address that was extracted from an X-Forwarded-For (XFF), True-Client-IP, or custom-defined HTTP header.

To display a value for this field, you must enable the HTTP preprocessor **Extract Original Client IP Address** option in the network analysis policy. Optionally, in the same area of the network analysis policy, you can also specify up to six custom client IP headers, as well as set the priority order in which the system selects the value for the Original Client IP event field.

Priority (Syslog: Priority)

The event priority as determined by the Talos Intelligence Group. The priority corresponds to either the value of the `priority` keyword or the value for the `classtype` keyword. For other intrusion events, the priority is determined by the decoder or preprocessor. Valid values are high, medium, and low.

Protocol (Syslog: Protocol)

In the Secure Firewall Management Center web interface, this field is a search field only.

The name or number of the transport protocol used in the connection as listed in http://www.iana.org/assignments/protocol-numbers. This is the protocol associated with the source and destination port/ICMP column.

Reviewed By

The name of the user who reviewed the event. When searching this field, you can enter **`unreviewed`** to search for events that have not been reviewed.

Revision (Syslog Only)

The version of the signature that was used to generate the event.

See also information about the following intrusion event fields: Generator, GID, Message, SID, and Snort ID.

Rule Group

A count of non-MITRE rule groups that you can click to bring up a modal, which shows the full list of rule groups.

Security Context (Syslog: Context)

The metadata identifying the virtual firewall group through which the traffic passed. The system only populates this field for ASA FirePOWER in multiple context mode.

SID (Syslog Only)

The signature ID (also known as the Snort ID) of the rule that generated the event.

See also information about the following intrusion event fields: Generator, GID, Message, Revision, and Snort ID.

Snort ID

This field is a search field only.

(For the syslog field, see SID.)

When performing your search: Specify the Snort ID (SID) of the rule that generated the event or, optionally, specify the combination Generator ID (GID) and SID of the rule, where the GID and SID are separated with a colon (:) in the format GID:SID. You can specify any of the values in the following table:

Table 110: Snort ID Search Values

Value	Example
a single SID	10000
a SID range	10000-11000
greater than a SID	>10000
greater than or equal to a SID	>=10000
less than a SID	<10000
less than or equal to a SID	<=10000
a comma-separated list of SIDs	10000,11000,12000
a single GID:SID combination	1:10000
a comma-separated list of GID:SID combinations	1:10000,1:11000,1:12000
a comma-separated list of SIDs and GID:SID combinations	10000,1:11000,12000

The SID of the events you are viewing is listed in the Message column. For more information, see the description in this section for the Message field.

Source Continent

The continent of the sending host involved in the intrusion event.

Source Country

The country of the sending host involved in the intrusion event.

Source Host Criticality

The source host criticality (value of the Host Criticality attribute for that corresponding host) when the event is generated.

Keep in mind that this field is not updated when the criticality of the host changes. However, new events will have the new criticality value.

Source IP (Syslog: SrcIP)

The IP address used by the sending host involved in the intrusion event.

See also Initiator and responder, source and destination, and sender and receiver fields, on page 773.

Source Port / ICMP Type (Syslog: SrcPort, ICMPType)

The port number on the sending host. For ICMP traffic, where there is no port number, this field displays the ICMP type.

Source User (Syslog: User)

The username associated with the IP address of the host that initiated the connection, which may or may not be the source host of the exploit. This user value is typically known only for users on your network.

If applicable, the username is preceded by `<realm>\`.

SSL Actual Action (Syslog: SSLActualAction)

In the Secure Firewall Management Center web interface, this field is a search field only.

The action the system applied to encrypted traffic:

Block/Block with reset

Represents blocked encrypted connections.

Decrypt (Resign)

Represents an outgoing connection decrypted using a re-signed server certificate.

Decrypt (Replace Key)

Represents an outgoing connection decrypted using a self-signed server certificate with a substituted public key.

Decrypt (Known Key)

Represents an incoming connection decrypted using a known private key.

Default Action

Indicates the connection was handled by the default action.

Do not Decrypt

Represents a connection the system did not decrypt.

Field values are displayed in the **SSL Status** field on the search workflow pages.

SSL Certificate Information

This field is a search field only.

The information stored on the public key certificate used to encrypt traffic, including:

- Subject/Issuer Common Name
- Subject/Issuer Organization
- Subject/Issuer Organization Unit
- Not Valid Before/After
- Serial Number
- Certificate Fingerprint
- Public Key Fingerprint

SSL Failure Reason

This field is a search field only.

The reason the system failed to decrypt encrypted traffic:

- Unknown
- No Match
- Success
- Uncached Session
- Unknown Cipher Suite
- Unsupported Cipher Suite
- Unsupported SSL Version
- SSL Compression Used
- Session Undecryptable in Passive Mode
- Handshake Error
- Decryption Error
- Pending Server Name Category Lookup
- Pending Common Name Category Lookup
- Internal Error
- Network Parameters Unavailable

- Invalid Server Certificate Handle
- Server Certificate Fingerprint Unavailable
- Cannot Cache Subject DN
- Cannot Cache Issuer DN
- Unknown SSL Version
- External Certificate List Unavailable
- External Certificate Fingerprint Unavailable
- Internal Certificate List Invalid
- Internal Certificate List Unavailable
- Internal Certificate Unavailable
- Internal Certificate Fingerprint Unavailable
- Server Certificate Validation Unavailable
- Server Certificate Validation Failure
- Invalid Action

Field values are displayed in the **SSL Status** field on the search workflow pages.

SSL Status

The action associated with the **SSL Actual Action** (SSL rule, default action, or undecryptable traffic action) that logged the encrypted connection.

If the system fails to decrypt an encrypted connection, it displays the **SSL Actual Action** (undecryptable traffic action) taken, as well as the **SSL Failure Reason**. For example, if the system detects traffic encrypted with an unknown cipher suite and allows it without further inspection, this field displays `Do Not Decrypt (Unknown Cipher Suite)`.

Click the **Lock icon** to view certificate details.

When searching this field, enter one or more of the **SSL Actual Action** and **SSL Failure Reason** values to view encrypted traffic the system handled or failed to decrypt.

SSL Subject/Issuer Country

This field is a search field only.

A two-character ISO 3166-1 alpha-2 country code for the subject or issuer country associated with the encryption certificate.

Time

The date and time of the event. This field is not searchable.

VLAN ID (Syslog: VLAN_ID)

The innermost VLAN ID associated with the packet that triggered the intrusion event.

Web Application (Syslog: WebApplication)

The web application, which represents the content or requested URL for HTTP traffic detected in the traffic that triggered the intrusion event.

If the system detects an application protocol of HTTP but cannot detect a specific web application, the system supplies a generic web browsing designation instead.

Web Application Category and Tag

Criteria that characterize the application to help you understand the application's function.

Intrusion event impact levels

To help you evaluate the impact an event has on your network, the Secure Firewall Management Center displays an impact level in the table view of intrusion events. For each event, the system adds an impact level icon whose color indicates the correlation between intrusion data, network discovery data, and vulnerability information.

Note Because no operating system information is available for hosts added to the network map from NetFlow data, the system cannot assign Vulnerable (impact level 1: red) impact levels for intrusion events involving those hosts. In such cases, use the host input feature to manually set the operating system identity for the hosts.

The following table describes the possible values for intrusion event impact levels:

Table 111: Impact levels

Impact Level	Vulnerability	Color	Description
Vulnerable (1)	Vulnerable	red	Either: • the source or the destination host is in the network map, and a vulnerability is mapped to the host • the source or destination host is potentially compromised by a virus, trojan, or other piece of malicious software
Potentially Vulnerable (2)	Potentially Vulnerable	orange	Either the source or the destination host is in the network map and one of these is true: • for port-oriented traffic, the port is running a server application protocol • for non-port-oriented traffic, the host uses the protocol

Impact Level	Vulnerability	Color	Description
Currently Not Vulnerable (3)	Currently Not Vulnerable	yellow	Either the source or the destination host is in the network map and one of these is true: • for port-oriented traffic (for example, TCP or UDP), the port is not open • for non-port-oriented traffic (for example, ICMP), the host does not use the protocol
Unknown Target (4)	Unknown Target	blue	Either the source or destination host is on a monitored network, but there is no entry for the host in the network map.
Unknown (0)	Unknown	gray	Neither the source nor the destination host is on a network that is monitored by network discovery.

View connection data associated with intrusion events

The system can log the connections where intrusion events are detected. Although this logging is automatic for intrusion policies associated with access control rules, you must manually enable connection logging to see associated connection data for the default action.

Viewing associated data is most useful when navigating between table views of events.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Intrusions** > **Events**.

Step 2 Choose the intrusion events using the check boxes in the table, then choose **Connections** from the **Jump to** drop-down list.

Tip
You can view the intrusion events associated with particular connections in a similar way. For more information, refer to Navigate between workflows, on page 701.

Mark intrusion events reviewed

If you are confident that an intrusion event is not malicious, you can mark the event reviewed.

If you have examined an intrusion event and are confident that the event does not represent a threat to your network security (for example, because you know that none of the hosts on your network are vulnerable to the detected exploit), you can mark the event reviewed.

Reviewed events are stored in the event database and are included in the event summary statistics, but no longer appear in the default intrusion event pages. Your name appears as the reviewer.

In a multidomain deployment, if you mark an event reviewed, the system marks it reviewed in all domains that can view that event.

If you perform a backup and then delete reviewed intrusion events, restoring your backup restores the deleted intrusion events but does not restore their reviewed status. You view those restored intrusion events under **Intrusion Events**, not under **Reviewed Events**.

Procedure

On a page that displays intrusion events, you have two options:

- To mark one or more intrusion events from the list of events, check the check boxes next to the events and click **Review**.
- To mark all intrusion events from the list of events, click **Review All**.

View previously reviewed intrusion events

In a multidomain deployment, if you mark an event reviewed, the system marks it reviewed in all domains that can view that event.

Procedure

Step 1 Choose **Analysis** > **Intrusions** > **Reviewed Events**.

Step 2 You have the following choices:

- Adjust the time range as described in Change the time window, on page 695.
- If you are using a custom workflow that does not include the table view of intrusion events, choose any of the system-provided workflows by clicking **(switch workflow)** next to the workflow title.
- To learn more about the events that appear, refer to Intrusion event fields, on page 790.

Mark reviewed intrusion events unreviewed

You can return a reviewed event to the default intrusion events view by marking the event unreviewed.

In a multidomain deployment, if you mark an event reviewed, the system marks it reviewed in all domains that can view that event.

Procedure

On a page that displays reviewed events, you have two choices:

- To remove individual intrusion events from the list of reviewed events, check the check boxes next to specific events and click **Unreview**.
- To remove all intrusion events from the list of reviewed events, click **Unreview All**.

Preprocessor events

A preprocessor event is a security detection mechanism that

- reports the execution of specified preprocessor options by generating an event whenever a packet triggers that preprocessor option and the associated preprocessor rule is enabled,
- helps detect anomalous protocol exploits by analyzing packet characteristics and protocol behaviors, and
- differs from rule events in that the packet display does not include a detailed rule description for the event.

Preprocessor event characteristics

Preprocessors provide two functions: performing the specified action on the packet (for example, decoding and normalizing HTTP traffic) and reporting the execution of specified preprocessor options by generating an event whenever a packet triggers that preprocessor option and the associated preprocessor rule is enabled.

Preprocessor events differ from rule events in that the packet display does not include a detailed rule description for the event. Instead, the packet display shows the event message, the GID, the SID, the packet header data, and the packet payload. This allows you to analyze the packet's header information, determine if its header options are being used and if they can exploit your system, and inspect the packet payload. After the preprocessors analyze each packet, the rules engine executes appropriate rules against it (if the preprocessor was able to defragment it and establish it as part of a valid session) to further analyze potential content-level threats and report on them.

Double encoding detection example

You can enable the `Double Encoding` HTTP Inspect option and the associated preprocessor rule with the HTTP Inspect Generator (GID) 119 and the Snort ID (SID) 2 to generate an event when the preprocessor encounters IIS double-encoded traffic.

Attackers can craft overlapping IP fragments to cause a DoS attack on a host. The IP defragmentation preprocessor can detect this type of attack and generate an intrusion event for it.

Preprocessor generator IDs

Each preprocessor has its own Generator ID number, or GID, that indicates which preprocessor was triggered by the packet. Some of the preprocessors also have related SIDs, which are ID numbers that classify potential attacks. This helps you analyze events more effectively by categorizing the type of event much the way a

rule's Snort ID (SID) can offer context for packets triggering rules. You can list preprocessor rules by preprocessor in the Preprocessors filter group on the intrusion policy Rules page; you can also list preprocessor rules in the preprocessor and packet decoder sub-groupings in the Category filter group.

Note Events generated by standard text rules have a generator ID of 1 (Global domain or legacy GID) or 1000 - 2000 (descendant domains). For shared object rules, the events have a generator ID of 3. For both, the event's SID indicates which specific rule triggered.

Table 112: Generator IDs

ID	Component	Description
1	Standard Text Rule	The event was generated when the packet triggered a standard text rule (Global domain or legacy GID).
2	Tagged Packets	The event was generated by the Tag generator, which generates packets from a tagged session. This occurs when the `tag` rule option is used.
3	Shared Object Rule	The event was generated when the packet triggered a shared object rule.
102	HTTP Decoder	The decoder engine decoded HTTP data within the packet.
105	Back Orifice Detector	The Back Orifice Detector identified a Back Orifice attack associated with the packet.
106	RPC Decoder	The RPC decoder decoded the packet.
116	Packet Decoder	The event was generated by the packet decoder.
119, 120	HTTP Inspect Preprocessor	The event was generated by the HTTP Inspect preprocessor. GID 120 rules relate to server-specific HTTP traffic.
122	Portscan Detector	The event was generated by the portscan flow detector.
123	IP Defragmentor	The event was generated when a fragmented IP datagram could not be properly reassembled.
124	SMTP Decoder	The event was generated when the SMTP preprocessor detected an exploit against an SMTP verb.
125	FTP Decoder	The event was generated when the FTP/Telnet decoder detected an exploit within FTP traffic.
126	Telnet Decoder	The event was generated when the FTP/Telnet decoder detected an exploit within telnet traffic.
128	SSH Preprocessor	The event was generated when the SSH preprocessor detected an exploit within SSH traffic.

ID	Component	Description
129	Stream Preprocessor	The event was generated during stream preprocessing by the stream preprocessor.
131	DNS Preprocessor	The event was generated by the DNS preprocessor.
133	DCE/RPC Preprocessor	The event was generated by the DCE/RPC preprocessor.
134	Rule Latency Packet Latency	The event was generated when rule latency suspended (134:1) or re-enabled (134:2) a group of intrusion rules, or when the system stopped inspecting a packet because the packet latency threshold was exceeded (134:3).
135	Rate-Based Attack Detector	The event was generated when a rate-based attack detector identified excessive connections to hosts on the network.
137	SSL Preprocessor	The event was generated by the TLS/SSL preprocessor.
138, 139	Sensitive Data Preprocessor	The event was generated by the sensitive data preprocessor.
140	SIP Preprocessor	The event was generated by the SIP preprocessor.
141	IMAP Preprocessor	The event was generated by the IMAP preprocessor.
142	POP Preprocessor	The event was generated by the POP preprocessor.
143	GTP Preprocessor	The event was generated by the GTP preprocessor.
144	Modbus Preprocessor	The event was generated by the Modbus SCADA preprocessor.
145	DNP3 Preprocessor	The event was generated by the DNP3 SCADA preprocessor.
148	CIP Preprocessor	The event was generated by the CIP SCADA preprocessor.
149	S7Commplus preprocessor	The event was generated by the S7Commplus SCADA preprocessor.
1000 - 2000	Standard Text Rule	The event was generated when the packet triggered a standard text rule (descendant domains).

Intrusion event workflow pages

Intrusion event workflow pages are a set of predefined workflows that

- provide event data populated interfaces to view and analyze intrusion events generated by preprocessor, decoder, and intrusion rules violations,
- step you through a series of pages to help pinpoint specific intrusion events for evaluation, and
- contain three different types of pages or event views for comprehensive analysis.

Workflow page types and functionality

The preprocessor, decoder, and intrusion rules that are enabled in the current intrusion policy generate intrusion events whenever the traffic that you monitor violates the policy.

The predefined intrusion event workflows contain three different types of pages, or event views:

- one or more drill-down pages
- the table view of intrusion events
- a packet view

Drill-down pages generally include two or more columns in a table (and, for some drill-down views, more than one table) that allow you to view one specific type of information.

When you "drill down" to find more information for one or more destination ports, you automatically select those events and the next page in the workflow appears. In this way, drill-down tables help you reduce the number of events you are analyzing at one time.

The initial *table view* of intrusion events lists each intrusion event in its own row. The columns in the table list information such as the time, the source IP address and port, the destination IP address and port, the event priority, the event message, and more.

When you select events on a table view, instead of selecting events and displaying the next page in the workflow, you add to what are called *constraints*. Constraints are limits that you impose on the types of events that you want to analyze.

For example, if you click **Close** ( ) in any column and clear **Time** from the drop-down list, you can remove Time as one of the columns. To narrow the list of events in your analysis, you can click the link for a value in one of the rows in the table view. For example, to limit your analysis to the events generated from one of the source IP addresses (presumably, a potential attacker), click the IP address in the **Source IP Address** column.

If you select one or more rows in a table view and then click **View**, the packet view appears. A *packet view* provides information about the packet that triggered the rule or the preprocessor that generated the event. Each section of the packet view contains information about a specific layer in the packet. You can expand collapsed sections to see more information.

Note Because each portscan event is triggered by multiple packets, portscan events use a special version of the packet view.

If the predefined workflows do not meet your specific needs, you can create custom workflows that display only the information you are interested in. Custom intrusion event workflows can include drill-down pages, a table view of events, or both; the system automatically includes a packet view as the last page. You can easily switch between the predefined workflows and your own custom workflows depending on how you want to investigate events.

Use intrusion event workflows

To avoid displaying the same intrusion events on different workflow pages, the time range pauses when you click a link at the bottom of the page to display another page of events, and resumes when you click to take any other action on the subsequent page.

Tip At any point in the process, you can save the constraints as a set of search criteria. For example, if you find that over the course of a few days your network is being probed by an attacker from a single IP address, you can save your constraints during your investigation and then use them again later. You cannot, however, save compound constraints as a set of search criteria.

Procedure

Step 1 Access an intrusion event workflow using **Analysis** > **Intrusions** > **Events**.

Step 2 Optionally, constrain the number of intrusion events that appear on the event views as described in Intrusion event drill-down page constraints, on page 810 or Intrusion event table view constraints, on page 811.

Step 3 Choose from the available workflow options.

Event analysis and intelligence

- To learn more about the columns that appear, refer to Intrusion event fields, on page 790.
- To view a host's profile, click **Host Profile** that appears next to the host IP address.
- To view geolocation details, click flag that appears in the Source Country or Destination Country columns.
- To view data in available sources external to your system, right-click an event value. The options you see depend on the data type and include public sources; other sources depend on the resources you have configured. For information, refer to Event investigation using web-based resources, on page 638
- To gather general intelligence about an event, right-click an event value in the table and choose from a Cisco or third-party intelligence source. For example, you can get details about a suspicious IP address from Cisco Talos. The options you see depend on the data type and the integrations that are configured on your system. For more information, refer to Event investigation using web-based resources, on page 638.

Time and data management

- To modify the time and date range for displayed events, refer to Change the time window, on page 695.

  Tip
  If no intrusion events appear on the event views, adjusting the specified time range might return results. If you specified an older time range, events in that time range might have been deleted. Adjusting the rule thresholding configuration might generate events.

  Note
  Events generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.

Event actions

- To delete events from the event database, check the check boxes next to events you want to delete, then click **Delete**, or click **Delete All**.

- To mark events reviewed to remove them from intrusion event pages, but not the event database, refer to Mark intrusion events reviewed, on page 803.
- To download a local copy of the packet (a packet capture file in libpcap format) that triggered each selected event, check the check boxes next to events triggered by the packets you want to download, then click **Download Packets**, or click **Download All Packets**. Captured packets are saved in libpcap format. This format is used by several popular protocol analyzers.

Navigation and workflow

- To sort events on the current workflow page or navigate within the current workflow page, refer to Use workflows, on page 676.
- To navigate between pages in the current workflow, keeping the current constraints, click the appropriate page link at the top left of the workflow page.
- To navigate to other event views to view associated events, refer to Navigate between workflows, on page 701.
- To temporarily use a different workflow, click **(switch workflow)**.
- To bookmark the current page so that you can quickly return to it, click **Bookmark This Page**.
- To view the Intrusion Events section of the Summary Dashboard, click **Dashboards**.
- To navigate to the bookmark management page, click **View Bookmarks**.
- To generate a report based on the data in the current view, refer to Create a report template from an event view, on page 536.

Intrusion event drill-down page constraints

The table describes how to use drill-down pages to constrain events in intrusion event workflows.

Table 113: Constraining events on drill-down pages

To...	You can...
drill down to the next workflow page constraining on a specific value	click the value. For example, on the Destination Port workflow, to constrain the events to those with a destination of port 80, click **80/tcp** in the **DST Port/ICMP Code** column. The next page of the workflow, Events, appears and contains only port 80/tcp events.

To...	You can...
drill down to the next workflow page constraining on selected events	select the check boxes next to the events you want to view on the next workflow page, then click **View**. For example, on the Destination Port workflow, to constrain the events to those with destination ports 20/tcp and 21/tcp, select the check boxes next to the rows for those ports and click **View**. The next page of the workflow, Events, appears and contains only port 20/tcp and 21/tcp events. Note that if you constrain on multiple rows and the table has more than one column (not including a Count column), you build what is called a compound constraint. Compound constraints ensure that you do not include more events in your constraint than you mean to. For example, if you use the Event and Destination workflow, each row that you select on the first drill-down page creates a compound constraint. If you pick event 1:100 with a destination IP address of 10.10.10.100 and you also pick event 1:200 with a destination IP address of 192.168.10.100, the compound constraint ensures that you do not also select events with 1:100 as the event type and 192.168.10.100 as the destination IP address or events with 1:200 as the event type and 10.10.10.100 as the destination IP address.
drill down to the next workflow page keeping the current constraints	click **View All**.

Intrusion event table view constraints

The table describes how to use the table view.

Table 114: Constraining events on the table view of events

To...	You can...
constrain the view to events with a single attribute	click the attribute. For example, to constrain the view to events with a destination of port 80, click **80/tcp** in the **DST Port/ICMP Code** column.
remove a column from the table	click **Close** ( ) in the column heading that you want to hide. In the pop-up window that appears, click **Apply**. If you want to hide or show other columns, select or clear the appropriate check boxes before you click **Apply**. To add a disabled column back to the view, click the **expand arrow** to expand the search constraints, then click the column name under **Disabled Columns**.
view the packets associated with one or more events	either: • click the **down arrow** next to the event whose packets you want to view. • select one or more events whose packets you want to view, and, at the bottom of the page, click **View.** • at the bottom of the page, click **View All**to view the packets for all events that match the current constraints.

Using the intrusion event packet view

A packet view provides information about the packet that triggered the rule that generated an intrusion event.

Tip The packet view on a Secure Firewall Management Center does not contain packet information when the **Transfer Packet**option is disabled for the device detecting the event.

The packet view indicates why a specific packet was captured by providing information about the intrusion event that the packet triggered, including the event's time stamp, message, classification, priority, and, if the event was generated by a standard text rule, the rule that generated the event. The packet view also provides general information about the packet, such as its size.

In addition, the packet view has a section that describes each layer in the packet: data link, network, and transport, as well as a section that describes the bytes that comprise the packet. If the system decrypted the packet, you can view the decrypted bytes. You can expand collapsed sections to display detailed information.

Note Because each portscan event is triggered by multiple packets, portscan events use a special version of the packet view.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 On the table view of intrusion events, choose packets to view as described in Intrusion event table view constraints, on page 811.

Step 2 Optionally, if you chose more than one event, you can page through the packets in the packet view by using the page numbers at the bottom of the page.

Step 3 You also have these options:

- Adjust: To modify the date and time range in the packet views, refer to Change the time window, on page 695.
- Configure: To configure the intrusion rule that triggered the event, click the arrow next to Actions and continue as described in Configure intrusion rules within the packet view, on page 816.
- Delete: To delete an event from the database, click **Delete** to delete the event whose packet you are viewing or click **Delete All** to delete all the events whose packets you previously selected.
- Download: To download a local copy of the packet (a packet capture file in libpcap format) that triggered the event, click **Download Packet**to save a copy of the captured packet for the event you are viewing or click **Download All Packets** to save copies of the captured packets for all the events whose packets you previously selected. The captured packet is saved in libpcap format. This format is used by several popular protocol analyzers.

  Note
  You cannot download a portscan packet because single portscan events are based on multiple packets; however, the portscan view provides all usable packet information. You must have at least 15% available disk space in order to download.

- Mark reviewed: To mark an event reviewed to remove it from event views, but not the event database, click **Review** to mark the event whose packet you are viewing or click **Review All** to mark all the events whose packets you previously selected. For more information, refer to Mark intrusion events reviewed, on page 803.
- View additional information: To expand or collapse a page section, click the arrow next to the section. For details, refer to Event information fields, on page 813, Frame information fields, on page 819, and Data link layer information fields, on page 819.
- View network layer information: Refer to View network layer information, on page 820.
- View packet byte information: Refer to View packet byte information, on page 825.
- View transport layer information: Refer to View transport layer information, on page 823

Event information fields

On the packet view, you can view information about the packet in the event information section.

Event

The event message. For rule-based events, this corresponds to the rule message. For other events, this is determined by the decoder or preprocessor.

The ID for the event is appended to the message in the format (GID:SID:Rev). GID is the generator ID of the rules engine, the decoder, or the preprocessor that generated the event. SID is the identifier for the rule, decoder message, or preprocessor message. Rev is the revision number of the rule.

Timestamp

The time that the packet was captured, in UTC time zone.

Classification

The event classification. For rule-based events, this corresponds to the rule classification. For other events, this is determined by the decoder or preprocessor.

Priority

The event priority. For rule-based events, this corresponds to either the value of the `priority` keyword or the value for the `classtype` keyword. For other events, this is determined by the decoder or preprocessor.

Ingress security zone

The ingress security zone of the packet that triggered the event. Only this security zone field is populated in a passive deployment.

Egress security zone

The egress security zone of the packet that triggered the event. This field is not populated in a passive deployment.

Domain

The domain where the managed device belongs. This field is only present if you have ever configured the Firewall Management Center for multitenancy.

Device

The managed device where the access control policy was deployed.

Security context

The metadata identifying the virtual firewall group through which the traffic passed. Note that the system only populates this field for ASA FirePOWER in multiple context mode.

Ingress interface

The ingress interface of the packet that triggered the event. Only this interface column is populated for a passive interface.

Egress interface

For an inline set, the egress interface of the packet that triggered the event.

Source/Destination IP

The host IP address or domain name where the packet that triggered the event (source) originated, or the target (destination) host of the traffic that triggered the event.

Source Port/ICMP type

Source port of the packet that triggered the event. For ICMP traffic, where there is no port number, the system displays the ICMP type.

Destination Port/ICMP code

The port number for the host receiving the traffic. For ICMP traffic, where there is no port number, the system displays the ICMP code.

Email headers

The data that was extracted from the email header. Note that email headers do not appear in the table view of intrusion events, but you can use email header data as a search criterion.

To associate email headers with intrusion events for SMTP traffic, you must enable the SMTP preprocessor **Log Headers** option. For rule-based events, this row appears when email data is extracted.

HTTP hostname

The host name, if present, extracted from the HTTP request Host header. This row displays the complete host name, up to 256 bytes. You can expand the complete host name if it is longer than a single row.

To display host names, you must enable the HTTP Inspect preprocessor **Log Hostname** option.

Note that HTTP request packets do not always include a host name. For rule-based events, this row appears when the packet contains the HTTP host name or the HTTP URI.

HTTP URI

The normalized URI, if present, associated with the HTTP request packet that triggered the intrusion event. This row displays the complete URI, up to 2048 bytes. You can expand the complete URI if it is longer than a single row.

To display the URI, you must enable the HTTP Inspect preprocessor **Log URI** option.

Note that HTTP request packets do not always include a URI. For rule-based events, this row appears when the packet contains the HTTP host name or the HTTP URI.

To see the associated HTTP URI in intrusion events triggered by HTTP responses, you should configure HTTP server ports in the **Perform Stream Reassembly on Both Ports** option; note, however, that this increases resource demands for traffic reassembly.

Intrusion policy

The intrusion policy, if present, where the intrusion, preprocessor, or decoder rule that generated the intrusion event was enabled. You can choose an intrusion policy as the default action for an access control policy or associate an intrusion policy with an access control rule.

Access control policy

The access control policy that includes the intrusion policy where the intrusion, preprocessor, or decoder rule that generated the event is enabled.

Access control rule

The access control rule associated with an intrusion rule that generated the event. Default Action indicates that the intrusion policy where the rule is enabled is not associated with an access control rule but, instead, is configured as the default action of the access control policy.

Rule

For standard text rule events, the rule that generated the event.

Note that if the event is based on a shared object rule, a decoder, or a preprocessor, the rule is not available.

Because rule data may contain sensitive information about your network, administrators may toggle users' ability to view rule information in the packet view with the View Local Rules permission in the user role editor.

Actions

For standard text and custom rule events, expand **Actions** to take any of these actions on the rule that triggered the event:

- edit the rule
- view documentation for the revision of the rule; for standard text rules only, after clicking **View Documentation** under Actions, you can click **Rule Documentation** in the documentation pop-up window to view more-specific rule details.
- add a comment to the rule
- change the state of the rule
- set a threshold for the rule

- suppress the rule

If the event is based on a shared object rule, a decoder, or a preprocessor, the rule is not available.

Configure intrusion rules within the packet view

Within the packet view of an intrusion event, you can take several actions on the rule that triggered the event. Note that if the event is based on a shared object rule, a decoder, or a preprocessor, the rule is not available.

Procedure

Step 1 Within the packet view of an intrusion event that was generated by an intrusion rule, expand **Actions** in the Event Information section.

Step 2 You have the following choices:

- Comment: For standard text rule events, click **Rule Comment** to add a text comment to the rule that generated the event. This allows you to provide additional context and information about the rule and the exploit or policy violation it identifies. You can also add and view rule comments in the intrusion rules editor.
- Disable: To disable this rule, click one of the following options:
  - **Disable this rule in the current Snort 2 policy (<policy_name>)**
  - **Disable this rule in all locally created Snort 2 policies**

  If this event is generated by a standard text rule, you can disable the rule, if necessary. You can set the rule in all policies that you can edit locally. Alternately, you can set the rule only in the current policy (that is, the policy that generated the event) if you can edit the current policy locally.

  Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by the system.

  Note
  You cannot disable shared object rules from the packet view, nor can you disable rules in the default policies.

- Drop packets and generate an event: To set the rule to drop packets that trigger it and generate an event, click one of the following options:
  - **Set this rule to drop the triggering packet and generate an event in the current Snort 2 policy (<policy_name>)**
  - **Set this rule to drop the triggering packet and generate an event in all locally created Snort 2 inline policies**

  If your managed device is deployed inline on your network, you can set the rule that triggered the event to drop packets that trigger the rule in all policies that you can edit locally. Alternately, you can set the rule only in the current policy (that is, the policy that generated the event) if you can edit the current policy locally.

  Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by the system. Note also that this option appears only when **Drop when Inline** is enabled in the current policy.

- Edit: For standard text rule events, click **Edit** (to edit the Snort 2 rule) or **Edit Snort 3 Rule** to modify the rule that generated the event. If the event is based on a shared object rule, a decoder, or a preprocessor, the rule is not available.

  Note
  If you edit a system-provided rule (as opposed to a custom standard text rule), you actually create a new local rule. Make sure you set the local rule to generate events and also disable the original rule in the current intrusion policy. Note, however, that you cannot enable local rules in the default policies.

- Generate events: Click **Set this rule to generate events in all locally created Snort 2 policies** to set the rule to generate events.

  If this event is generated by a standard text rule, you can set the rule to generate events in all policies that you can edit locally.

  Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by the system.

  Note
  You cannot set shared object rules to generate events from the packet view, nor can you disable rules in the default policies.

- Set suppression options: Expand **Set Suppression Options** and continue as described in Set suppression options within the packet view, on page 818.

  You can use this option to suppress the rule that triggered this event in all policies that you can edit locally. Alternately, you can suppress the rule only in the current policy (that is, the policy that generated the event) if you can edit the current policy locally.

  Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by Cisco.

- Set threshold options: Expand **Set Thresholding Options** and continue as described in Set threshold options within the packet view, on page 817.

  You can use this option to create a threshold for the rule that triggered this even in all policies that you can edit locally. Alternately, you create a threshold only for the current policy (that is, the policy that generated the event) if you can edit the current policy locally.

  Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default intrusion policy provided by the system.

- View documentation: Click **View Documentation** to learn more about the rule that generated the event. Optionally, then click **Rule Documentation** to view more-specific rule details.

Set threshold options within the packet view

You can set threshold options in all policies that you can edit locally or, when it can be edited locally, only in the current policy (that is, the policy that caused the event to be generated).

Procedure

Step 1 Within the packet view of an intrusion event that was generated by an intrusion rule, expand **Actions** in the Event Information section.

Step 2 Expand **Set Thresholding Options** and choose one of the two possible options:

- **in the current Snort 2 policy (<policy_name>)**
- **in all locally created Snort 2 policies**

Step 3 Choose the type of threshold you want to set:

- Click **limit** to limit notification to the specified number of event instances per time period.
- Click **threshold** to provide notification for each specified number of event instances per time period.
- Click **both** to provide notification once per time period after a specified number of event instances.

Step 4 Click the appropriate threshold to indicate whether you want the event instances tracked by **Source** or **Destination** IP address.

Step 5 In the **Count** field, enter the number of event instances you want to use as your threshold.

Step 6 In the **Seconds** field, enter a number between 1 and 86400 that specifies the time period for which event instances are tracked.

Step 7 If you want to override any current thresholds for this rule in existing intrusion policies, check the **Override any existing settings for this rule** check box.

Step 8 Click **Save Thresholding**.

Set suppression options within the packet view

This task is performed within the packet view of an intrusion event that was generated by an intrusion rule.

Before you begin

Follow these steps to set suppression options within the packet view:

Procedure

Step 1 Within the packet view of an intrusion event that was generated by an intrusion rule, expand **Actions** in the Event Information section.

Step 2 Expand **Set Suppression Options** and click one of the two possible options:

- **in the current Snort 2 policy (<policy_name>)**
- **in all locally created Snort 2 policies**

Note
The current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by Cisco.

Step 3 Choose one of the following **Track By** options:

- Click **Source** to suppress events generated by packets originating from a specified source IP address.
- Click **Destination** to suppress events generated by packets going to a specified destination IP address.
- Click **Rule** to completely suppress events for the rule that triggered this event.

Step 4 In the **IP address or CIDR block** field, enter the IP address or CIDR block/prefix length you want to specify as the source or destination IP address.

Step 5 Click **Save Suppression**.

Frame information fields

On the packet view, click the arrow next to **Frame** to view information about the captured frame. The packet view may display a single frame or multiple frames. Each frame provides information about an individual network packet. You would see multiple frames, for example, in the case of tagged packets or packets in reassembled TCP streams.

Frame information fields provide these details:

Frame identification and timing

- **Frame n**: The captured frame, where *n* is 1 for single-frame packets and the incremental frame number for multi-frame packets. The number of captured bytes in the frame is appended to the frame number.
- **Arrival Time**: The date and time the frame was captured.
- **Time delta from previous captured frame**: For multi-frame packets, the elapsed time since the previous frame was captured.
- **Time delta from previous displayed frame**: For multi-frame packets, the elapsed time since the previous frame was displayed.
- **Time since reference or first frame**: For multi-frame packets, the elapsed time since the first frame was captured.
- **Frame Number**: The incremental frame number.

Frame characteristics

- **Frame Length**: The length of the frame in bytes.
- **Capture Length**: The length of the captured frame in bytes.
- **Frame is marked**: Whether the frame is marked (true or false).
- **Protocols in frame**: The protocols included in the frame.

Data link layer information fields

On the packet view, click the arrow next to the data link layer protocol (for example, **Ethernet II**) to view the data link layer information about the packet, which contains the 48-bit media access control (MAC) addresses for the source and destination hosts. It may also display other information about the packet, depending on the hardware protocol.

Note Note that this example discusses Ethernet link layer information; other protocols may also appear.

The packet view reflects the protocol used at the data link layer. This reference describes the information you might see for an Ethernet II or IEEE 802.3 Ethernet packet in the packet view.

Destination

The MAC address for the destination host.

Note Ethernet can also use multicast and broadcast addresses as the destination address.

Source

The MAC address for the source host.

Type

For Ethernet II packets, the type of packet that is encapsulated in the Ethernet frame; for example, IPv6 or ARP datagrams. Note that this item only appears for Ethernet II packets.

Length

For IEEE 802.3 Ethernet packets, the total length of the packet, in bytes, not including the checksum. Note that this item only appears for IEEE 802.3 Ethernet packets.

View network layer information

When analyzing network packets, you need to examine the network layer protocol details to understand packet routing, addressing, and protocol-specific information. This view provides comprehensive data about IP packets and other network layer protocols.

Procedure

On the packet view, click the arrow next to the network layer protocol (for example, **Internet Protocol**) to view more detailed information about network layer information related to the packet.

Note
Note that this example discusses IP packets; other protocols may also appear.

IPv4 network layer information fields

This listing describes protocol-specific information that might appear in an IPv4 packet.

Version

The Internet Protocol version number.

Header length

The number of bytes in the header, including any IP options. An IP header with no options is 20 bytes long.

Differentiated services field

The values for differentiated services that indicate how the sending host supports Explicit Congestion Notification (ECN):

- `0x0`: does not support ECN-Capable Transport (ECT)
- `0x1` and `0x2`: supports ECT
- `0x3`: Congestion Experienced (CE)

Total length

The length of the IP packet, in bytes, minus the IP header.

Identification

The value that uniquely identifies an IP datagram sent by the source host. This value is used to trace fragments of the same datagram.

Flags

The values that control IP fragmentation, where:

values for the Last Fragment flag indicate whether there are more fragments associated with the datagram:

- `0`: there are no more fragments associated with the datagram
- `1`: there are more fragments associated with the datagram

values for the Don't Fragment flag control whether the datagram can be fragmented:

- `0`: the datagram can be fragmented
- `1`: the datagram must **not** be fragmented

Fragment offset

The value for the fragment offset from the beginning of the datagram.

Time to live (ttl)

The remaining number of hops that the datagram can make between routers before the datagram expires.

Protocol

The transport protocol that is encapsulated in the IP datagram; for example, ICMP, IGMP, TCP, or UDP.

Header checksum

The indicator for whether the IP checksum is valid. If the checksum is invalid, the datagram may have been corrupted during transit or may be being used in an intrusion evasion attempt.

Source/Destination

The IP address or domain name for the source (or destination) host.

Note that to display the domain name, you must enable IP address resolution.

Click the address or domain name to view the context menu, then select **Whois** to do a whois search on the host, **View Host Profile** to view host information, or choose an option to add the address to a global Block list or Do-Not-Block list.

IPv6 network layer information fields

This listing describes protocol-specific information that might appear in an IPv6 packet.

Traffic class

An experimental 8-bit field in the IPv6 header for identifying IPv6 packet classes or priorities similar to the differentiated services functionality provided for IPv4. When unused, this field is set to zero.

Flow label

A optional 20-bit IPv6 hexadecimal value 1 to FFFFF that identifies a special flow such as non-default quality of service or real-time service. When unused, this field is set to zero.

Payload length

A 16-bit field identifying the number of octets in the IPv6 payload, which is comprised of all of the packet following the IPv6 header, including any extension headers.

Next header

An 8-bit field identifying the type of header immediately following the IPv6 header, using the same values as the IPv4 Protocol field.

Hop limit

An 8-bit decimal integer that each node that forwards the packet decrements by one. The packet is discarded if the decremented value reaches zero.

Source

The 128-bit IPv6 address for the source host.

Destination

The 128-bit IPv6 address for the destination host.

View transport layer information

Use this procedure when analyzing network packets to examine transport layer details including TCP, UDP, and ICMP protocol information and payload data.

Procedure

Step 1 On the packet view, click the arrow next to the transport layer protocol (for example, **TCP**, **UDP**, or **ICMP**).

Step 2 Optionally, click **Data** when present to view the first twenty-four bytes of the payload for the protocol immediately above it in the Packet Information section of the packet view.

Step 3 View the contents of the transport layer for TCP, UDP, and ICMP protocols as described in TCP packet view fields, on page 823, UDP packet view fields, on page 824, or ICMP packet view fields, on page 824.

Note
These examples discuss TCP, UDP, and ICMP packets; other protocols may also appear.

TCP packet view fields

This section describes the protocol-specific information for a TCP packet.

Source port

The number that identifies the originating application protocol.

Destination port

The number that identifies the receiving application protocol.

Sequence number

The value for the first byte in the current TCP segment, keyed to initial sequence number in the TCP stream.

Next sequence number

In a response packet, the sequence number of the next packet to send.

Acknowledgment number

The TCP acknowledgment, which is keyed to the sequence number of the previously accepted data.

Header length

The number of bytes in the header.

Flags

The six bits that indicate the TCP segment's transmission state:

- `U`: The urgent pointer is valid

- A: The acknowledgment number is valid
- P: The receiver should push data
- R: Reset the connection
- S: Synchronize sequence numbers to start a new connection
- F: The sender has finished sending data

Window size

The amount of unacknowledged data, in bytes, that the receiving host will accept.

Checksum

The indicator for whether the TCP checksum is valid. If the checksum is invalid, the datagram may have been corrupted during transit or may be being used in an in evasion attempt.

Urgent pointer

The position, if present, in the TCP segment where the urgent data ends. Used in conjunction with the U flag.

Options

The values, if present, for TCP options.

UDP packet view fields

This section describes the protocol-specific information for a UDP packet.

Source port

The number that identifies the originating application protocol.

Destination port

The number that identifies the receiving application protocol.

Length

The combined length of the UDP header and data.

Checksum

The indicator for whether the UDP checksum is valid. If the checksum is invalid, the datagram may have been corrupted during transit.

ICMP packet view fields

The Internet Control Message Protocol (ICMP) packets contain several protocol-specific fields. These fields help identify the type of message, additional codes for certain message types, and ensure packet integrity.

Type

The type field indicates the kind of ICMP message:

Echo and error reporting

- 0: Echo reply
- 3: Destination unreachable
- 8: Echo request
- 11: Time exceeded
- 12: Parameter problem

Routing and flow control

- 4: Source quench
- 5: Redirect
- 9: Router advertisement
- 10: Router solicitation

Timestamp and address management

- 13: Timestamp request
- 14: Timestamp reply
- 17: Address mask request
- 18: Address mask reply

Obsolete messages

- 15: Information request (obsolete)
- 16: Information reply (obsolete)

Code

This field provides additional context for certain ICMP message types. For example, ICMP message types 3, 5, 11, and 12 use codes as defined in RFC 792 to further specify the message's meaning.

Checksum

This field indicates whether the ICMP packet is valid. If the checksum is invalid, the datagram may have been corrupted while being transmitted.

View packet byte information

Use this procedure when you need to analyze packet contents at the byte level for troubleshooting or network analysis purposes. If the system has decrypted traffic, you can also view the decrypted packet bytes.

Procedure

On the packet view, click the arrow next to **Packet Bytes** to view hexadecimal and ASCII versions of the bytes that comprise the packet. If the system decrypted traffic, you can view the decrypted packet bytes.

Internally sourced intrusion events

An internally sourced intrusion event is a security incident that originates from internal sources and indicates a compromised host on your network.

Investigation indicators

If the source IP address is on your network, this is a sign that you should investigate this host.

View intrusion event statistics

The Intrusion Event Statistics page provides you with a quick summary of the current state of your appliance and any intrusion events generated for your network.

Each of the IP addresses, ports, protocols, event messages, and so on shown on the page is a link. Click any link to view the associated event information. For example, if one of the top 10 destination ports is `80 (http)/tcp`, clicking that link displays the first page in the default intrusion events workflow, and lists the events targeting that port. Note that only the events (and the managed devices that generate events) in the current time range appear. Also, intrusion events that you have marked reviewed continue to appear in the statistics. For example, if the current time range is the past hour but the first event was generated five hours ago, when you click the **First Event** link, the resulting event pages will not show the event until you change the time range.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Overview** > **Summary** > **Intrusion Event Statistics**.

Step 2 From the two selection boxes at the top of the page, choose the zones and devices whose statistics you want to view, or choose **All Security Zones** and **All Devices** to view statistics for all the devices that are collecting intrusion events.

Step 3 Click **Get Statistics**.

Tip
To view data from a custom time range, click the link in the upper right page area and follow the directions in Change the time window, on page 695.

Host statistics

The Host Statistics section of the intrusion event statistics page provides information about the appliance itself. On the Secure Firewall Management Center, this section also provides information about any managed devices.

This information includes these metrics:

- **Time**: The current time on the appliance.
- **Uptime**: The number of days, hours, and minutes since the appliance itself was restarted. On the Secure Firewall Management Center, the uptime also shows the last time each managed device was rebooted, the number of users logged in, and the load average.
- **Disk Usage**: The percentage of the disk that is being used.
- **Memory Usage**: The percentage of system memory that is being used.
- **Load Average**: The average number of processes in the CPU queue for the past 1 minute, 5 minutes, and 15 minutes.

Event overview

The event overview section of the intrusion event statistics page provides an overview of the information in the intrusion event database.

These statistics include the following:

- **Events**: The number of events in the intrusion event database.
- **Events in Time Range**: The currently selected time range as well as the number and percentage of events from the database that fall within the time range.
- **First Event**: The event message for the first event in the event database.
- **Last Event**: The event message for the last event in the event database.



Note If you select a managed device while viewing intrusion event data on the Secure Firewall Management Center, the Event Overview section for that device appears instead.

Event statistics

The Event Statistics section of the Intrusion Event Statistics page provides more specific information about the information in the intrusion event database.

This information includes details on:

- the top 10 event types
- the top 10 source IP addresses
- the top 10 destination IP addresses

- the top 10 destination ports
- the protocols, ingress and egress security zones, and devices with the greatest number of events

Note In a multidomain deployment, the system builds a separate network map for each leaf domain. As a result, a leaf domain can contain an IP address that is unique within its network, but identical to an IP address in another leaf domain. When you view event statistics in an ancestor domain, the system may display multiple instances of that repeated IP address. At first glance, they might appear to be duplicate entries. However, if you drill down to the host profile information for each IP address, the system shows that they belong to different leaf domains.

View intrusion event performance graphs

Graphs can be generated to reflect number of intrusion events per second, number of megabits per second, average number of bytes per packet, the percent of packets uninspected by Snort, and the number of packets blocked as the result of TCP normalization. These graphs can show statistics for the last hour, last day, last week, or last month of operation.

Note New data is accumulated for statistics graphs every five minutes. Therefore, if you reload a graph quickly, the data may not change until the next five-minute increment occurs. Each graph displays *average* values in the intervals shown (day, hour, or five minutes) for the selected time period (last month, week, day, or hour). Decimal values are displayed when the average is less than one.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Overview** > **Summary** > **Intrusion Event Performance**.

Step 2 From the **Select Device** list, choose the devices whose data you want to view.

Step 3 From the **Select Graph(s)** list, choose the type of graph you want to create as described in Intrusion event performance statistics graph types, on page 828.

Step 4 From the **Select Time Range** list, choose the time range you would like to use for the graph.

Step 5 Click **Graph**.

Step 6 To save the graph, right-click it and follow the instructions for your browser to save the image.

Intrusion event performance statistics graph types

The table lists the available graph types for intrusion event performance statistics. Graph types may display differently if they are populated with data affected by the network analysis policy **Inline Mode** setting. If

Inline Mode is disabled, graph types marked with an asterisk (*) in the web interface (a `yes` in the column below) populate with data about the traffic the system would have modified or dropped if **Inline Mode** were enabled.

Table 115: Intrusion event performance graph types

To generate data for...	You must...	Which represents...	Affected by Inline Mode?
Avg Bytes/Packet	n/a	the average number of bytes included in each packet.	no
ECN Flags Normalized in TCP Traffic/Packet	enable **Explicit Congestion Notification** and select **Packet**	the number of packets for which ECN flags have been cleared on a per-packet basis regardless of negotiation.	yes
ECN Flags Normalized in TCP Traffic/Session	enable **Explicit Congestion Notification** and select **Stream**	the number of times that ECN flags have been cleared on a per-stream basis when ECN use was not negotiated.	yes
Events/Sec	n/a	the number of events per second generated on the device.	no
ICMPv4 Echo Normalizations	enable **Normalize ICMPv4**	the number of ICMPv4 packets for which the 8-bit Code field in Echo (Request) or Echo Reply messages were cleared.	yes
ICMPv6 Echo Normalizations	enable **Normalize ICMPv6**	the number of ICMPv6 packets for which the 8-bit Code field in Echo (Request) or Echo Reply messages was cleared.	yes
IPv4 DF Flag Normalizations	enable **Normalize IPv4** and **Normalize Don't Fragment Bit**	the number of IPv4 packets for which the single-bit Don't Fragment subfield of the IPv4 Flags header field was cleared.	yes
IPv4 Options Normalizations	enable **Normalize IPv4**	the number of IPv4 packets for which the option octet was set to 1 (No Operation).	yes
IPv4 Reserved Flag Normalizations	enable **Normalize IPv4** and **Normalize Reserved Bit**	the number of IPv4 packets for which the single-bit Reserved subfield of the IPv4 Flags header field was cleared.	yes
IPv4 Resize Normalizations	enable **Normalize IPv4**	the number of IPv4 packets with excessive-length payload that have been truncated to the datagram length specified in the IP header.	yes
IPv4 TOS Normalizations	enable **Normalize IPv4** and **Normalize TOS Bit**	the number of IPv4 packets for which the one-byte Differentiated Services (DS) field (formerly known as the Type of Service (TOS) field) was cleared.	yes
IPv4 TTL Normalizations	enable **Normalize IPv4** , **Maximum TTL** , and **Reset TTL**	the number of IPv4 Time to Live normalizations.	yes

To generate data for...	You must...	Which represents...	Affected by Inline Mode?
IPv6 Options Normalizations	enable **Normalize IPv6**	the number of IPv6 packets for which the Option Type field in the Hop-by-Hop Options or Destination Options extension header was set to `00` (Skip and continue processing).	yes
IPv6 TTL Normalizations	enable **Normalize IPv6** , **Minimum TTL** , and **Reset TTL**	the number of IPv6 Hop Limit (TTL) normalizations.	yes
Mbits/Sec	n/a	the number of megabits per second of traffic that passes through the device.	no
Packet Resized to Fit MSS Normalizations	enable **Trim Data to MSS**	the number of packets for which the payload was longer than the TCP Data field, so the payload was trimmed to the Maximum Segment Size.	yes
Packet Resized to Fit TCP Window Normalizations	enable **Trim Data to Window**	the number of packets for which the TCP Data field was trimmed to fit the receiving host's TCP window.	yes
Percent Packets Dropped	n/a	the average percentage of uninspected packets across all selected devices. For example, if you select two devices, then an average of 50% may indicate that one device has a 90% drop rate and the other has a 10% drop rate. It may also indicate that both devices have a drop rate of 50%. The graph only represents the total % drop when you select a single device.	no
RST Packets With Data Stripped Normalizations	enable **Remove Data on RST**	the number of packets for which data was removed from a TCP reset (RST) packet.	yes
SYN Packets With Data Stripped Normalizations	enable **Remove Data on SYN**	the number of packets for which data was removed from SYN packets when the TCP operating system was not Mac OS.	yes
TCP Header Padding Normalizations	enable **Normalize/Clear Option Padding Bytes**	the number of TCP packets in which option padding bytes were set to 0.	yes
TCP No Option Normalizations	enable **Allow These TCP Options** and set to an option other than `any`	the number of packets from which the Time Stamp option was stripped.	yes
TCP NS Flag Normalizations	enable **Explicit Congestion Notification** and select **Packet**	the number of ECN Nonce Sum (NS) option normalizations.	yes
TCP Options Normalizations	enable **Allow These TCP Options** and set to an option other than `any`	the number of options (excluding MSS, Window Scale, Time Stamp, and explicitly allowed options) for which the option field is set to No Operation (TCP Option 1).	yes

To generate data for...	You must...	Which represents...	Affected by Inline Mode?
TCP Packets Blocked By Normalizations	enable **Normalize TCP Payload** (segment reassembly must fail)	the number of packets dropped because the TCP segments could not be properly reassembled.	yes
TCP Reserved Flags Normalizations	enable **Normalize/Clear Reserved Bits**	the number of TCP packets where the Reserved bits have been cleared.	yes
TCP Segment Reassembly Normalizations	enable **Normalize TCP Payload** (segment reassembly must be successful)	the number of packets for which the TCP Data field was normalized to ensure consistency in retransmitted data (any segments that cannot be properly reassembled are dropped).	yes
TCP SYN Option Normalizations	enable **Allow These TCP Options** and set to an option other than `any`	the number of options for which the Maximum Segment Size or Window Scale option was set to No Operation (TCP Option 1) because the SYN control bit was not set.	yes
TCP Timestamp ECR Normalizations	enable **Allow These TCP Options** and set to an option other than `any`	the number of packets for which the Time Stamp Echo Reply (TSecr) option field was cleared because the Acknowledgment (ACK) control bit was not set.	yes
TCP Urgent Pointer Normalizations	enable **Normalize Urgent Pointer**	the number of packets for which the two-byte TCP header Urgent Pointer field was greater than the payload length and was set to the payload length.	yes
Total Blocked Packets	configure **Inline Mode** or **Drop when Inline**	the total number of dropped packets, including rule, decoder, and preprocessor drops.	no
Total Injected Packets	configure **Inline Mode**	the number of packets that were resized before being retransmitted.	no
Total TCP Filtered Packets	configure TCP Stream Preprocessing	the number of packets skipped by the stream because of TCP port filtering.	no
Total UDP Filtered Packets	configure UDP Stream Preprocessing	the number of packets skipped by the stream because of UDP port filtering.	no
Urgent Flag Cleared Normalizations	enable **Clear URG if Urgent Pointer is Not Set**	the number of packets for which the TCP header URG control bit was cleared because the urgent pointer was not set.	yes
Urgent Pointer and Urgent Flag Cleared Normalizations	enable **Clear Urgent Pointer/URG on Empty Payload**	the number of packets for which the TCP header Urgent Pointer field and the URG control bit have been cleared because there was no payload.	yes
Urgent Pointer Cleared Normalizations	enable **Clear Urgent Pointer if URG=0**	the number of packets for which the 16-bit TCP header Urgent Pointer field was cleared because the urgent (URG) control bit was not set.	yes

View intrusion event graphs

The system provides graphs that show you intrusion event trends over time. You can generate intrusion event graphs over time ranging from the last hour to the last month, for one or all managed devices.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Overview** > **Summary** > **Intrusion Event Graphs**.

Step 2 Under **Select Device**, choose **all** to include all devices, or choose the specific device you want to include in the graph.

Step 3 Under **Select Graph(s)**, choose the type of graph you want to generate:

- Top 10 Destination Ports
- Top 10 Source IP Addresses
- Top 10 Event Messages

Step 4 Under **Select Time Range**, choose the time range for the graph:

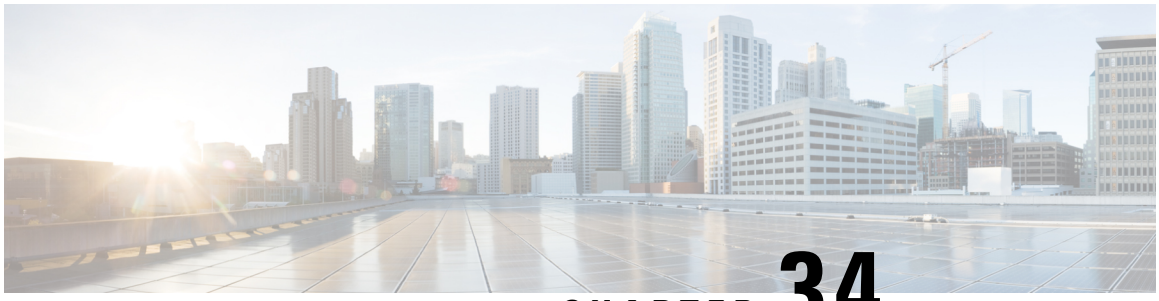
- Last Hour
- Last Day
- Last Week
- Last Month

Step 5 Click **Graph**.

History for intrusion events

This reference provides the historical changes and feature additions for intrusion events across different software versions.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
IPS Events Datastore Replacement	7.1	Any	• Intrusion incidents, the intrusion event clipboard, and default custom tables (that use the intrusion event columns - **Intrusion Events with Source Criticality** and **Intrusion Events with Destination Criticality**) are deprecated. You can no longer add events to the clipboard using the **Copy** and **Copy All** buttons. Deprecated pages: • **Analysis** > **Intrusions** > **Clipboard** • **Analysis** > **Intrusions** > **Incidents** • Two new fields are added to the main intrusion event table **- Source Host Criticality** and **Destination Host Criticality**. Supported platforms: Secure Firewall Management Center
Unique identifier for connection event in syslogs	6.4.0.4	Any	These syslog fields collectively uniquely identify a connection event and appear in syslogs for intrusion events: DeviceUUID, First Packet Time, Connection Instance ID, and Connection Counter.
IntrusionPolicy field is now included in syslog	6.4	Any	Intrusion event syslogs now specify the intrusion policy that triggered the event.
New intrusion event search field: CVE ID	6.4	Any	You can now search by MITRE's Common Vulnerabilities and Exposures identification number Modified screens: **Analysis** > **Intrusions** > **Events** > **Edit Search** Supported Platforms: All.



CHAPTER 34

File/Malware Events and Network File Trajectory

The following topics provide an overview of file and malware events, local malware analysis, dynamic analysis, captured files, and network file trajectories.

- File and malware events and network file trajectory, on page 835
- File and malware events, on page 836
- View details about analyzed files, on page 855
- Using captured file workflows, on page 858
- Submit files for analysis manually, on page 862
- Network file trajectory, on page 863
- History for file and malware events and network file trajectory, on page 869

File and malware events and network file trajectory

File and malware events are security data logs that

- are automatically generated and logged when matched traffic triggers a file policy,
- trigger automatic logging of the associated connection to the Secure Firewall Management Center database when generated, and
- provide analysis data to address negative impacts and block future attacks.

File analysis and network file trajectory features

Based on file analysis results, you can review captured files and generated malware and file events using tables on pages available under the **Analysis** > **Files** > **File Events**. These include details such as file composition, disposition, threat score, and dynamic analysis summary reports.

To further target your analysis, you can use a malware file's *network file trajectory* (a map of how the file traversed your network, passing among hosts, as well as various file properties) to track the spread of an individual threat across hosts over time, allowing you to concentrate outbreak control and prevention efforts where most useful.

If you configure local malware analysis or dynamic analysis in a file rule, the system preclassifies files matching the rule and generates a file composition report.

If your organization has deployed *Secure Endpoint* and integrated that deployment with your Secure Firewall Management Center, you can also import records of scans, malware detections, and quarantines, as well as

indications of compromise (IOC) identified by that product. This data is displayed alongside event data gathered by Secure Firewall for a more complete picture of malware on your network.

The Context Explorer, dashboards, and reporting features can also aid a deeper understanding of the files and malware detected, captured, and blocked. You can also use events to trigger correlation policy violations, or alert you via email, SMTP, or syslog.



Note To configure your system to detect malware and generate file and malware events, refer to *Network Malware Protection and File Policies* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

File and malware events

File and malware events are logging records that

- represent files and malware detected by the Secure Firewall Management Center
- provide different information depending on how and why they were generated, and
- include file events, malware events, and retrospective malware events with varying field availability.

Event types

The system generates different types of file and malware events:

- *File events*: Events that represent files, including malware, detected by the system (malware defense). File events do not contain Secure Endpoint-related fields.
- *Malware events*: Events that represent malware detected by either malware defense or Secure Endpoint. Malware events can also record data other than threats from your Secure Endpoint deployment, such as scans and quarantines.
- *Retrospective malware events*: Events that represent files detected by malware defense whose dispositions (whether the files are malware) have changed.



Note

- Files identified as malware by malware defense generate both a file event and a malware event. Malware events generated by Secure Endpoint do not have corresponding file events.
- File events generated by inspecting NetBIOS-ssn (SMB) traffic do not immediately generate connection events because the client and server establish a persistent connection. The system generates connection events after the client or server ends the session.
- The system supports the display and input of file names that use Unicode (UTF-8) characters. However, Unicode file names appear in PDF reports in transliterated form. Additionally, the SMB protocol replaces unprintable characters in file names with periods.

File and malware event types

File and malware event types are security classifications that identify and categorize different kinds of file-related activities and malicious software incidents detected within network traffic for analysis and threat response.

File events

The system logs the file events generated when a managed device detects or blocks a file in network traffic, according to the rules in currently deployed file policies.

When the system generates a file event, it also logs the end of the associated connection to the Secure Firewall Management Center database, regardless of the logging configuration of the invoking access control rule.

Malware events

The system (specifically the malware defense feature) generates malware events when it detects malware in network traffic as part of your overall access control configuration. Malware events contain the disposition of the resulting event and contextual data about how, where, and when the malware was detected.

Table 116: Malware event generation scenarios

When the system detects a file and...	Disposition
successfully queries the AMP cloud (performs a malware cloud lookup) for the file's disposition	Malware, Clean, or Unknown
queries the AMP cloud but cannot establish a connection or the cloud is otherwise unavailable	Unavailable You may see a small percentage of events with this disposition; this is expected behavior.
the threat score associated with a file exceeds the malware threshold threat score defined in the file policy that detected the file, or local malware analysis identifies malware	Malware
it is on the custom detection list (manually marked as malware)	Custom Detection
it is on the on the clean list (manually marked as clean),	Clean

File rule actions and file disposition in malware events

Each file rule defines an action that determines how the system manages traffic matching the conditions of the rule. If you select *Block Malware* or *Malware Cloud Lookup* as file rule action, the system queries the AMP cloud to assess if files traversing the network contain malware, and blocks files identified as threats. A cloud lookup allows the system to retrieve and log the file's disposition based on its SHA-256 hash value.

The table maps file rule actions to file disposition and the resulting action in malware events.

Table 117: File disposition and file action in malware events

File Rule Action Selected	File Disposition	File Action in the Malware Event
• Block Malware • Malware Cloud Lookup	Malware	Block
	• Clean • Unknown • Unavailable • NA	Cloud Lookup **Note** Under the file policy editor Advanced Settings, you can set a threshold threat score for **If AMP Cloud disposition is Unknown, override disposition based upon threat score** option. If you set a threshold threat score, files with an AMP cloud verdict of Unknown are considered malware if their dynamic analysis score is equal to or worse than the threshold.

Retrospective malware events

A retrospective malware event is a security event that

- represents a disposition change for files detected in the last week that have the same SHA-256 hash value,
- contains limited information including the date and time of notification, the new disposition, the SHA-256 hash value, and the threat name, and
- does not contain IP addresses or other contextual information.

Retrospective malware event processing

For malware detected in network traffic, dispositions can change. For example, the AMP cloud can determine that a file that was previously thought to be clean is now identified as malware, or the reverse—that a malware-identified file is actually clean. When the disposition changes for a file you queried in the last week, the AMP cloud notifies the system. Then, two things happen:

- The Secure Firewall Management Center generates a new retrospective malware event.

  This new retrospective malware event represents a disposition change for all files detected in the last week that have the same SHA-256 hash value. For that reason, these events contain limited information: the date and time the Secure Firewall Management Center was notified of the disposition change, the new disposition, the SHA-256 hash value of the file, and the threat name. They do not contain IP addresses or other contextual information.

- The Secure Firewall Management Center changes the file disposition for previously detected files with the retrospective event's associated SHA-256 hash value.

  If a file's disposition changes to Malware, the Secure Firewall Management Center logs a new malware event to its database. Except for the new disposition, the information in this new malware event is identical to that in the file event generated when the file was initially detected.

  If a file's disposition changes to Clean, the Secure Firewall Management Center does not delete the malware event. Instead, the event reflects the change in disposition. This means that files with clean dispositions can appear in the malware table, but only if they were originally thought to be malware. Files that were never identified as malware appear only in the files table.

Malware events generated by Secure Endpoint

A malware event generated by Secure Endpoint is a security event that

- occurs when lightweight connectors on endpoints inspect files and communicate with the AMP cloud to determine if files contain malware,
- includes threat identifications, scans, quarantines, blocked executions, and cloud recalls sent from the AMP cloud to the Secure Firewall Management Center, and
- is logged by the Secure Firewall Management Center when files are positively identified as malware.

Endpoint connector functionality

Organizations using Secure Endpoint install lightweight connectors on *endpoints*: computers and mobile devices. Connectors can inspect files upon upload, download, execution, open, copy, move, and so on. These connectors communicate with the AMP cloud to determine if inspected files contain malware.

When a file is positively identified as malware, the AMP cloud sends the threat identification to the Secure Firewall Management Center. The AMP cloud can also send other kinds of information to the Secure Firewall Management Center, including data on scans, quarantines, blocked executions, and cloud recalls. The Secure Firewall Management Center logs this information as malware events.



Note The IP addresses reported in malware events generated by Secure Endpoint may not be in your network map—and may not even be in your monitored network at all. Depending on your deployment, level of compliance, and other factors, endpoints in your organization monitored by Secure Endpoint may not be the same hosts as those monitored by malware defense.

Malware event analysis with Secure Endpoint

Malware event analysis with Secure Endpoint is a security feature that

- displays malware events detected by Secure Endpoint on Firewall Management Center event pages alongside events detected by Malware Defense,
- enables viewing of file trajectory and other information about a particular SHA in Secure Endpoint when using the Cisco public cloud, and
- provides integrated malware detection capabilities across endpoint and network security solutions.

If your organization has deployed Cisco Secure Endpoint:

- You can configure the system to display malware events detected by Secure Endpoint on Firewall Management Center event pages, alongside events detected by Malware Defense.
- If you are using the AMP public cloud, you can view file trajectory and other information about a particular SHA in Secure Endpoint. Simply right-click a file's SHA hash in a table on an event page.

To configure this functionality, refer to *Integrate DeviceUUID (Syslog Only)Secure Firewall and Secure Endpoint* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Event data from Secure Endpoint

Event data from Secure Endpoint is file and malware detection information that

- provides in-depth malware event information with endpoint context such as file path and invoking client application,
- differs from network-based detection in timing and data type, as malware detection occurs at the endpoint during download or execution, and
- integrates with Firewall Management Center to provide comprehensive malware protection data.

Data differences between detection methods

Malware events detected by Secure Endpoint contain different information than malware events detected by managed devices in network traffic.

Because Secure Endpoint malware detection is performed at the endpoint at download or execution time, while managed devices detect malware in network traffic, the information in the two types of malware events is different. For example, malware events detected by Secure Endpoint ("endpoint-based malware") contain information on file path, invoking client application, and so on, while malware detections in network traffic contain port, application protocol, and originating IP address information about the connection used to transmit the file.

As another example, for malware events detected by Malware Defense ("network-based malware events"), user information represents the user most recently logged into the host where the malware was destined, as determined by network discovery. But Secure Endpoint-reported users represent the user currently logged into the endpoint where the malware was detected.



Note Depending on your deployment, endpoints monitored by Secure Endpoint may not be the same hosts as those monitored by malware defense. For this reason, malware events generated by Secure Endpoint do not add hosts to the network map. However, the system uses IP and MAC address data to tag monitored hosts with indications of compromise obtained from your Secure Endpoint deployment. If two different hosts monitored by different malware solutions have the same IP and MAC address, the system can incorrectly tag monitored hosts with Secure Endpoint IOCs.

This table summarizes the differences between the event data generated by Firepower when using a Malware license, and event data generated by Secure Endpoint.

Table 118: Summary of data differences between AMP products

Feature	malware defense	Secure Endpoint
Events generated	File events, captured files, malware events, and retrospective malware events	Malware events
Information in malware events	Basic malware event information, plus connection data (IP address, port, and application protocol)	In-depth malware event information; no connection data
Network file trajectory	Firewall Management Center-based	Firewall Management Center and the Secure Endpoint management console each have a network file trajectory. Both are useful.

Related topics

Integrate Firepower and Secure Endpoint in the *Cisco Secure Firewall Management Center Device Configuration Guide*

Using file and malware event workflows

This procedure allows you to view file and malware events in a table format and manipulate the event view to display information relevant to your analysis.

The page you see when you access events differs depending on the workflow, which is simply a series of pages you can use to evaluate events by moving from a broad to a more focused view. You can also create a custom workflow that displays only the information that matches your specific needs.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

You must be an Admin or Security Analyst user to perform this task.

Procedure

Choose one of the following:

- **Analysis** > **Files** > **File Events**
- **Analysis** > **Files** > **Malware Events**

Tip
In the table view of events, several fields are hidden by default. To show a hidden field in an event view, expand the search constraints, then click the field name under **Disabled Columns**.

Tip
To quickly view the connections where specific files were detected, choose the files using the check boxes in the table, then choose **Connections Events** from the **Jump to** drop-down list.

Tip
Right-click an item in the table to see options. (Not every column offers options.)

Related Topics

File and malware event fields, on page 841
Predefined file workflows, on page 669
Predefined malware workflows, on page 669
Configure event view settings, on page 234

File and malware event fields

File and malware events, which you can view and search using workflows, contain the fields listed in this section. Keep in mind that the information available for any individual event can vary depending on how and why it was generated.

Note Files identified as malware by malware defense generate both a file event and a malware event. Malware events generated by Secure Endpoint do not have corresponding file events, and file events do not have Secure Endpoint-related fields.

Syslog messages are populated with initial values and do not update, even if the equivalent field in the Firewall Management Center web interface is updated, for example with a retrospective verdict.

Action (Syslog: FileAction)

The action associated with file policy rule that detected the file, and any associated file rule action options.

AMP Cloud

The name of the AMP cloud where the AMP for Endpoints event originated.

Application File Name

The client application accessing the malware file when AMP for Endpoints detection occurred. These applications are **not** tied to network discovery or application control.

Application File SHA256

The SHA-256 hash value of the parent file accessing the AMP for Endpoints-detected or quarantined file when detection occurred.

In the unified event viewer, this field appears as **Application File SHA-256**.

Application Protocol (Syslog: ApplicationProtocol)

The application protocol used by the traffic in which a managed device detected the file.

Application Protocol Category or Tag

The criteria that characterize the application to help you understand the application's function.

Application Risk

The risk associated with the application traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of application detected in the connection has an associated risk; this field displays the highest of those.

Archive Depth (Syslog: ArchiveDepth)

The level (if any) at which the file was nested in an archive file.

Archive Name (Syslog: ArchiveFileName)

The name of the archive file (if any) which contained the malware file.

To view the contents of an archive file, go to any table under **Analysis** > **Files** > **File Events** that lists the archive file, right-click the archive file's table row to open the context menu, then click **View Archive Contents**.

Archive SHA256 (Syslog: ArchiveSHA256)

The SHA-256 hash value of the archive file (if any) which contains the malware file.

To view the contents of an archive file, go to any table under **Analysis** > **Files** > **File Events** that lists the archive file, right-click that archive file's table row to open the context menu, then click **View Archive Contents**.

ArchiveFileStatus (Syslog Only)

The status of an archive being inspected. Can have the following values:

- Pending — Archive is being inspected
- Extracted — Successfully inspected without any problems
- Failed — Failed to inspect, insufficient system resources
- Depth Exceeded — Successful, but archive exceeded the nested inspection depth
- Encrypted — Partially successful, archive was or contains an archive that is encrypted
- Not Inspectable — Partially successful, file is possibly malformed or corrupt

Business Relevance

The business relevance associated with the application traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of application detected in the connection has an associated business relevance; this field displays the lowest (least relevant) of those.

Category / File Type Category

The general categories of file type, for example: Office Documents, Archive, Multimedia, Executables, PDF files, Encoded, Graphics, or System Files.

Client (Syslog: Client)

The client application that runs on one host and relies on a server to send a file.

Client Category or Tag

The criteria that characterize the application to help you understand the application's function.

Connection Counter (Syslog Only)

A counter that distinguishes one connection from another simultaneous connection. This field has no significance on its own.

The following fields collectively uniquely identify the connection event associated with a particular file or malware event: DeviceUUID, First Packet Time, Connection Instance ID, and Connection Counter.

Connection Instance ID (Syslog Only)

The Snort instance that processed the connection event. This field has no significance on its own.

The following fields collectively uniquely identify the connection event associated with a particular file or malware event: DeviceUUID, First Packet Time, Connection Instance ID, and Connection Counter.

Count

After you apply a constraint that creates two or more identical rows, the number of events that match the information in each row.

Detection Name

The name of the detected malware.

Note Malware events involving archive files occur when malware is detected inside archive files, such as ZIP or other compressed files. The top-level archive event may not show a detection name because the archive itself is not malicious. Detection names are assigned only to the infected child files within the archive. Use the archive's SHA-256 hash or the archive file name to map and correlate the parent archive event with its infected child files.

Detector

The AMP for Endpoints detector that identified the malware, such as ClamAV, Spero, or SHA.

Device

For file events and for malware events generated by firewall devices, the name of the device that detected the file.

For malware events generated by AMP for Endpoints and for retrospective malware events generated by the AMP cloud, the name of the Firewall Management Center.

DeviceUUID (Syslog Only)

The unique identifier of the firewall device that generated an event.

The following fields collectively uniquely identify the connection event associated with a particular file or malware event: DeviceUUID, First Packet Time, Connection Instance ID, and Connection Counter.

Disposition / File Disposition (Syslog: SHA_Disposition)

The file's disposition:

Malware

Indicates that the AMP cloud categorized the file as malware, local malware analysis identified malware, or the file's threat score exceeded the malware threshold defined in the file policy.

Clean

Indicates that the AMP cloud categorized the file as clean, or that a user added the file to the clean list. Clean files appear in the malware table only if they were changed to clean.

Unknown

Indicates that the system queried the AMP cloud, but the file has not been assigned a disposition; in other words, the AMP cloud has not categorized the file.

Custom Detection

Indicates that a user added the file to the custom detection list.

Unavailable

Indicates that the system could not query the AMP cloud. You may see a small percentage of events with this disposition; this is expected behavior.

N/A

Indicates a Detect Files or Block Files rule handled the file and the Secure Firewall Management Center did not query the AMP cloud.

File dispositions appear only for files for which the system queried the AMP cloud.

Syslog fields reflect only the initial disposition; they do not update to reflect retrospective verdicts.

Domain

For file events and for malware events generated by firewall devices, the domain of the device that detected the file. For malware events generated by AMP for Endpoints and for retrospective malware events generated by the AMP cloud, the domain associated with the AMP cloud connection that reported the event.

This field is only present if you have ever configured the Firewall Management Center for multitenancy.

DstIP (Syslog Only)

The IP address of the host that responded to the connection. This may be the IP address of the sender or the recipient of the file, depending on the value in the FileDirection field:

If FileDirection is `Upload`, then this is the IP address of the file recipient.

If FileDirection is `Download`, then this is the IP address of the file sender.

See also **SrcIP**.

See also Initiator and responder, source and destination, and sender and receiver fields, on page 773.

DstPort (Syslog Only)

The port used in the connection described under **DstIP**.

Egress Virtual Router

In networks using virtual routing, the name of the virtual router through which traffic exited the network.

Event Subtype

The AMP for Endpoints action that led to malware detection, for example, Create, Execute, Move, or Scan.

Event Type

The sub-type of malware event.

File Name (Syslog: FileName)

The name of the file.

File Path

The file path of the malware file detected by AMP for Endpoints, not including the file name.

File Policy (Syslog: FilePolicy)

The file policy that detected the file.

File Storage / Stored (Syslog: FileStorageStatus)

The storage status of the file associated with the event:

Stored

Returns all events where the associated file is currently stored.

Stored in connection

Returns all events where the system captured and stored the associated file, regardless of whether the associated file is currently stored.

Failed

Returns all events where the system failed to store the associated file.

Syslog fields contain only the initial status; they do not update to reflect changed status.

File Timestamp

The time and date that AMP for Endpoints detected the malware file was created.

FileDirection (Syslog Only)

Whether the file was downloaded or uploaded during the connection. Possible values are:

- Download — the file was transferred from the DstIP to the SrcIP.
- Upload — the file was transferred from the SrcIP to the DstIP.

FileSandboxStatus (Syslog Only)

Indicates whether the file was sent for dynamic analysis and if so, the status.

First Packet Time (Syslog Only)

The time the system encountered the first packet.

The following fields collectively uniquely identify the connection event associated with a particular file or malware event: DeviceUUID, First Packet Time, Connection Instance ID, and Connection Counter.

FirstPacketSecond (Syslog Only)

The time at which the file download or upload flow started.

The time the event occurred is captured in the message header timestamp.

HTTP Response Code

The HTTP status code sent in response to a client's HTTP request when a file is transferred.

Ingress Virtual Router

In networks using virtual routing, the name of the virtual router through which traffic entered the network.

IOC

Whether the malware event triggered an indication of compromise (IOC) against a host involved in the connection. When AMP for Endpoints data triggers an IOC rule, a full malware event is generated, with the type AMP IOC.

Message

Additional information associated with a malware event. For file events and for malware events generated by firewall devices, this field is populated only for files whose disposition has changed, that is, that have an associated retrospective event.

Protocol (Syslog Only)

The protocol used for the connection, for example TCP or UDP.

Receiving Continent

The continent of the host receiving the file.

Receiving Country

The country of the host receiving the file.

Receiving IP

In the Firewall Management Center web interface, for file events and for malware events generated by firewall devices, the IP address of the host receiving the file. See also Initiator and responder, source and destination, and sender and receiver fields, on page 773.

For malware events generated by AMP for Endpoints, the IP address of the endpoint whose connector reported the event.

For syslog equivalents (events generated by firewall devices only), see **DstIP** and **SrcIP**.

Receiving Port

In the Firewall Management Center web interface, the destination port used by the traffic where the file was detected.

For syslog equivalents, see **DstIP** and **SrcIP** and **DstPort** and **SrcPort**.

Security Context (Syslog: Context)

The metadata identifying the virtual firewall group through which the traffic passed. Note that the system only displays this field when managing at least one ASA FirePOWER device that is running in multiple context mode.

Sending Continent

The continent of the host sending the file.

Sending Country

The country of the host sending the file.

Sending IP

In the Firewall Management Center web interface, the IP address of the host sending the file. See also Initiator and responder, source and destination, and sender and receiver fields, on page 773.

For syslog equivalents, see **DstIP** and **SrcIP**.

Sending Port

In the Firewall Management Center web interface, the source port used by the traffic where the file was detected.

For syslog equivalents, see **DstIP** and **SrcIP** and **DstPort** and **SrcPort**.

SHA256 / File SHA256 (Syslog: FileSHA256)

The SHA-256 hash value of the file.

To have a SHA256 value, the file must have been handled by one of:

- a Detect Files file rule with **Store files** enabled
- a Block Files file rule with **Store files** enabled
- a Malware Cloud Lookup file rule
- a Block Malware file rule
- AMP for Endpoints

This column also displays a network file trajectory icon that represents the most recently detected file event and file disposition, and that links to the network file trajectory.

Size (KB) / File Size (KB) (Syslog: FileSize)

In the Firewall Management Center web interface, the size of the file, in kilobytes.

In syslog messages: The size of the file, in bytes.

Note that if the system determines the file type of a file before the file is fully received, the file size may not be calculated. In this case, this field is blank.

SperoDisposition (Syslog Only)

Indicates whether the SPERO signature was used in file analysis. Possible values:

- Spero detection performed on file
- Spero detection not performed on file

SrcIP (Syslog Only)

The IP address of the host that initiated the connection. This may be the IP address of the sender or the recipient of the file, depending on the value in the FileDirection field:

If FileDirection is **Upload**, this is the IP address of the file sender.

If FileDirection is **Download**, this is the IP address of the file recipient.

See also **DstIP**.

See also Initiator and responder, source and destination, and sender and receiver fields, on page 773.

SrcPort (Syslog Only)

The port used in the connection described under **SrcIP**.

SSL Actual Action (Syslog: SSLActualAction)

The action the system applied to encrypted traffic:

Block or Block with reset

Represents blocked encrypted connections.

Decrypt (Resign)

Represents an outgoing connection decrypted using a re-signed server certificate.

Decrypt (Replace Key)

Represents an outgoing connection decrypted using a self-signed server certificate with a substituted public key.

Decrypt (Known Key)

Represents an incoming connection decrypted using a known private key.

Default Action

Indicates the connection was handled by the default action.

Do not Decrypt

Represents a connection the system did not decrypt.

Field values are displayed in the **SSL Status** field on the search workflow pages.

SSL Certificate Information

The information stored on the public key certificate used to encrypt traffic, including:

- Subject/Issuer Common Name
- Subject/Issuer Organization
- Subject/Issuer Organization Unit
- Not Valid Before/After
- Serial Number, Certificate Fingerprint
- Public Key Fingerprint

For syslog, see **SSLCertificate**.

SSL Failure Reason (Syslog: SSLFlowStatus)

The reason the system failed to decrypt encrypted traffic:

- Unknown
- No Match
- Success
- Uncached Session
- Unknown Cipher Suite
- Unsupported Cipher Suite
- Unsupported SSL Version
- SSL Compression Used
- Session Undecryptable in Passive Mode
- Handshake Error
- Decryption Error
- Pending Server Name Category Lookup
- Pending Common Name Category Lookup
- Internal Error
- Network Parameters Unavailable
- Invalid Server Certificate Handle
- Server Certificate Fingerprint Unavailable
- Cannot Cache Subject DN
- Cannot Cache Issuer DN
- Unknown SSL Version
- External Certificate List Unavailable
- External Certificate Fingerprint Unavailable
- Internal Certificate List Invalid
- Internal Certificate List Unavailable
- Internal Certificate Unavailable
- Internal Certificate Fingerprint Unavailable
- Server Certificate Validation Unavailable
- Server Certificate Validation Failure
- Invalid Action

Field values are displayed in the **SSL Status** field on the search workflow pages.

SSL Status

The action associated with the **SSL Actual Action** (SSL rule, default action, or undecryptable traffic action) that logged the encrypted connection. The **Lock icon** links to TLS/SSL certificate details. If the certificate is unavailable (for example, for connections blocked due to TLS/SSL handshake error), the lock icon is grayed out.

If the system fails to decrypt an encrypted connection, it displays the **SSL Actual Action** (undecryptable traffic action) taken, as well as the **SSL Failure Reason**. For example, if the system detects traffic encrypted with an unknown cipher suite and allows it without further inspection, this field displays `Do Not Decrypt (Unknown Cipher Suite)`.

When searching this field, type one or more of the **SSL Actual Action** and **SSL Failure Reason** values to view encrypted traffic the system handled or failed to decrypt.

SSL Subject/Issuer Country

The two-character ISO 3166-1 alpha-2 country code for the subject or issuer country associated with the encryption certificate.

SSLCertificate (Syslog Only)

The certificate fingerprint of the TLS/SSL server.

Threat Name (Syslog: ThreatName)

The name of the detected malware.

Threat Score (Syslog: ThreatScore)

The threat score most recently associated with this file. This is a value from 0 to 100 based on the potentially malicious behaviors observed during dynamic analysis.

The threat score icon links to the Dynamic Analysis Summary report.

Time

The date and time the event was generated. This field is not searchable.

In syslog messages, see **FirstPacketSecond**.

Type / File Type (Syslog: FileType)

The type of file, for example, HTML or MSEXE.

URI / File URI (Syslog: URI)

The URI of the connection associated with the file transaction, for example, the URL from which a user downloaded the file.

User (Syslog: User)

The username associated with the IP address that initiated the connection. If this IP address is external to your network, the associated username is typically unknown.

If applicable, the username is preceded by `<realm>\`.

For file events and for malware events generated by firewall devices, this field displays the username that was determined by an identity policy or authoritative logins. In absence of an identity policy, it displays No Authentication Required.

For malware events generated by AMP for Endpoints, AMP for Endpoints determines user names. These users *cannot* be tied to user discovery or control. They do not appear in the Users table, nor can you view details for these users.

Web Application (Syslog: WebApplication)

The application that represents the content or requested URL for HTTP traffic detected in the connection.

Web Application Category or Tag

Criteria that characterize the application to help you understand the application's function.

Malware event sub-types

The table lists the malware event subtypes, whether a malware event generated by Malware Defense (a "network-based malware event") or Secure Endpoint (an "endpoint-based malware event") can have that subtype, and whether the system uses that subtype to build network file trajectories.

Table 119: Malware event types

Malware Event Subtype/Search Value	malware defense	Secure Endpoint	File Trajectory
Threat Detected in Network File Transfer	yes	no	yes
Threat Detected in Network File Transfer (retrospective)	yes	no	yes
Threat Detected	no	yes	yes
Threat Detected in Exclusion	no	yes	yes
Threat Quarantined	no	yes	yes
AMP IOC (Indications of compromise)	no	yes	no
Blocked Execution	no	yes	no
Cloud Recall Quarantine	no	yes	no
Cloud Recall Quarantine Attempt Failed	no	yes	no
Cloud Recall Quarantine Started	no	yes	no
Cloud Recall Restore from Quarantine	no	yes	no
Cloud Recall Restore from Quarantine Failed	no	yes	no
Cloud Recall Restore from Quarantine Started	no	yes	no
Quarantine Failure	no	yes	no

Malware Event Subtype/Search Value	malware defense	Secure Endpoint	File Trajectory
Quarantined Item Restored	no	yes	no
Quarantine Restore Failed	no	yes	no
Quarantine Restore Started	no	yes	no
Scan Completed, No Detections	no	yes	no
Scan Completed With Detections	no	yes	no
Scan Failed	no	yes	no
Scan Started	no	yes	no

Information available in file and malware event fields

Lists the information available in file and malware event fields for system displays and integrations.

If your organization has deployed Secure Endpoint and integrated that product with your Secure Firewall deployment:

- Malware events and indications of compromise (IOCs) imported from your Secure Endpoint deployment do not contain contextual connection information, but they do include information obtained at download or execution time, such as file path, invoking client application, and so on.
- File event table views do not display Secure Endpoint-related fields.

The table shows which fields are available for various types of file and malware events.

Table 120: Information available in file and malware event fields

Field	File Event	Malware Events Detected by the System	Retrospective Events Detected by the System	Malware Events Detected by Secure Endpoint
Action	yes	yes	yes	no
AMP Cloud	no	no	no	yes
Application File Name	no	no	no	yes
Application File SHA256	no	no	no	yes
Application Protocol	yes	yes	no	no
Application Protocol Category or Tag	yes	yes	yes	no
Application Risk	yes	yes	yes	no
Archive Depth	yes	yes	no	yes
Archive Name	yes	yes	no	yes

Field	File Event	Malware Events Detected by the System	Retrospective Events Detected by the System	Malware Events Detected by Secure Endpoint
Archive SHA256	yes	yes	no	yes
Business Relevance	yes	yes	yes	no
Category / File Type Category	yes	yes	no	yes
Client	yes	yes	yes	no
Client Category or Tag	yes	yes	yes	no
Count	yes	yes	yes	yes
Detection Name	no	yes	no	no
Detector	no	no	no	yes
Device	yes	yes	yes	yes
Disposition / File Disposition	yes	yes	yes	no
Domain	yes	yes	yes	yes
Event Subtype	no	no	no	yes
Event Type	no	yes	yes	yes
File Name	yes	yes	no	yes
File Path	no	no	no	yes
File Policy	yes	no	no	no
File Timestamp	no	no	no	yes
HTTP Response Code	yes	yes	no	no
IOC (Indication of Compromise)	no	yes	yes	yes
Message	yes	yes	no	yes
Receiving Continent	yes	yes	yes	no
Receiving Country	yes	yes	no	no
Receiving IP	yes	yes	no	yes
Receiving Port	yes	yes	no	no
Security Context	yes	yes	yes	yes
Sending Continent	yes	yes	yes	no
Sending Country	yes	yes	no	no

Field	File Event	Malware Events Detected by the System	Retrospective Events Detected by the System	Malware Events Detected by Secure Endpoint
Sending IP	yes	yes	no	no
Sending Port	yes	yes	no	no
SHA256 / File SHA256	yes	yes	yes	yes
Size (KB) / File Size (KB)	yes	yes	no	yes
SSL Actual Action (search only)	yes	yes	no	no
SSL Certificate Information (search only)	yes	yes	no	no
SSL Failure Reason (search only)	yes	yes	no	no
SSL Status	yes	yes	no	no
SSL Subject/Issuer Country (search only)	yes	yes	no	no
File Storage / Stored (search only)	yes	yes	no	no
Threat Name	no	yes	yes	yes
Threat Score	yes	yes	no	no
Time	yes	yes	yes	yes
Type / File Type	yes	yes	no	yes
URI / File URI	yes	yes	no	no
User	yes	yes	no	yes
Web Application	yes	yes	yes	no
Web Application Category or Tag	yes	yes	yes	no

View details about analyzed files

To see additional options, right-click a file SHA in a table on an event page. For more information, refer to Event investigation using web-based resources, on page 638.

File composition reports

A file composition report is an analysis document that

- helps identify potentially embedded malware within analyzed files,
- lists file properties and any detected viruses, and

- is pruned when the system prunes stored files.

File composition report access

If you configure local malware analysis or dynamic analysis, the system generates a file composition report after analyzing a file. The report may also list additional information specific to that file type.

To view file composition information, refer to Using a network file trajectory, on page 867.

View file details in AMP private cloud

If you have deployed an AMP private cloud, you can view additional details about analyzed files in your private cloud.

For more information, refer to the documentation for your private cloud.

Procedure

Sign in directly to your AMP private cloud console.

Threat scores and dynamic analysis summary reports

The Secure Firewall Management Center caches a file's threat score for the same amount of time as the file's disposition. If the system later detects these files, it displays the cached threat scores instead of re-querying the Secure Malware Analytics Cloud or Secure Malware Analytics Appliance. You can automatically assign a malware file disposition to any file with a threat score that exceeds the defined malware threshold threat score.

Threat scores

Table 121: Threat score ratings

Threat Score	Numeric Score	Icon
`Low`	1-25	**Low**
`Medium`	26-50	**Medium**
`High`	51-75	**High**
`Very High`	76-100	**Very High**

The Secure Firewall Management Center caches a file’s threat score for the same amount of time as the file’s disposition. If the system later detects these files, it displays the cached threat scores instead of re-querying the Secure Malware Analytics Cloud or Secure Malware Analytics Appliance. You can automatically assign a malware file disposition to any file with a threat score that exceeds the defined malware threshold threat score.

Dynamic analysis summary

If a dynamic analysis summary is available, you can click the threat score icon to view it. If multiple reports exist, this summary is based on the most recent report matching the exact threat score. If none match the exact threat score, the report with the highest threat score is displayed. If more than one report exists, you can select a threat score to view each separate report.

The summary lists each component threat comprising the threat score. Each component threat is expandable to list the AMP cloud findings, as well as any processes related to this component threat.

The process tree shows the processes that started when the Secure Malware Analytics Cloud attempted to run the file. This can help identify whether a file that contains malware is attempting to access processes and system resources beyond what is expected (for example, running a Word document opens Microsoft Word, then starts Internet Explorer, then runs the Java Runtime Environment).

Each listed process contains a process identifier you can use to verify the actual process. Child nodes in the process tree represent processes started as a result of parent processes.

From the dynamic analysis summary, you can click **View Full Report** to view the full Analysis Report, detailing the AMP cloud's full analysis, including general file information, a more in-depth review of all detected processes, a breakdown of the file analysis, and other relevant information.

View Dynamic Analysis Results in the Cisco Secure Malware Analytics Cloud

Access more detailed reporting on analyzed files than is available in the Firewall Management Center by viewing additional details about files sent for analysis from your managed devices through the Secure Malware Analytics portal.

Secure Malware Analytics offers more detailed reporting on analyzed files than is available in the Firewall Management Center. If your organization has a Secure Malware Analytics Cloud account, you can access the Secure Malware Analytics portal directly to view additional details about files sent for analysis from your managed devices.

Before you begin

- Associate your Firewall Management Center with your Secure Malware Analytics Cloud account. See *Enabling Access to Dynamic Analysis Results in the Public Cloud* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.
- License requirement: Malware
- You must be in the global domain for this task.
- You must have one of these user roles: Admin, Access Admin, Network Admin

Follow these steps to view dynamic analysis results in the Cisco Secure Malware Analytics Cloud:

Procedure

Step 1 Access the Secure Malware Analytics Cloud portal at the address provided in your Secure Malware Analytics documentation.

Step 2 Sign in with the account credentials that you used to create the association in the prerequisites to this task.

Step 3 View files submitted by your organization, or search for a particular file using its SHA.

If you have questions, refer to Secure Malware Analytics documentation.

Using captured file workflows

This task allows you to view a list of captured files in a table and manipulate the event view depending on the information relevant to your analysis. The page you see when you access captured files differs depending on the workflow, which is simply a series of pages you can use to evaluate events by moving from a broad to a more focused view. You can also create a custom workflow that displays only the information that matches your specific needs.

When a managed device captures a file detected in network traffic, it logs an event.



Note If a device captures a file containing malware, the device generates two events: a file event when it detects the file, and a malware event when it identifies malware.

If the system recaptures a file after a configuration change, such as an updated file policy, it updates existing information for that file.

For example, if you configure a file policy to capture files with a **Malware Cloud Lookup** action, the system stores the file disposition and threat score along with the file. Then, if you update your file policy, and the system recaptures the same file due to a new **Detect Files** action, the system updates the file's **Last Changed** value. However, the system does not remove the existing disposition and threat score, even though you did not perform another malware cloud lookup.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

You must be an Admin or Security Analyst user to perform this task.

Procedure

Choose **Analysis** > **Files** > **Captured Files**.

Tip
In the table view of events, several fields are hidden by default. To show a hidden field in an event view, expand the search constraints, then click the field name under **Disabled Columns**.

Related Topics

Captured file fields, on page 859
Predefined captured file workflows, on page 670
Configure event view settings, on page 234

Captured file fields

The table view of captured files, which is the final page in predefined captured file workflows, and which you can add to custom workflows, includes a column for each field in the captured files table.

When searching this table keep in mind that your search results depend on the available data in the events you are searching; depending on the available data, your search constraints may not apply. For example, if a file has never been submitted for dynamic analysis, it may not have an associated threat score.

Table 122: Captured file fields

Field	Description
Archive Inspection Status	For archive files, the status of archive inspection: • Pending indicates that the system is still inspecting the archive file and its contents. If the file passes through your system again, complete information becomes available. • Extracted indicates that the system was able to extract and inspect the archive's contents. • Failed may, in rare cases, occur if the system is unable to process an extraction. • Depth Exceeded indicates that the archive contains further nested archive files beyond the maximum allowed depth. • Encrypted indicates that the archive file's contents are encrypted and could not be inspected. • Not Inspectable indicates that the system did not extract and inspect the archive's contents. Policy rule actions, policy configuration, and corrupted files are three major reasons for this status. To view the contents of an archive file, right-click its row in the table to bring up the context menu, then choose **View Archive Contents**.
Category	The general categories of file type, for example: Office Documents, Archive, Multimedia, Executables, PDF files, Encoded, Graphics, or System Files.
Detection Name	The name of the detected malware.
Disposition	The file's malware defense disposition: • Malware indicates that local malware analysis identified malware, the AMP cloud categorized the file as malware, or that the file's threat score exceeded the malware threshold defined in the file policy. • Clean indicates that the AMP cloud categorized the file as clean, or that a user added the file to the clean list. • Unknown indicates that the system queried the AMP cloud, but the file has not been assigned a disposition; in other words, the AMP cloud has not categorized the file. • Custom Detection indicates that a user added the file to the custom detection list. • Unavailable indicates that the system could not query the AMP cloud. You may see a small percentage of events with this disposition; this is expected behavior. • N/A indicates a Detect Files or Block Files rule handled the file and the Secure Firewall Management Center did not query the AMP cloud.

Field	Description
Domain	The domain where the captured file was detected. This field is only present if you have ever configured the Firewall Management Center for multitenancy.
Dynamic Analysis Status	One or more values indicating whether the file was submitted for dynamic analysis: • Analysis Complete: file submitted for dynamic analysis that received a threat score and dynamic analysis summary report. • Capacity Handled: file stored because it could not be submitted currently. • Capacity Handled (Network Issue): file stored because it could not be submitted due to a network connectivity issue. • Capacity Handled (Rate Limit): file stored because it could not be submitted due to the maximum number of submissions reached. • Device Not Activated: file not submitted because the device is not activated on the on-premises Secure Malware Analytics Appliance. If you see this status, contact Support. • Failure (Analysis Timeout): file submitted for which the AMP cloud has yet to return a result. • Failure (Cannot Run File): file submitted that the AMP cloud could not run in the test environment. • Failure (Network Issue): file that did not get submitted due to a network connectivity failure. • Failed to send: file could not be submitted for dynamic analysis due to an internal submission failure before reaching the Secure Malware Analytics cloud. • Not Sent for Analysis: file not submitted. • Not Suspicious (Not Sent For Analysis): file pre-classified as non-malware. • Previously Analyzed: file with a cached threat score, indicating that it has been previously sent. • Rejected for Analysis: based on static analysis, the file is unlikely to pose a risk, for example because it includes no dynamic elements. • Sent for Analysis: file pre-classified as malware and queued for dynamic analysis.
Dynamic Analysis Status Changed	The last time the file's dynamic analysis status changed.
File Name	The most recently detected file name associated with the file's SHA-256 hash value.
Last Changed	The last time the information associated with this file was updated.
Last Sent	The time the file was most recently submitted to the AMP cloud for dynamic analysis.

Field	Description
Local Malware Analysis Status	One of these values indicating whether the system performed local malware analysis on a file: • Analysis Complete: the system inspected the file using local malware analysis and pre-classified the file • Analysis Failed: the system attempted to inspect the file using local malware analysis and failed • Manual Request Submitted: a user submitted a file for local malware analysis • Not Analyzed: the system did not inspect the file with local malware analysis
SHA256	The SHA-256 hash value of the file, as well as a network file trajectory icon representing the most recently detected file event and file disposition. To view the network file trajectory, click the trajectory icon.
Storage Status	Indicates whether the file is stored on a managed device: • File Stored • Not Stored (Disposition Was Pending)
Threat Score	The threat score most recently associated with this file. To view the Dynamic Analysis Summary report, click the threat score icon.
Type	The type of file; for example, HTML or MSEXE.

Stored files download

Stored files download is a security feature that

- enables downloading files from devices to a local host for long-term storage and analysis,
- archives files in a `.zip` package with disposition, file type, and SHA-256 hash value information, and
- requires confirmation by default due to potential malware risks.

Download process details

Once a device stores a file, as long as the Secure Firewall Management Center can communicate with that device and it has not deleted the file, you can download the file to a local host for long-term storage and analysis, and manually analyze the file. You can download a file from any associated file event, malware event, captured file view, or the file's trajectory.

Because malware is harmful, by default, you must confirm every file download. However, you can disable the confirmation in your User Preferences.

Because files with a disposition of Unknown may contain malware, when you download a file, the system first archives the file in a `.zip` package. The `.zip` file name contains the file disposition and file type, if available, and SHA-256 hash value. You can password-protect the `.zip` file to prevent accidental unpacking. You can edit or remove the default `.zip` file password in your User Preferences.

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

Submit files for analysis manually

You can manually submit captured files for analysis to help identify threats and gain faster insights.

When files are submitted manually, the system performs local analysis first. Then it submits them to the cloud for dynamic analysis. If local analysis is not enabled in a file policy, files are sent only for cloud-based dynamic analysis.

Manual submission allows you to analyze a range of file types (including those not eligible for automatic submission, such as .swf, .jar, and others), helping you quickly pinpoint causes of incidents.

Note The system checks the AMP cloud for updates (no more than once a day) to the list of file types eligible for dynamic analysis. It also checks for the minimum and maximum file sizes you can submit.

Before you begin

- Ensure at least one file rule is configured to store captured files.
- For more information, refer to the *Network Malware Protection and File Policies* chapter in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Depending on the situation, you can submit files for analysis in two ways.

Procedure

Step 1 To submit a single file for analysis:

a) Select one of these options:

- **Analysis** > **Files** > **File Events**
- **Analysis** > **Files** > **Malware Events**
- **Analysis** > **Files** > **Captured Files**

b) Click **Table View of <Event type or files>**.

c) Right-click a file in the table and select **Analyze File**.

Step 2 To submit multiple captured files for analysis (up to 25 at a time):

a) Select **Analysis** > **Files** > **Captured Files**

b) Select the checkbox beside each file to analyze.

c) Click **Analyze**.

Network file trajectory

A network file trajectory is a security feature that

- maps how hosts transferred files, including malware files, across your network,
- charts file transfer data, the disposition of the file, and if a file transfer was blocked or the file was quarantined, and
- enables determination of which hosts and users may have transferred malware, which hosts are at risk, and observation of file transfer trends.

File tracking capabilities

You can track the transmission of any file with an AMP cloud-assigned disposition. The system can use information related to detecting and blocking malware from both malware defense and Secure Endpoint to build a trajectory.

Recently detected malware and analyzed trajectories

Recently detected malware and analyzed trajectories are network security information displays that

- show the malware most recently detected on your network,
- display the files whose trajectory maps you have most recently viewed, and
- provide detailed information about each file including SHA-256 hash values, names, types, and current file dispositions.

Network file trajectory list page information

The Network File Trajectory List page displays the malware most recently detected on your network, as well as the files whose trajectory maps you have most recently viewed. From these lists, you can view when each file was most recently seen on the network, the file's SHA-256 hash value, name, type, current file disposition, contents (for archive files), and the number of events associated with the file.

The page also contains a search box that lets you locate files, either based on SHA-256 hash value or file name, or by the IP address of the host that transferred or received a file. After you locate a file, you can click the **File SHA256** value to view the detailed trajectory map.

Network file trajectory

A network file trajectory is a file analysis detail that

- tracks a file's movement across the network by searching for its SHA-256 value or by clicking a **File SHA 256** link,
- provides summary information about the file's identity, activity, and disposition, and

- shows related events, hosts, and users associated with the file.

Network file trajectory detailed view components

The network file trajectory details page comprises three main parts:

- Summary Information: A file's trajectory page displays summary information about the file, including file identification information; when the file was first seen and most recently seen on the network, and by what user; the number of related events and hosts associated with the file; and the file's current disposition. From this section, if the managed device stored the file, you can download it locally, submit the file for dynamic analysis, or add the file to a file list.
- Trajectory Map: A file's trajectory map visually tracks a file from the first detection on your network to the most recent. The map shows when hosts transferred or received the file, how often they transferred the file, and when the file was blocked or quarantined. Vertical lines between data points represent file transfers between hosts. Horizontal lines connecting the data points show a host's file activity over time.

  The map also shows how often file events occurred for the file and when the system assigned the file a disposition or retrospective disposition. You can select a data point in the map and highlight a path that traces back to the first instance the host transferred that file; this path also intersects with every occurrence involving the host as either sender or receiver of the file, and identifies the user involved.
- Related Events: The Events table lists event information for each data point in the map. Using the table and the map, you can pinpoint specific file events, hosts and users on the network that transferred or received this file, related events in the map, and other related events in a table constrained on selected values.

Network file trajectory summary information

The summary information appears at the top of the details page for a file that appears in the network file trajectory list.



Tip To view related file events, click a field value link. The first page in the File Events default workflow opens in a new window, displaying all file events that also contain the selected value.

Table 123: Network file trajectory summary information fields

Name	Description
Archive Contents	For inspected archive files, the number of files the archive contains.

Name	Description
Current Disposition	One of the malware defense file dispositions: • **Malware** indicates that the AMP cloud categorized the file as malware, local malware analysis identified malware, or the file's threat score exceeded the malware threshold defined in the file policy. • **Clean** indicates that the AMP cloud categorized the file as clean, or that a user added the file to the clean list. • **Unknown** indicates that the system queried the AMP cloud, but the file has not been assigned a disposition; in other words, the AMP cloud has not categorized the file. • **Custom Detection** indicates that a user added the file to the custom detection list. • **Unavailable** indicates that the system could not query the AMP cloud. You may see a small percentage of events with this disposition; this is expected behavior. • **N/A** indicates a Detect Files or Block Files rule handled the file and the Secure Firewall Management Center did not query the AMP cloud.
Detection Name	Name of the malware detected by local malware analysis.
Event Count	The number of events seen on the network associated with the file, and the number of events displayed in the map if there are more than 250 detected events.
File Category	The general categories of file type, for example, **Office Documents** or **System Files**.
File Names	The names of the file associated with the event, as seen on the network. If multiple file names are associated with a SHA-256 hash value, the most recent detected file name is listed. You can expand this to view the remaining file names by clicking **more**.
File SHA256	The SHA-256 hash value of the file. The hash is displayed by default in a condensed format. To view the full hash value, hover your pointer over it. If multiple SHA-256 hash values are associated with a file name, hover your pointer over the link to view all of the hash values.
File Size (KB)	The size of the file, in kilobytes.
File Type	The file type of the file, for example, **HTML** or **MSEXE**.
First Seen	The first time malware defense or Secure Endpoint detected the file, as well as the IP address of the host that first uploaded the file and identifying information for the user involved.
Last Seen	The most recent time malware defense or Secure Endpoint detected the file, as well as the IP address of the host that last downloaded the file and identifying information for the user involved.
Parent Application	The client application accessing the malware file when detection occurred by Secure Endpoint. These applications are **not** tied to network discovery or application control.
Seen On	The number of hosts that either sent or received the file. Because one host can upload and download a file at different times, the total number of hosts may not match the total number of senders plus the total number of receivers in the **Seen On Breakdown** field.

Name	Description
Seen On Breakdown	The number of hosts that sent the file, followed by the number of hosts that received the file.
Threat Name	Name of the threat associated with the detected malware by Secure Endpoint.
Threat Score	The file's threat score.

Network file trajectory maps and related events lists

A network file trajectory map is a visual representation that

- displays file movement across network hosts using IP addresses on the y-axis and timestamps on the x-axis,
- shows up to 250 events associated with a file SHA-256 hash in chronological order, and
- combines file events, file transfers, and retrospective events with disposition information in a single interface.

Map layout and event display

The file trajectory map's y-axis contains a list of all host IP addresses that have interacted with the file. The IP addresses are listed in descending order based on when the system first detected the file on that host. Each row contains all events associated with that IP address, whether a single file event, file transfer, or retrospective event. The x-axis contains the date and time the system detected each event. The timestamps are listed in chronological order. If multiple events occurred within a minute, all are listed within the same column. You can scroll the map horizontally and vertically to view additional events and IP addresses.

The map displays up to 250 events associated with the file SHA-256 hash. If there are more than 250 events, the map displays the first 10, then truncates extra events with an **Arrow**. The map then displays the remaining 240 events.

The first page of the File Events default workflow appears in a new window with all the extra events constrained based on the file type. If malware events generated by Secure Endpoint are not displayed, you must switch to the Malware Events table to view these.

Each data point represents an event plus the file disposition, as described in the legend below the map. For example, a Malware Block event icon combines the Malicious Disposition icon and the Block Event icon.

Malware events generated by Secure Endpoint ("endpoint-based malware events") include one icon. A retrospective event displays an icon in the column for each host on which the file is detected. File transfer events always include two icons, one file send icon and one file receive icon, connected by a vertical line. Arrows indicate the file transfer direction from sender to receiver.

To track a file's progress through the network, you can click any data point to highlight a path that includes all data points related to the selected data point. This includes data points associated with these types of events:

- any file transfers in which the associated IP address was either sender or receiver
- any malware events generated by Secure Endpoint ("endpoint-based malware events") involving the associated IP address
- if another IP address was involved, all file transfers in which that associated IP address was either sender or receiver

- if another IP address was involved, any malware events generated by Secure Endpoint ("endpoint-based malware events") involving the other IP address

All IP addresses and timestamps associated with any highlighted data point are also highlighted. The corresponding event in the Events table is also highlighted. If a path includes truncated events, the path itself is highlighted with a dotted line. Truncated events might intersect the path, but are not displayed in the map.

Using a network file trajectory

Use network file trajectory to track files as they move through your network, analyze their behavior, and identify potential security threats by examining file events and associated metadata.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Tip If your organization has deployed Secure Endpoint, that product also has a network file trajectory feature. To pivot from Firewall Management Center to Secure Endpoint, refer to Work with Event Data in the Secure Endpoint Console, on page 868. For details about the file trajectory feature in Secure Endpoint, refer to Secure Endpoint documentation.

Before you begin

If you are using malware defense, you need the Malware license.

You must be an Admin or Security Analyst user to perform this task.

Follow these steps to use a network file trajectory.

Procedure

Step 1 Choose **Analysis** > **Files** > **Network File Trajectory**.

Tip
You can also access a file's trajectory from the Context Explorer, dashboard, or event views with file information.

Step 2 Click a **File SHA 256** link in the list.

Step 3 Optionally, enter a complete SHA-256 hash value, the host IP address, or the name of a file you want to track into the search field, and press Enter.

Tip
If only one result matches, the Network File Trajectory page for that file appears.

Step 4 In the Summary Information section, you can:

- Add a file to a file list: To add a file or remove a file from the clean list or custom detection list, click **Edit** ( ).
- Download a file: To download a file, click **Download** ( ), and if prompted, confirm you want to download the file. If the file is unavailable for download, this download file is dimmed.

- Report: Click threat score to view the Dynamic Analysis Summary report.
- Submit for dynamic analysis: Click **AMP Cloud** to submit the file for dynamic analysis. If the file is unavailable for submission or you cannot connect to the AMP cloud, this AMP cloud is dimmed.
- View archive contents: To view information about an archive file's contents, click **View** ( ).
- View file composition: To view a file's composition, click **File List**. If the system has not generated a file composition report, this file list is dimmed.
- View captured files with same threat score: Click the threat score link to view all captured files with that threat score.

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

Step 5 On the trajectory map, you can:

- Locate the first instance: Click an IP address to locate the first time a file event occurred involving an IP address. This highlights a path to that data point, as well as any intervening file events and IP addresses related to the first file event. The corresponding event in the Events table is also highlighted. The map scrolls to that data point if not currently visible.
- Track: Click any data point to highlight a path that includes all data points related to the selected data point, tracking a file's progress through the network.
- View hidden events: Click the arrow to view all events not displayed in the File Summary event view.
- View matching file events: Hover your pointer over the **Matching File Event** to view summary information for the event. If you click any event summary information link, the first page of the File Events default workflow appears in a new window with all the extra events constrained based on the file type. The File Summary event view opens in a new window, displaying all file events that match on the criteria value you clicked.

Step 6 In the Events table, you can:

- Highlight: Choose a table row to highlight a data point in the map. The map scrolls to display the selected file event if not currently visible.
- Sort: Click the column headers to sort events in ascending or descending order.

Work with Event Data in the Secure Endpoint Console

This task allows you to access the Secure Endpoint console to view malware event data and use the global network file trajectory tool.

If your organization has deployed Secure Endpoint, you can view malware event data in the Secure Endpoint console, and use that application's global network file trajectory tool.

Tip For information about using Secure Endpoint and its console, refer to the online help in the console or other documentation available from https://www.cisco.com/c/en/us/support/security/fireamp-endpoints/tsd-products-support-series-home.html.

Before you begin

- The connection to Secure Endpoint must be configured (see *Integrate Secure Firewall and Secure Endpoint* in the *Cisco Secure Firewall Management Center Device Configuration Guide*) and Secure Firewall Management Center must be able to connect to the AMP cloud.
- You will need your Secure Endpoint credentials.
- You must be an Admin user to perform this task.
- If you want to pivot from a malware event in Firewall Management Center, ensure that the Secure Endpoint contextual cross-launch options are properly enabled. Refer to the topics under Event investigation using web-based resources, on page 638.

Follow these steps to access the Secure Endpoint console from the Secure Firewall Management Center:

Procedure

Step 1 Method 1:

a) Choose **Integration** > **AMP** > **AMP Management**.
b) Click the cloud name in the table.

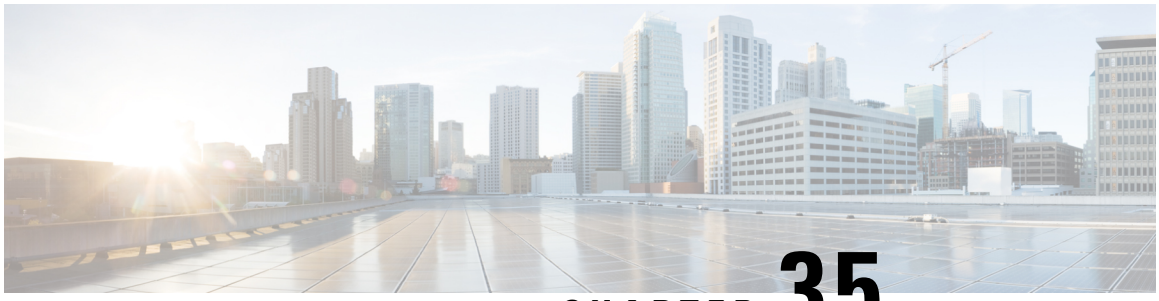
Step 2 Method 2:

a) Navigate to a malware event in a table under **Analysis** > **Files** > **Captured Files**.
b) Right-click a file SHA and choose the Secure Endpoint option.

History for file and malware events and network file trajectory

This reference tracks the chronological development and enhancement of file and malware events and network file trajectory features across different software versions, documenting when specific capabilities were introduced and their implementation details.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Improved preclassification of files for dynamic analysis.	6.7	Any	Additional assessment avoids unnecessary sending of files for dynamic analysis. The new Dynamic Analysis Status for files not sent to the cloud based on this assessment is **Rejected for Analysis**. New/modified screens: **Analysis** > **Captured Files** > **Table View of Captured Files**.
Unique identifier for connection events in syslogs.	6.4.0.4	Any	These syslog fields collectively uniquely identify a connection event and appear in syslogs for file and malware events: DeviceUUID, First Packet Time, Connection Instance ID, and Connection Counter.
Send file and malware events via syslog.	6.4	Any	Field descriptions in this chapter specify the fields included in syslog messages. For configuration information, refer to Configuration locations for syslogs for file and malware events, on page 651.



CHAPTER 35

Host Profiles

The following topics describe how to use host profiles:

- Requirements and prerequisites for host profiles, on page 871
- Host profiles, on page 872
- Basic host information in the host profile, on page 874
- Operating systems in the host profile, on page 876
- Servers in the host profile, on page 880
- Web applications in the host profile, on page 884
- Host protocols in the host profile, on page 886
- Indications of compromise in the host profile, on page 886
- VLAN tags in the host profile, on page 887
- User history in host profiles, on page 887
- Host attributes in the host profile, on page 888
- Allow list violations in the host profile, on page 892
- Malware detections in the host profile, on page 893
- Vulnerabilities in the host profile, on page 894
- Scan results in the host profile, on page 896
- History for host profiles, on page 897

Requirements and prerequisites for host profiles

This reference provides the requirements and prerequisites necessary for configuring and managing host profiles in the system.

Model support

Any.

Supported domains

Any

User roles

- Admin

- Security Analyst

Host profiles

A host profile is a system view that provides a complete view of all the information the system has gathered about a single host.

Host profile access and information

To access a host profile:

- navigate from any network map view.
- navigate from any event view that includes the IP addresses of hosts on monitored networks.

Host profiles provide basic information about detected hosts or devices, such as the host name or MAC addresses. Depending on your licenses and system configuration, host profiles can also provide you with this information:

Device and application insights

- the operating system running on a host
- the servers running on a host
- the clients and web applications running on a host
- the protocols running on a host

Security and network insights

- the indications of compromise (IOC) tags on a host
- the VLAN tags on a host
- the last twenty-four hours of user activity on your network
- the compliance allow violations associated with a host
- the most recent malware events for a host
- the vulnerabilities associated with a host
- the Nmap scan results for a host

Host attributes are also listed in the profile. You can use host attributes to classify hosts in ways that are important to your network environment. For example, you can:

- assign a host attribute that indicates the building where the host is located
- use the *host criticality* attribute to designate the business criticality of a given host and tailor correlation policies and alerts based on host criticality

From a host profile, you can view the existing host attributes applied to that host and modify the host attribute values.

If you use adaptive profile updates as part of a passive intrusion prevention deployment, you can tailor the way the system processes traffic so it best fits the type of operating system on the host and the servers and clients the host is running.

Optionally, you can perform an Nmap scan from the host profile to augment the server and operating system information in your host profile. The Nmap scanner actively probes the host to obtain information about the operating system and servers running on the host. The results of the scan are added to the list of operating system and server identities for the host.

Host profile limitations

This reference describes limitations that may affect host profile availability and the completeness of information displayed in host profiles. Understanding these limitations helps explain why some hosts may not have profiles or why information may vary between different host types.

Unavailable Hosts

A host profile may not be available for every host on your network. Possible reasons include:

- The host was deleted from the network map because it timed out.
- You have reached your host limit.
- The host resides in a network segment that is not monitored by the network discovery policy.

Unavailable Information

The information displayed in a host profile may vary according to the type of host and the information available about the host.

For example:

- If your system detects a host using a non-IP-based protocol like STP, SNAP, or IPX, the host is added to the network map as a MAC host and much less information is available than for an IP host.
- The system can add hosts to the network map from exported NetFlow records, but the available information for these hosts is limited; see Differences between NetFlow and managed device data.

(Deployments Running VRF) A Single IP Address May Represent Multiple Hosts

If a host was reported by a device running VRF, a single IP address may actually represent multiple hosts. VRF can monitor multiple networks that have overlapping IP addresses, so the same IP address can exist on different networks.

View host profiles

Host profiles provide detailed information about devices and their status on your network. Accessing this information helps identify compromised hosts and assess network security events.

Procedure

You have two choices:

- On any network map, drill down to the IP address of the host whose profile you want to view.

- On any event view, click **Host Profile** or **Compromised Host** next to the IP address of the host whose profile you want to view.

Basic host information in the host profile

Each host profile provides basic information about a detected host or other device. Descriptions of each of the basic host profile fields follow.

Domain

The domain associated with the host.

IP addresses

All IP addresses (both IPv4 and IPv6) associated with the host. The system detects IP addresses associated with hosts and, where supported, groups multiple IP addresses used by the same host. IPv6 hosts often have at least two IPv6 addresses (local-only and globally routable), and may also have IPv4 addresses. IPv4-only hosts may have multiple IPv4 addresses.

The host profile lists all detected IP addresses associated with that host. Where available, routable host IP addresses also include a flag icon and country code indicating the geolocation data associated with that address.

Note that only the first three addresses are shown by default. Click **show all** to show all addresses for a host.

Hostname

The fully qualified domain name of the host, if known.

NetBIOS name

The NetBIOS name of the host, if available. Microsoft Windows hosts, as well as Macintosh, Linux, or other platforms configured to use NetBIOS, can have a NetBIOS name. For example, Linux hosts configured as Samba servers have NetBIOS names.

Device (hops)

Either:

- the reporting device for the network where the host resides, as defined in the network discovery policy, or
- the device that processed the NetFlow data that added the host to the network map

The number of network hops between the device that detected the host and the host itself follows the device name, in parentheses. If multiple devices can see the host, the reporting device is displayed in bold.

If this field is blank, either:

- the host was added to the network map by a device that is not explicitly monitoring the network where the host resides, as defined in the network discovery policy, or
- the host was added using the host input feature and has not also been detected by the system.

MAC addresses (TTL)

The host's detected MAC address or addresses and associated NIC vendors, with the NIC's hardware vendor and current time-to-live (TTL) value in parentheses.

If multiple devices detected the host, the Firewall Management Center displays all MAC addresses and TTL values associated with the host, regardless of which device reported them.

If the MAC address is displayed in bold font, the MAC address is the actual/true/primary MAC address of the host, definitively tied to the IP address by detection through ARP and DHCP traffic.

MAC addresses that are not displayed in bold font are secondary addresses, which cannot be definitively associated with the IP address of the host. For example, since the firewall device can obtain MAC addresses only for hosts on its own network segments, if traffic originates from a network segment to which the firewall device is not directly connected, the observed MAC address (i.e. the router MAC address) will be displayed as a secondary MAC address for the host.

Host type

The type of device that the system detected: host, mobile device, jailbroken mobile device, router, bridge, NAT device, or load balancer.

The methods the system uses to distinguish network devices include:

- the analysis of Cisco Discovery Protocol (CDP) messages, which can identify network devices and their type (Cisco devices only)
- the detection of the Spanning Tree Protocol (STP), which identifies a device as a switch or bridge
- the detection of multiple hosts using the same MAC address, which identifies the MAC address as belonging to a router
- the detection of TTL value changes from the client side, or TTL values that change more frequently than a typical boot time, which identify NAT devices and load balancers
- The methods the system uses to distinguish mobile devices include:
- analysis of User-Agent strings in HTTP traffic from the mobile device's mobile browser
- monitoring of HTTP traffic of specific mobile applications

If a device is not identified as a network device or a mobile device, it is categorized as a host.

Last seen

The date and time that any of a host's IP addresses was last detected.

Current user

The user most recently logged into this host.

Note that a non-authoritative user logging into a host only registers as the current user on the host if the existing current user is not an authoritative user.

View

Links to views of connection, discovery, malware, and intrusion event data, using the default workflow for that event type and constrained to show events related to the host; where possible, these events include all IP addresses associated with the host.

Operating systems in the host profile

The system passively detects the identity of the operating system running on a host by analyzing the network and application stack in traffic generated by the host or by analyzing host data reported by the User Agent. The system also collates operating system information from other sources, such as the Nmap scanner or application data imported through the host input feature. The system considers the priority assigned to each identity source when determining which identity to use. By default, user input has the highest priority, followed by application or scanner sources, followed by the discovered identity.

Sometimes the system supplies a general operating system definition rather than a specific one because the traffic and other identity sources do not provide sufficient information for a more focused identity. The system collates information from the sources to use the most detailed definition possible.

Because the operating system affects the vulnerabilities list for the host and the event impact correlation for events targeting the host, you may want to manually supply more specific operating system information. In addition, you can indicate that fixes have been applied to the operating system, such as service packs and updates, and invalidate any vulnerabilities addressed by the fixes.

For example, if the system identifies a host's operating system as Microsoft Windows 2003, but you know that the host is actually running Microsoft Windows XP Professional with Service Pack 2, you can set the operating system identity accordingly. Setting a more specific operating system identity refines the list of vulnerabilities for the host, so your impact correlation for that host is more focused and accurate.

If the system detects operating system information for a host and that information conflicts with a current operating system identity that was supplied by an active source, an identity conflict occurs. When an identity conflict is in effect, the system uses both identities for vulnerabilities and impact correlation.

You can configure the network discovery policy to add discovery data to the network map for hosts monitored by NetFlow exporters. However, there is no operating system data available for these hosts, unless you set the use the host input feature to set the operating system identity.

If a host is running an operating system that violates a compliance allow list in an activated network discovery policy, the Firewall Management Center marks the operating system information with the allow list **Violation**. In addition, if a jailbroken mobile device violates an active allow list, the icon appears next to the operating system for the device.

You can set a custom display string for the host's operating system identity. That display string is then used in the host profile.

Note Changing the operating system information for a host may change its compliance with a compliance allow list.

In the host profile for a network device, the label for the Operating Systems section changes to Systems and an additional Hardware column appears. If a value for a hardware platform is listed under Systems, that system represents a mobile device or devices detected behind the network device. Note that mobile devices may or

may not have hardware platform information, but hardware platform information is never detected for systems that are not mobile devices.

Descriptions of the operating system information fields displayed in the host profile follow.

Hardware

The hardware platform for a mobile device.

OS Vendor/Vendor

The operating system vendor.

OS Product/Product

One of these values:

- the operating system determined most likely to be running on the host, based on the identity data collected from all sources
- `Pending` if the system has not yet identified an operating system and no other identity data is available
- `unknown` if the system cannot identify the operating system and no other identity data is available for the operating system



Note If the host's operating system is not one the system is capable of detecting, see .

OS Version/Version

The operating system version. If a host is a jailbroken mobile device, `Jailbroken` is indicated in parentheses after the version.

Source

One of these values:

- User: user_name
- Application: app_name
- Scanner: scanner_type (Nmap or other scanner)
- Firepower

The system may reconcile data from multiple sources to determine the identity of an operating system.

View operating system identities

You can view the specific operating system identities discovered or added for a host. The system uses source prioritization to determine the current identity for the host. In the list of identities, the current identity is highlighted by boldface text.

Note that the **View** is only available if multiple operating system identities exist for the host.

Procedure

Step 1 Click **View** in the **Operating System** or **Operating System Conflicts** section of the host profile.

Step 2 Review the list of operating system identities. The current identity is shown in bold text, determined by source prioritization.

For detailed information about identity sources and prioritization, refer to Operating systems in the host profile, on page 876.

Step 3 Optionally, click **Delete** ( ) next to any operating system identity to remove it.

Note
You cannot delete Cisco-detected operating system identities.

This system removes the identity from the Operating System Identity Information pop-up window and, if applicable, updates the current identity for the operating system in the host profile.

Set the current operating system identity

You can set the current operating system identity for a host using the web interface. Setting the identity through the web interface overrides all other identity sources so that identity is used for vulnerability assessment and impact correlation. However, if the system detects a conflicting operating system identity for the host after you edit the operating system, an operating system conflict occurs. Both operating systems are then considered current until you resolve the conflict.

Procedure

Step 1 Click **Edit** in the **Operating System** section of a host profile.

Step 2 You have several options:

- Choose **Current Definition** from the **OS Definition** drop-down list to confirm the current operating system identity through host input, then skip to step 6.
- Choose a variation on the current operating system identity from the **OS Definition** drop-down list, then skip to step 6.
- Choose **User-Defined** from the **OS Definition** drop-down list, then continue with step 3.

Step 3 Optionally, choose **Use Custom Display String** and modify the custom strings you want to display in the **Vendor String**, **Product String**, and **Version String** fields.

Step 4 Optionally, to change to an operating system from a different vendor, choose from the **Vendor** and **Product** drop-down lists.

Step 5 Optionally, to configure the operating system product release level, choose from the **Major**, **Minor**, **Revision**, **Build**, **Patch**, and **Extension** drop-down lists.

Step 6 Optionally, if you want to indicate that fixes for the operating system have been applied, click **Configure Fixes**.

Step 7 Choose the applicable fixes in the drop-down list, and click **Add**.

Step 8 Optionally, add the relevant patches and extensions using the **Patch** and **Extension** drop-down lists.

Step 9 Click **Finish**.

Operating system identity conflicts

An operating system identity conflict is a host profile condition that

- occurs when the system detects a new operating system identity that conflicts with a current identity from an active source,
- highlights conflicting identities in bold within the host profile, and
- can be resolved by manually setting the operating system identity through the system web interface, overriding other sources for vulnerability assessment and impact correlation.

Identity conflict resolution

You can resolve an identity conflict and set the current operating system identity for a host through the system web interface. Setting the identity through the web interface overrides all other identity sources so that identity is used for vulnerability assessment and impact correlation.

Make a conflicting operating system identity current

When a host has multiple operating system identities that conflict with each other, you need to specify which identity should be considered the current operating system for that host.

Procedure

Step 1 Navigate to the **Operating System** section of a host profile.

Step 2 Choose one of the following options:

- Click **Make Current** next to the operating system identity you want to set as the operating system for the host.
- If the identity that you *do not* want as the current identity came from an active source, delete the unwanted identity.

Resolving an operating system identity conflict

Operating system conflicts arise when there are discrepancies between the detected operating system and the configured identity in the host profile. This task allows you to select the correct operating system definition or create a custom one.

Procedure

Step 1 Click **Resolve** in the **Operating System Conflicts** section of a host profile.

Step 2 You have the following choices:

- Choose **Current Definition** from the **OS Definition** drop-down list to confirm the current operating system identity through host input, then skip to step 6.
- Choose a variation on one of the conflicting operating system identities from the **OS Definition** drop-down list, then skip to step 6.
- Choose **User-Defined** from the **OS Definition** drop-down list, then continue with step 3.

Step 3 Optionally, choose **Use Custom Display String** and enter the custom strings you want to display in the **Vendor String**, **Product String**, and **Version String** fields.

Step 4 Optionally, to change to an operating system from a different vendor, choose from the **Vendor** and **Product** drop-down lists.

Step 5 Optionally, to configure the operating system product release level, choose from the **Major**, **Minor**, **Revision**, **Build, Patch,** and **Extension** drop-down lists.

Step 6 Optionally, if you want to indicate that fixes for the operating system have been applied, click **Configure Fixes**.

Step 7 Add the fixes you have applied to the fixes list.

Step 8 Click **Finish**.

Servers in the host profile

The Servers Section of the host profile lists servers either detected on hosts on your monitored network, added from exported NetFlow records, or added through an active source like a scanner or the host input feature.

The list can include up to 100 servers per host. After that limit is reached, new server information from any source, whether active or passive, is discarded until you delete a server from the host or a server times out.

If you scan a host using Nmap, Nmap adds the results of previously undetected servers running on open TCP ports to the Servers list. If you perform an Nmap scan or import Nmap results, an expandable Scan Results section also appears in the host profile, listing the server information detected on the host by the Nmap scan. In addition, if the host is deleted from the network map, the Nmap scan results for that server for the host are discarded.

Note The system can add hosts to the network map from exported NetFlow records, but the available information for these hosts is limited; see Differences between NetFlow and managed device data.

The process for working with servers in the host profile differs depending on how you access the profile:

- If you access the host profile by drilling down through the network map, the details for that server appear with the server name highlighted in bold. If you want to view the details for any other server on the host, click **View** ( ) next to that server name.
- If you access the host profile in any other way, expand the Servers section and click **View** ( ) next to the server whose details you want to see.

Note If the host is running a server that violates a compliance allow list in an activated correlation policy, the Firewall Management Center marks the non-compliant server with the allow list **Violation**.

Descriptions of the columns in the Servers list follow.

Protocol

The name of the protocol the server uses.

Port

The port where the server runs.

Application protocol

One of:

- the name of the application protocol
- `pending` if the system cannot positively or negatively identify the application protocol for one of several reasons
- `unknown` if the system cannot identify the application protocol based on known application protocol fingerprints, or if the server was added through host input by adding a vulnerability with port information without adding a corresponding server

When you hover the mouse on an application protocol name, the tags display.

Vendor and version

The vendor and version identified by the system, Nmap, or another active source, or acquired via the host input feature. The field is blank if none of the available sources provides an identification.

Server details in the host profile

The Firewall Management Center lists up to 16 passively detected identities per server. Passive detection sources include network discovery data and NetFlow records. A server can have multiple passive identities if the system detects multiple vendors or versions of that server. For example, a load balancer between your managed device and your web server farm may cause your system to identify multiple passive identities for HTTP if your web servers are not running the same version of the server software. Note that the Firewall Management Center does not limit the number of server identities from active sources such as user input, scanners, or other applications.

The Firewall Management Center displays the current identity in bold. The system uses the current identity of a server for multiple purposes, including assigning vulnerabilities to a host, impact assessment, evaluating correlation rules written against host profile qualifications and compliance allow lists, and so on.

The server detail may also display updated sub-server information known about the selected server.

The server detail may also display the server banner, which appears below the server details when you view a server from the host profile. Server banners provide additional information about a server that may help you identify the server. The system cannot identify or detect a misidentified server when an attacker purposely

alters the server banner string. The server banner displays the first 256 bytes of the first packet detected for the server. It is collected only once, the first time the server is detected by the system. Banner content is listed in two columns, with a hexadecimal representation on the left and a corresponding ASCII representation on the right.



Note To view server banners, you must enable the **Capture Banners** check box in the network discovery policy. This option is disabled by default.

The server details section of the host profile includes the following information:

Protocol

The name of the protocol the server uses.

Port

The port where the server runs.

Hits

The number of times the server was detected by a managed device or an Nmap scanner. The number of hits is 0 for servers imported through host input, unless the system detects traffic for that server.

Last Used

The time and date the server was last detected. The last used time for host input data reflects the initial data import time unless the system detects new traffic for that server. Scanner and application data imported through the host input feature times out according to settings in the Firewall Management Center configuration, but user input through the Firewall Management Center web interface does not time out.

Application Protocol

The name of the application protocol used by the server, if known.

Vendor

The server vendor. This field does not appear if the vendor is unknown.

Version

The server version. This field does not appear if the version is unknown.

Source

One of these values:

- User: user_name
- Application: app_name
- Scanner: scanner_type (Nmap or other scanner)
- `Firepower`, `Firepower Port Match`, or `Firepower Pattern Match` for applications detected by the system
- `NetFlow` for servers added to the network map from NetFlow records

The system may reconcile data from multiple sources to determine the identity of a server.

View server details

Use this task to investigate server-specific data, such as configuration settings or operational status, for troubleshooting or validating setups.

Before you begin

Ensure you have a host profile open in the user interface.

Procedure

In a host profile, click **View** ( ) next to a server in the **Servers** section.

Editing server identities

You can manually update the identity settings for a server on a host and configure any fixes that you have applied to the host to remove the vulnerabilities addressed by the fixes. You can also delete server identities.

Deleting an identity does not delete the server, even if you delete the only identity. Deleting an identity does remove the identity from the Server Detail pop-up window and, if applicable, updates the current identity for the server in the host profile.

You cannot edit or delete server identities added by a Cisco-managed device.

Procedure

Step 1 Navigate to the **Servers** section of a host profile.

Step 2 Click **View** to open the Server Detail pop-up window.

Step 3 To delete a server identity, click **Delete** ( ) next to the server identity you want to remove.

Step 4 To modify a server identity, click **Edit** ( ) next to the server in the servers list.

You have two choices:

- Choose the current definition from the **Select Server Type** drop-down list.
- Choose the type of server from the **Select Server Type** drop-down list.

Step 5 Optionally, to only list vendors and products for that server type, choose the **Restrict by Server Type** check box.

Step 6 Optionally, to customize the name and version of the server, choose the **Use Custom Display String**, and enter a **Vendor String** and **Version String**.

Step 7 In the **Product Mappings** section, choose the operating system, product, and versions you want to use.

Example:

For example, if you want the server to map to Red Hat Linux 9, choose **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the version.

Step 8 If you want to indicate that fixes for the server have been applied, click **Configure Fixes**, and add the patches you want to apply for that server to the fixes list.

Step 9 Click **Finish**.

Resolving server identity conflicts

A server identity conflict occurs when an active source, such as an application or scanner, adds identity data for a server to a host, after which the system detects traffic for that port that indicates a conflicting server identity.

Procedure

Step 1 In a host profile, navigate to the **Servers** section.

Step 2 Click resolve next to a server.

Step 3 Choose the type of server from the **Select Server Type** drop-down list.

Step 4 Optionally, to only list vendors and products for that server type, choose the **Restrict by Server Type** check box.

Step 5 Optionally, to customize the name and version of the server, choose **Use Custom Display String**, and enter a **Vendor String** and **Version String**.

Step 6 In the **Product Mappings** section, choose the operating system, product, and versions you want to use.

Example:

For example, if you want the server to map to Red Hat Linux 9, choose **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the version.

Step 7 If you want to indicate that fixes for the server have been applied, click **Configure Fixes**, and add the patches you want to apply for that server to the fixes list.

Step 8 Click **Finish**.

Web applications in the host profile

The Web Application section of the host profile displays the clients and web applications that the system identifies as running on the hosts on your network. The system can identify client and web application information from both passive and active detection sources, although the information for hosts added from NetFlow records is limited.

Details in this section include the product and version of the detected applications on a host, any available client or web application information, and the time that the application was last detected in use.

The section lists up to 16 clients running on the host. After that limit is reached, new client information from any source, whether active or passive, is discarded until you delete a client application from the host or the system deletes the client from the host profile due to inactivity (the client times out).

Additionally, for each detected web browser, the system displays the first 100 web applications accessed. After that limit is reached, new web applications associated with that browser from any source, whether active or passive, are discarded until either:

- the web browser client application times out, or
- you delete application information associated with a web application from the host profile

If the host is running an application that violates a compliance allow list in an activated correlation policy, the Firewall Management Center marks the non-compliant application with the allow list **Violation**.

Tip To analyze the connection events associated with a particular application on the host, click **Logging** ( ) next to the application. The first page of your preferred workflow for connection events appears, showing connection events constrained by the type, product, and version of the application, as well as the IP address(es) of the host. If you do not have a preferred workflow for connection events, you must select one.

Descriptions of the application information that appears in a host profile follow.

Application protocol

Displays the application protocol used by the application (HTTP browser, DNS client, and so on).

Client

Client information derived from payload if identified by the system, captured by Nmap, or acquired via the host input feature. The field is blank if none of the available sources provides an identification.

Version

Displays the version of the client.

Web application

For web browsers, the content detected by the system in the HTTP traffic. Web application information indicates the specific type of content (for example, WMV or QuickTime) identified by the system, captured by Nmap, or acquired via the host input feature. The field is blank if none of the available sources provides an identification.

Delete web applications from the host profile

If the system detects the application again, it re-adds it to the network map and the host profile.

Procedure

Step 1 In a host profile, navigate to the **Applications** section.

Step 2 Click **Delete** ( ) next to the application you want to delete.

Host protocols in the host profile

Each host profile contains details about the protocols detected in network traffic associated with the host. The recorded protocol information includes:

- **Protocol**: The name of a protocol used by the host.
- **Layer**: The network layer where the protocol runs (**Network** or **Transport**).

If a protocol displaying in the host profile violates a compliance allow list in an activated correlation policy, the Firewall Management Center marks the non-compliant protocol with the allow list **violation**.

If the host profile lists protocols that you know are not running on the host, you can delete those protocols. Deleting a protocol from a host may bring the host into compliance with a compliance allow list.

Note If the system detects the protocol again, it automatically re-adds the protocol to the network map and the host profile.

Delete a protocol from the host profile

You can remove protocols from host profiles when they are no longer required for your network configuration or security policies.

Procedure

Step 1 Navigate to the **Protocols** section of a host profile.

Step 2 Click **Delete** ( ) next to the protocol you want to delete.

Indications of compromise in the host profile

The system correlates various types of data (intrusion events, Security Intelligence, connection events, and file or malware events) to determine whether a host on your monitored network is likely to be compromised by malicious means. Certain combinations and frequencies of event data trigger indications of compromise (IOC) tags on affected hosts. The Indications of Compromise section of the host profile displays all indication of compromise tags for a host.

To configure the system to tag indications of compromise, see *Enabling Indications of Compromise Rules* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

For more information about working with indications of compromise, see Indications of compromise data, on page 922 and the subtopics under that topic.

VLAN tags in the host profile

The VLAN Tag section of the host profile appears if the host is a member of a Virtual LAN (VLAN).

Physical network equipment often uses VLANs to create logical network segments from different network blocks. The system detects 802.1q VLAN tags and displays the information for each.

- **VLAN ID** identifies the VLAN where the host is a member. This can be any integer between zero and 4095 for 802.1q VLANs.
- **Type** identifies the encapsulated packet containing the VLAN tag, which can be either Ethernet or Token Ring.
- **Priority** identifies the priority in the VLAN tag, which can be any integer from zero to 7, where 7 is the highest priority.

If VLAN tags are nested within the packet, the system processes and the Firewall Management Center displays the innermost VLAN tag. The system collects and displays VLAN tag information only for MAC addresses that it identifies through ARP and DHCP traffic.

VLAN tag information can be useful, for example, if you have a VLAN composed entirely of printers and the system detects a Microsoft Windows 2000 operating system in that VLAN. VLAN information also helps the system generate more accurate network maps.

User history in host profiles

User history in the host profile is a network monitoring feature that

- provides a graphic representation of the last twenty-four hours of user activity,
- displays user identities with bar graphs to indicate when user logins are detected, and
- associates user logins with host IP addresses for tracking purposes.

User login types and associations

The system handles different types of user logins and their associations with host resources:

- **Authoritative logins:** Display as colored bar graphs and take precedence in host associations
- **Non-authoritative logins:** Display as gray bar graphs and do not disrupt existing authoritative user associations
- **Failed logins:** Appear in the user list when capture of failed logins is configured in the network discovery policy

Periodic login requests, such as those made to check email, appear as short regular bars. When an authoritative user login is detected for a host, that user takes over the association with the host IP address, and subsequent non-authoritative user logins do not disrupt the association.

Typical user activity patterns

A typical user logs off in the evening and may share the host resource with another user. The user history displays this pattern through bar graph representations showing login duration and frequency throughout the twenty-four hour period.

Host attributes in the host profile

A host attribute is a classification criterion that

- organizes hosts into categories relevant to network management,
- provides both predefined and user-defined attribute options, and
- enables assignment of values to hosts for more granular control.

Host attribute types

Three types of attributes are present in the system:

- *predefined host attributes*
- *compliance allow list host attributes*
- *user-defined host attributes*

After an attribute is set, whether predefined or user-defined, a value must be assigned to the associated host.



Note Host attributes can be defined at any domain level. You can assign host attributes created in current and ancestor domains.

Predefined Host Attributes

The Firewall Management Center provides two predefined host attributes:

Host Criticality

Use this attribute to designate the business criticality of a given host and to tailor correlation responses to host criticality. For example, if you consider your organization's mail servers more critical to your business than a typical user workstation, you can assign a value of High to your mail servers and other business-critical devices and Medium or Low to other hosts. You can then create a correlation policy that launches different alerts based on the criticality of an affected host.

Notes

Use this host-specific attribute to record information about the host that you want other analysts to view. For example, if you have a computer on the network that has an older, unpatched version of an operating system that you use for testing, you can use the Notes feature to indicate that the system is intentionally unpatched.

Allow list host attributes

Allow list host attributes are automatically generated attributes that

- are created with the same name as the compliance allow list,
- identify the compliance status of hosts against the allow list, and
- cannot be edited or deleted by users.

Host attribute values

Possible values for allow list host attributes are:

- Compliant: Identifies hosts that are compliant with the allow list.
- Non-Compliant: Identifies hosts that violate the allow list.
- Not Evaluated: Identifies hosts that are not valid targets of the allow list or have not been evaluated for any reason.

You cannot edit the value of an allow list host attribute or delete an allow list host attribute.

User-Defined Host Attributes

If you want to identify hosts using criteria that differs from those used in the predefined host attributes or compliance allow list host attributes, you can create user-defined host attributes. For example, you can:

- Assign physical location identifiers to hosts, such as a facility code, city, or room number.
- Assign a Responsible Party Identifier that indicates which system administrator is responsible for a given host. You can then craft correlation rules and policies to send alerts to the correct system administrator when problems related to a host are detected.
- Automatically assign values to hosts from a predefined list based on the hosts' IP addresses. This feature can be useful to assign values to new hosts when they appear on your network for the first time.

User-defined host attributes appear in the host profile page, where you can assign values on a per-host basis. You can also:

- Use the attributes in correlation policies and searches.
- View the attributes on the host attribute table view of events and generate reports based on them.

User-defined host attributes can be one of the following types:

Text

Allows you to manually assign a text string to a host.

Integer

Allows you to specify the first and last number of a range of positive integers, then manually assign one of these numbers to a host.

List

Allows you to create a list of string values, then manually assign one of the values to a host. You can also automatically assign values to hosts based on the host's IP addresses.

If you auto-assign values based on one IP address of a host with multiple IP addresses, those values will apply across all addresses associated with that host. Keep this in mind when you view the Host Attributes table.

When automatically assigning list values, consider using network objects rather than literal IP addresses. This approach can improve maintainability, particularly in a multidomain deployment where using override-enabled objects allows descendant domain administrators to tailor ancestor configurations to their local environments. In a multidomain deployment, be careful when defining auto-assigned lists at ancestor domain levels to avoid matching unintended hosts when the descendant domains use overlapping IP addresses.

URL

Allows you to manually assign a URL value to a host.

Deleting a user-defined host attribute removes it from every host profile where it is used.

Create text- or URL-based host attributes

Host attributes provide a way to classify and group hosts in your network based on custom criteria. This allows for more granular security policies and improved network management.

Procedure

Step 1 Choose **Analysis** > **Hosts heading** > **Host Attributes**.

Step 2 Click **Host Attribute Management**.

Step 3 Click **Create Attribute**.

Step 4 Enter a **Name**.

Step 5 Choose the **Type** of attribute that you want to create as described in User-Defined Host Attributes, on page 889.

Step 6 Click **Save**.

Create integer-based host attributes

When you define an integer-based host attribute, you must specify the range of numbers that the attribute accepts.

Procedure

Step 1 Choose **Analysis** > **Hosts heading** > **Host Attributes**.

Step 2 Click **Host Attribute Management**.

Step 3 Click **Create Attribute**.

Step 4 Enter a **Name**.

Step 5 Choose the **Type** of attribute that you want to create as described in User-Defined Host Attributes, on page 889.

Step 6 In the **Min** field, enter the minimum integer value that can be assigned to a host.

Step 7 In the **Max** field, enter the maximum integer value that can be assigned to a host.

Step 8 Click **Save**.

Create list-based host attributes

When you define a list-based host attribute, you must supply each of the values for the list. These values can contain alphanumeric characters, spaces, and symbols.

Before you begin

Follow these steps to create list-based host attributes:

Procedure

Step 1 Choose **Analysis** > **Hosts heading** > **Host Attributes**.

Step 2 Click **Host Attribute Management**.

Step 3 Click **Create Attribute**.

Step 4 Enter a **Name**, and choose the **Type** of attribute that you want to create as described in User-Defined Host Attributes, on page 889.

Step 5 To add a value to the list, click **Add Value**.

Step 6 In the **Name** field, enter the first value you want to add.

Optionally, to auto-assign the attribute value you just added to your hosts, click **Add Networks**.

Step 7 Choose the value you added from the **Value** drop-down list.

Step 8 In the **IP Address** and **Netmask** fields, enter the IP address and network mask (IPv4) that represent the IP address block where you want to auto-assign this value. Note that you must enter the **IP Address** and **Netmask** in CIDR notation, for example, *198.51.100.0/24*.

Repeat steps 6 through 10 to add additional values to the list and assign them automatically to new hosts that fall within an IP address block.

Step 9 Click **Save**.

Set host attribute values

You can set values for predefined and user-defined host attributes. You cannot set values for compliance allow list host attributes generated by the system.

Procedure

Step 1 Open the host profile you want to modify.

Step 2 In the **Attributes** section, click **Edit Attributes**.

Step 3 Update attribute as desired.

Step 4 Click **Save**.

Allow list violations in the host profile

A *compliance allow list* (or *allow list*) is a set of criteria that allows you to specify the operating systems, application protocols, clients, web applications, and protocols that are allowed to run on a specific subnet.

If you add an allow list to an active correlation policy, when the system detects that a host is violating the allow list, the Firewall Management Center logs an allow list event—which is a special kind of correlation event— to the database. Each of these allow list events is associated with an *allow list violation*, which indicates how and why a particular host is violating the allow list. If a host violates one or more allow lists, you can view these violations in its host profile in two ways.

First, the host profile lists all of the individual allow list violations associated with the host.

Second, in the sections associated with operating systems, applications, protocols, and servers, the Firewall Management Center marks non-compliant elements with the allow list **Violation**. For example, for an allow list that allows only Microsoft Windows hosts, the host profile displays the allow list violation icon next to the operating system information for that host.

Descriptions of the allow list violation information in the host profile follow.

Type

The type of the violation, that is, whether the violation occurred as a result of a non-compliant operating system, application, server, or protocol.

Reason

The specific reason for the violation. For example, if you have an allow list that allows only Microsoft Windows hosts, the host profile displays the current operating system running on the host (such as `Linux Linux 2.4, 2.6`)

Allow List

The name of the allow list associated with the violation.

Note You can use a host's profile to create a shared host profile for compliance allow lists.

Create Shared Allow List Host Profiles

Shared host profiles for compliance allow lists specify which operating systems, application protocols, clients, web applications, and protocols are allowed to run on target hosts across multiple allow lists. That is, if you create multiple allow lists but want to use the same host profile to evaluate hosts running a particular operating system across the allow lists, use a shared host profile.

You can use a host profile of any host with a known IP address to create a shared host profile that your compliance allow lists can use. However, note that you cannot create a shared host profile based on an individual host's host profile if the system has not yet identified the operating system of the host.

Procedure

Step 1 In a host profile, click **Generate Allow List Profile**.

Step 2 Modify and save the shared host profile according to your specific needs.

Malware detections in the host profile

The most recent malware detections section lists the most recent malware events where the host sent or received a malware file, up to 100 events. The host profile lists both network-based malware events (those generated by malware defense) and endpoint-based malware events (those generated by Secure Endpoint).

If the host is involved in a file event where the file is then retrospectively identified as malware, the original events where the file was transmitted appear in the malware detections list after the malware identification occurs. When a file identified as malware is retrospectively determined not to be malware, the malware events related to that file no longer appear in the list. For example, if a file has a disposition of `Malware` and that disposition changes to `Clean`, the event for that file is removed from the malware detections list on the host profile.

When viewing malware detections in the host profile, you can view malware events for that host by clicking the **Malware**.

Description of the columns in the Most Recent Malware Detections sections of the host profile follow.

The Most Recent Malware Detections section contains these columns:

- **Time**: The date and time the event was generated. For an event where the file was retrospectively identified as malware, note that this is the time of the original event, not the time when the malware was identified.
- **Host Role**: The host's role in the transmission of detected malware, either sender or receiver. Note that for malware events generated by Secure Endpoint ("endpoint-based malware events"), the host is always the receiver.
- **Threat Name**: The name of the detected malware.
- **File Name**: The name of the malware file.
- **File Type**: The type of file; for example, `PDF` or `MSEXE`.

Vulnerabilities in the host profile

The Vulnerabilities sections of the host profile list the vulnerabilities that affect that host. These vulnerabilities are based on the operating system, servers, and applications that the system detected on the host.

If there is an identity conflict for either the identity of the host's operating system or one of the application protocols on the host, the system lists vulnerabilities for both identities until the conflict is resolved.

Because no operating system information is available for hosts added to the network map from NetFlow data, the system cannot assign Vulnerable (impact level 1: red) impact levels for intrusion events involving those hosts. In such cases, use the host input feature to manually set the operating system identity for the hosts.

Server vendor and version information is often not included in traffic. By default, the system does not map the associated vulnerabilities for the sending and receiving hosts of such traffic. However, you can configure the system to map vulnerabilities for specific application protocols that do not have vendor or version information.

If you use the host input feature to add third-party vulnerability information for the hosts on your network, additional Vulnerabilities sections appear. For example, if you import vulnerabilities from a QualysGuard Scanner, host profiles on your include a QualysGuard Vulnerabilities section. For third-party vulnerabilities, the information in the corresponding Vulnerabilities section in the host profile is limited to the information that you provided when you imported the vulnerability data using the host input feature.

You can associate third-party vulnerabilities with operating systems and application protocols, but not clients. For information on importing third-party vulnerabilities, refer to the *Firepower System Host Input API Guide*.

Descriptions of the columns in the Vulnerabilities sections of the host profile follow.

Column	Description
Name	The name of the vulnerability.
Remote	Indicates whether the vulnerability can be remotely exploited. If this column is blank, the vulnerability definition does not include this information.
Component	The name of the operating system, application protocol, or client associated with the vulnerability.
Port	A port number, if the vulnerability is associated with an application protocol running on a specific port.

Download patches for vulnerabilities

You can download patches to mitigate the vulnerabilities discovered on the hosts on your network.

Procedure

Step 1 Access the host profile of a host for which you want to download a patch.

Step 2 Expand the **Vulnerabilities** section.

Step 3 Click the name of the vulnerability you want to patch.

Step 4 Expand the **Fixes** section to display the list of patches for the vulnerability.

Step 5 Click **Download** next to the patch you want to download.

Step 6 Download the patch and apply it to your affected systems.

Deactivate vulnerabilities for individual hosts

You can use the host vulnerability editor to deactivate vulnerabilities on a host-by-host basis. When you deactivate a vulnerability for a host, it is still used for impact correlations for that host, but the impact level is automatically reduced one level.

Procedure

Step 1 Navigate to the **Vulnerabilities** section of a host profile.

Step 2 Click **Edit Vulnerabilities**.

Step 3 Choose the vulnerability from the **Valid Vulnerabilities** list, and click the down arrow to move it to the **Invalid Vulnerabilities** list.

Tip
You can click and drag to choose multiple adjacent vulnerabilities; you can also double-click any vulnerability to move it from list to list.

Step 4 Click **Save**.

What to do next

- Optionally, activate the vulnerability for the host by moving it from the **Invalid Vulnerabilities** list to the **Valid Vulnerabilities** list.

Deactivate individual vulnerabilities

If you deactivate a vulnerability in a host profile, it deactivates it for all hosts in your network map. However, you can reactivate it at any time.

In a multidomain deployment, deactivating a vulnerability in an ancestor domain deactivates it in all descendant domains. Leaf domains can activate or deactivate a vulnerability for their devices if the vulnerability is activated in the ancestor domain.

Procedure

Step 1 Access the vulnerability detail:

- In an affected host profile, expand the **Vulnerabilities** section, and click the name of the vulnerability you want to enable or disable.

- In the predefined workflow, choose **Analysis** > **Hosts heading** > **Vulnerabilities**, and click **View** ( ) next to the vulnerability you want to enable or disable.

Step 2 Choose **Disabled** from the **Impact Qualification** drop-down list.

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

Step 3 Confirm that you want to change the **Impact Qualification** value for all hosts on the network map.

Step 4 Click **Done**.

What to do next

Optionally, activate the vulnerability by choosing **Enabled** from the **Impact Qualification** drop-down list while performing the steps above.

Scan results in the host profile

A scan result is a data record collected from network hosts that

- includes information about open and unfiltered ports,
- identifies the operating system and active servers running on those ports, and
- records additional details such as port state, server name, and detected vendor.

Scan result information

Each result indicates the following information:

- Source of the information
- Number and type of the scanned port
- Name of the server running on the port
- Additional information detected by Nmap, such as the state of the port or the vendor name for the server

If you scan for UDP ports, servers detected on those ports only appear in the Scan Results section.

Note

- You can run an Nmap scan from the host profile.
- The scan must find open ports on the host for Scan Results section to appear in the profile.

Scan a host from the host profile

You can perform a Nmap scan against a host from the host profile. After the scan completes, server and operating system information for that host are updated in the host profile. Any additional scan results are added to the Scan Results section of the host profile.

Caution Nmap-supplied server and operating system data remains static until you run another Nmap scan or override it with higher priority host input. If you plan to scan a host using Nmap, regularly schedule scans.

Before you begin

- Add an Nmap scan instance; see the *Host Identity Sources* chapter in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Procedure

Step 1 In the host profile, click **Scan Host**.

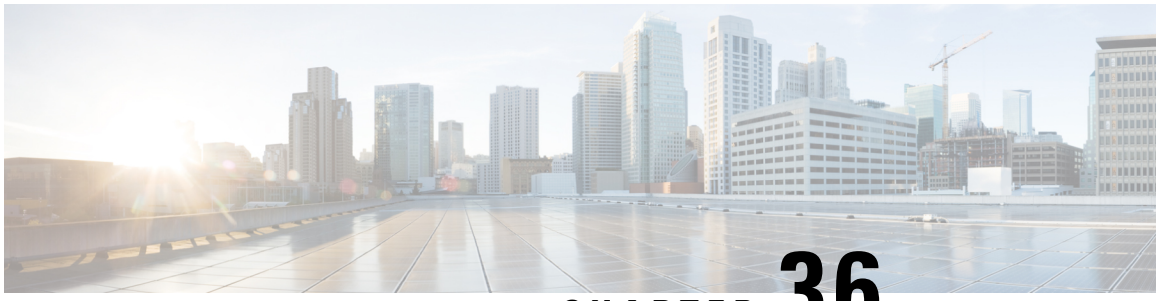
Step 2 Click **Scan** next to the scan remediation you want to use to scan the host.

The system scans the host and adds the results to the host profile.

History for host profiles

This reference provides the history of features and limitations for host profiles, specifically when virtual routing and forwarding (VRF) is used in the environment.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Limitation when using VRF	6.6	Any	If virtual routing and forwarding is used in your environment, a single IP address may represent multiple hosts because VRF may include overlapping network spaces. Supported Platforms: Firewall Management Center



CHAPTER 36

Discovery Events

The following topics describe how to work with discovery events:

- Requirements and prerequisites for discovery events, on page 899
- Discovery and identity data in discovery events, on page 899
- View discovery event statistics, on page 900
- View discovery performance graphs, on page 903
- Access discovery and identity workflows, on page 905
- History for working with discovery events, on page 956

Requirements and prerequisites for discovery events

This reference provides the requirements and prerequisites necessary for discovery events, including model support, supported domains, and required user roles.

Model support

Any.

Supported domains

Any

User roles

- Admin
- Security Analyst

Discovery and identity data in discovery events

Discovery and identity data in discovery events is network monitoring information that

- summarizes changes detected in network hosts, applications, and users,
- aggregates data based on configured network discovery and identity policies, and
- enables you to monitor threats and respond by reviewing event tables.

Discovery event data sources

The *network discovery* and *identity* policies specify the kinds of data you want to collect, the network segments you want to monitor, and the specific hardware interfaces you want to use to do it.

You can use discovery and identity event tables to identify threats associated with hosts, applications, and users on your network. The system provides a set of predefined workflows that you can use to analyze the events that your system generates. You can also create custom workflows that display only the information that matches your specific needs.

To collect and store network discovery and identity data for analysis, you must configure network discovery and identity policies. After you configure an identity policy, you must invoke it in your access control policy and deploy it to the devices you want to use to monitor traffic.

Your network discovery policy provides host, application, and non-authoritative user data. Your identity policy provides authoritative user data.

The following discovery event tables are located under the Analysis > Hosts and Analysis > Users menus.

Table 124: Discovery event table data types

Discovery Event Table	Populated With Discovery Data?	Populated With Identity Data?
Hosts	Yes	No
Host Indications of Compromise	Yes	No
Applications	Yes	No
Application Details	Yes	No
Servers	Yes	No
Host Attributes	Yes	No
Discovery Events	Yes	Yes
User Indications of Compromise	Yes	Yes
Active Sessions	Yes	Yes
User Activity	Yes	Yes
Users	Yes	Yes
Vulnerabilities	Yes	No
Third-Party Vulnerabilities	Yes	No

View discovery event statistics

The Discovery Statistics page displays a summary of the hosts, events, protocols, application protocols, and operating systems detected by the system.

The page lists statistics for the last hour and the total accumulated statistics. You can choose to view statistics for a particular device, or all devices. You can also view events that match the entries on the page by clicking the event, server, operating system, or operating system vendor listed within the summary.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Overview** > **Summary** > **Discovery Statistics**.

Step 2 From the **Select Device** list, choose the device whose statistics you want to view. Optionally, choose **All** to view statistics for all devices managed by the Firewall Management Center.

Step 3 You have the following options:

- In the Statistics Summary, view general statistics as described in Statistics summary section, on page 901.
- In the Event Breakdown, click the type of event you want to view. If no events appear, you may need to adjust the time range as described in Change the time window, on page 695.
- In the Protocol Breakdown, view the protocols currently in use by detected hosts.
- In the Application Protocol Breakdown, click the name of the application protocol you want to view.
- In the OS Breakdown, click the **OS Name** or **OS Vendor**.

Related Topics

The event breakdown section, on page 902
Protocol breakdown section, on page 903
The application protocol breakdown section, on page 903
OS breakdown section, on page 903

Statistics summary section

This reference explains how to read the Statistics Summary metrics for discovery events, hosts, and network activity.

The table summarizes each row in the Statistics Summary section:

Total events

Total number of discovery events stored on the Firewall Management Center.

Total events last hour

Total number of discovery events generated in the last hour.

Total events last day

Total number of discovery events generated in the last day.

Total application protocols

Total number of application protocols from servers running on detected hosts.

Total IP hosts

Total number of detected hosts identified by unique IP address.

Total MAC hosts

Total number of detected hosts not identified by IP address.

Note that the Total MAC Hosts statistic remains the same whether you are viewing discovery statistics for all devices or for a specific device. This is so because managed devices discover hosts based on their IP addresses. This statistic gives the total of all hosts that are identified by other means and is independent of a given managed device.

Total routers

Total number of detected nodes identified as routers.

Total bridges

Total number of detected nodes identified as bridges.

Host limit usage

Total percentage of the host limit currently in use. The host limit is defined by the model of your Firewall Management Center. Note that the host limit usage only appears if you are viewing statistics for all managed devices.



Note If the host limit is reached and a host is deleted, the host will not reappear on the network map unless you purge discovery data.

Last event received

The date and time that the most recent discovery event occurred.

Last connection received

The date and time that the most recent connection was completed.

The event breakdown section

The Event Breakdown section lists a count of each type of discovery event and host input event that occurred within the last hour, as well as a count of the total number of each event type stored in the database.

You can also use the Event Breakdown section to view details on discovery and host input events.

Related Topics

Discovery and host input events, on page 906

Protocol breakdown section

The Protocol Breakdown section lists the protocols currently in use by detected hosts. It displays each detected protocol name, its "layer" in the protocol stack, and the total number of hosts that communicate using the protocol.

The application protocol breakdown section

The Application Protocol Breakdown section lists the application protocols that are currently in use by detected hosts. It lists the protocol name, the total number of hosts running the application protocol in the past hour, and the total number of hosts that have been detected running the protocol at any point.

You can also use the Application Protocol Breakdown section to view details on servers using the detected protocols.

Related Topics

Server data, on page 926

OS breakdown section

The OS breakdown section lists the operating systems currently running on the monitored network, along with their vendors and the total number of hosts running each operating system.

You can use the OS breakdown section to view details on the detected operating systems.

A value of `unknown` for the operating system name or version means that the operating system or its version does not match any of the system's fingerprints. A value of `pending` means that the system has not yet gathered enough information to identify the operating system or its version.

Related Topics

Host data, on page 914

View discovery performance graphs

You can generate graphs that display performance statistics for managed devices with discovery events.

New data is accumulated for statistics graphs every five minutes. Therefore, if you reload a graph quickly, the data may not change until the next five-minute increment occurs.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

Edit the applicable network discovery policy to include applications, hosts, and users. (This may impact system performance.)

You must be an Admin or Maintenance user to perform this task.

Follow these steps to view discovery performance graphs:

Procedure

Step 1 Choose **Overview** > **Summary** > **Discovery Performance**.

Step 2 From the **Select Device** list, choose the Firewall Management Center or managed devices you want to include.

Step 3 From the **Select Graph(s)** list, choose the type of graph you want to create as described in Discovery performance graph types, on page 904.

Step 4 From the **Select Time Range** list, choose the time range you would like to use for the graph.

Step 5 Click **Graph** to graph the selected statistics.

Discovery performance graph types

This reference provides an overview of the available performance graphs designed to monitor system throughput and data processing metrics.

The following performance graph types are available for monitoring discovery operations:

Processed Events/Sec

Displays a graph that represents the number of events that the Data Correlator processes per second

Processed Connections/Sec

Displays a graph that represents the number of connections that the Data Correlator processes per second

Generated Events/Sec

Displays a graph that represents the number of events that the system generates per second

Mbits/Sec

Displays a graph that represents the number of megabits of traffic that are analyzed by the discovery process per second

Avg Bytes/Packet

Displays a graph that represents the average number of bytes included in each packet analyzed by the discovery process

K Packets/Sec

Displays a graph that represents the number of packets analyzed by the discovery process per second, in thousands

Access discovery and identity workflows

The Firewall Management Center provides a set of event workflows that you can use to analyze the discovery and identity data that is generated for your network. The workflows are, along with the network map, key sources of information about your network assets.

The Firewall Management Center provides predefined workflows for discovery and identity data, detected hosts and their host attributes, servers, applications, application details, vulnerabilities, user activities, and users. You can also create custom workflows.

Procedure

Step 1 To access a predefined workflow:

Data Type	Reference
Discovery and Host Input Data	Refer to View discovery and host input events, on page 912.
Host Data	Refer to View host data, on page 915.
Host Attributes Data	Refer to View host attributes, on page 920.
Host or User Indications of Compromise Data	Refer to View and work with indications of compromise data, on page 923.
Server Data	Refer to View server data, on page 927.
Application Data	Refer to View application data, on page 930.
Application Detail Data	Refer to View application detail data, on page 932.
Active Session Data	Refer to View active session data, on page 949.
User Data	Refer to View user data, on page 952.
User Activity Data	Refer to View user activity data, on page 954.
Network Map	Refer to View network maps, on page 624.

Step 2 To access a custom workflow, choose **Analysis** > **Advanced** > **Custom Workflows**.

Step 3 To access a workflow based on a custom table, choose **Analysis** > **Advanced** > **Custom Tables**.

Step 4 Perform any of the following actions, which are common to all of the pages accessed in the network discovery workflows:

- Constrain Columns — To constrain the columns that display, click **Close** ( ) in the column heading that you want to hide. In the pop-up window that appears, click **Apply**.

  Tip
  To hide or show other columns, check or clear the appropriate check boxes before you click **Apply**. To add a disabled column back to the view, click the expand arrow to expand the search constraints, then click the column name under Disabled Columns.

- Delete — To delete some or all items in the current constrained view, check the check boxes next to items you want to delete and click **Delete**, or click **Delete All**. These items remain deleted until the system's discovery function is restarted, when they may be detected again.

  Caution
  Before you delete a non-VPN session on the **Analysis** > **Users heading** > **Active Sessions** page, verify that the session is actually closed. After you delete the active session, an applicable policy will not be able to detect the session on the device, and therefore the session will not be monitored or blocked even if the policy was configured to perform those actions.

  Note
  For more information about VPN sessions on the **Analysis** > **Users heading** > **Active Sessions** page, refer to Viewing Remote Access VPN Current Users.

  Note
  You cannot delete Cisco (as opposed to third-party) vulnerabilities; you can; however, mark them reviewed.

- Drill Down — To drill down to the next page in the workflow, refer to Use drill-down pages, on page 684.
- Navigate Current Page — To navigate within the current workflow page, refer to Workflow page navigation tools, on page 681.
- Navigate within a Workflow — To navigate between pages in the current workflow, keeping the current constraints, click the appropriate page link at the top left of the workflow page.
- Navigate to Other Workflows — To navigate to other event views to examine associated events, refer to Navigate between workflows, on page 701.
- Sort Data — To sort data in a workflow, click the column title. Click the column title again to reverse the sort order.
- View Host Profile — To view the host profile for an IP address, click **Host Profile** or, for hosts with active indications of compromise (IOC) tags, the **Compromised Host** that appears next to the IP address.
- View User Profile: To view user identity information, click the user icon that appears next to the **User Identity**, or for users associated with IOCs, **Red User**.

Related Topics

Use workflows, on page 676
Purge data from the Firewall Management Center Database, on page 522

Discovery and host input events

Discovery and host input events are system-generated notifications that

- communicate the details of changes in monitored network segments through *new* events for newly discovered network features and change events for modifications in previously identified network assets,
- record network changes continuously after initial network mapping is complete by generating change events whenever the configuration of a previously discovered asset changes, and

- track manual data additions to the network map through host input events when operating system information, application protocols, clients, servers, host attributes, or vulnerability information is manually added, modified, or deleted.

Discovery event generation process

During its initial network discovery phase, the system generates new events for each host and any TCP or UDP servers discovered running on each host. Optionally, you can configure the system to use exported NetFlow records to generate these new host and server events.

In addition, the system generates new events for each network, transport, and application protocol running on each discovered host. You can disable detection of application protocols in discovery rules configured to monitor NetFlow exporters, but not in discovery rules configured to monitor managed devices. If you enable host or user discovery in non-NetFlow discovery rules, applications are automatically discovered.

After the initial network mapping is complete, the system continuously records network changes by generating change events. Change events are generated whenever the configuration of a previously discovered asset changes.

When a discovery event is generated, it is logged to the database. You can use the Firewall Management Center web interface to view, search, and delete discovery events. You can also use discovery events in correlation rules. Based on the type of discovery event generated as well as other criteria that you specify, you can build correlation rules that, when used in a correlation policy, launch remediations and syslog, SNMP, and email alert responses when network traffic meets your criteria.

You can add data to the network map using the host input feature. You can add, modify, or delete operating system information, which causes the system to stop updating that information for that host. You can also manually add, modify, or delete application protocols, clients, servers, and host attributes or modify vulnerability information. When you do this, the system generates host input events.

Discovery event types

You can configure the types of discovery events the system logs in your network discovery policy. When you view the discovery events table, the event type is listed in the **Event** column. Descriptions of the discovery event types follow.

Additional MAC Detected for Host

This event is generated when the system detects a new MAC address for a previously discovered host.

This event is often generated when the system detects hosts passing traffic through a router. While each host has a different IP address, they all appear to have the MAC address associated with the router. When the system detects the actual MAC address associated with the IP address, it displays the MAC address in bold text within the host profile and displays an “ARP/DHCP detected” message within the event description in the event view.

Client Timeout

This event is generated when the system drops a client from the database due to inactivity.

Client Update

This event is generated when the system detects a payload (that is, a specific type of content, such as audio, video, or webmail) in HTTP traffic.

DHCP: IP Address Changed

This event is generated when the system detects that a host IP address has changed due to DHCP address assignment.

DHCP: IP Address Reassigned

This event is generated when a host is reusing an IP address; that is, when a host obtains an IP address formerly used by another physical host due to DHCP IP address assignment.

Hops Change

This event is generated when the system detects a change in the number of network hops between a host and the device that detects the host. This may happen if:

- The device sees host traffic through different routers and is able to make a better determination of the host's location.
- The device detects an ARP transmission from the host, indicating that the host is on a local segment.

Host Deleted: Host Limit Reached

This event is generated when the host limit on the Firewall Management Center is exceeded and a monitored host is deleted from the network map.

Host Dropped: Host Limit Reached

This event is generated when the host limit on the Firewall Management Center is reached and a new host is dropped. Compare this with the previous event where old hosts are deleted from the network map when the host limit is reached. To drop new hosts when the host limit is reached, go to **Policies** > **Network Discovery**, then click the **Advanced** tab. In the **Advanced** tab, set **When Host Limit Reached** to **Drop hosts**.

Host IOC Set

This event is generated when an IOC (Indications of Compromise) is set for a host and generates an alert.

Host Timeout

This event is generated when a host is dropped from the network map because the host has not produced traffic within the interval defined in the network discovery policy. Note that individual host IP addresses and MAC addresses time out individually; a host does not disappear from the network map unless all of its associated addresses have timed out.

If you change the networks you want to monitor in your network discovery policy, you may want to manually delete old hosts from the network map so that they do not count against your host limit.

Host Type Changed to Network Device

This event is generated when the system detects that a detected host is actually a network device.

Identity Conflict

This event is generated when the system detects a new server or operating system identity that conflicts with a current active identity for that server or operating system.

If you want to resolve identity conflicts by rescanning the host to obtain newer active identity data, you can use Identity Conflict events to trigger an Nmap remediation.

Identity Timeout

This event is generated when server or operating system identity data from an active source times out.

If you want to refresh identity data by rescanning the host to obtain newer active identity data, you can use Identity Conflict events to trigger an Nmap remediation.

MAC Information Change

This event is generated when the system detects a change in the information associated with a specific MAC address or TTL value.

This event often occurs when the system detects hosts passing traffic through a router. While each host has a different IP address, they will all appear to have the MAC address associated with the router. When the system detects the actual MAC address associated with the IP address, it displays the MAC address in bold text within the host profile and displays an "ARP/DHCP detected" message within the event description in the event view. The TTL may change because the traffic may pass through different routers or if the system detects the actual MAC address of the host.

NETBIOS Name Change

This event is generated when the system detects a change to a host's NetBIOS name. This event will only be generated for hosts using the NetBIOS protocol.

New Client

This event is generated when the system detects a new client.



Note To collect and store client data for analysis, make sure that you enable application detection in your discovery rules in the network discovery policy.

New Host

This event is generated when the system detects a new host running on the network.

This event can also be generated when a device processes NetFlow data that involves a new host. To generate an event in this case, configure the network discovery rule that manages NetFlow data to discover hosts.

New Network Protocol

This event is generated when the system detects that a host is communicating with a new network protocol (IP, ARP, and so on).

New OS

This event is generated when the system either detects a new operating system for a host, or a change in a host's operating system.

New TCP Port

This event is generated when the system detects a new TCP server port (for example, a port used by SMTP or web services) active on a host. This event is not used to identify the application protocol or the server associated with it; that information is transmitted in the TCP Server Information Update event.

This event can also be generated when a device processes NetFlow data involving a server on your monitored networks that does not already exist in the network map. To generate an event in this case, configure the network discovery rule that manages NetFlow data to discover applications.

New Transport Protocol

This event is generated when the system detects that a host is communicating with a new transport protocol, such as TCP or UDP.

New UDP Port

This event is generated when the system detects a new UDP server port running on a host.

This event can also be generated when a device processes NetFlow data involving a server on your monitored networks that does not already exist in the network map. To generate an event in this case, configure the network discovery rule that manages NetFlow data to discover applications.

TCP Port Closed

This event is generated when the system detects that a TCP port has closed on a host.

TCP Port Timeout

This event is generated when the system has not detected activity from a TCP port within the interval defined in the system's network discovery policy.

TCP Server Information Update

This event is generated when the system detects a change in a discovered TCP server running on a host.

This event may be generated if a TCP server is upgraded.

UDP Port Closed

This event is generated when the system detects that a UDP port has closed on a host.

UDP Port Timeout

This event is generated when the system has not detected activity from a UDP port within the interval defined in the network discovery policy.

UDP Server Information Update

This event is generated when the system detects a change in a discovered UDP server running on a host.

This event may be generated if a UDP server is upgraded.

VLAN Tag Information Update

This event is generated when the system detects a change in the VLAN tag attributed to a host.

Related Topics

Host input event types, on page 911

Host input event types

When you view a table of discovery events, the event type is listed in the **Event** column.

Contrast host input events, which are generated when a user takes a specific action (such as manually adding a host), with discovery events, which are generated when the system itself detects a change in your monitored network (such as detecting traffic from a previously undetected host).

You can configure the types of host input events that the system logs by modifying your network discovery policy.

If you understand the information the different types of host input events provide, you can more effectively determine which events you want to log and alert on, and how to use these alerts in correlation policies. In addition, knowing the names of the event types can help you craft more effective event searches.

The following table describes each host input event type and what triggers it:

Add client

This event is generated when a user adds a client.

Add host

This event is generated when a user adds a host.

Add protocol

This event is generated when a user adds a protocol.

Add scan result

This event is generated when the system adds the results of an Nmap scan to a host.

Add port

This event is generated when a user adds a server port.

Delete client

This event is generated when a user deletes a client from the system.

Delete Host/Network

This event is generated when a user deletes an IP address or subnet from the system.

Delete protocol

This event is generated when a user deletes a protocol from the system.

Delete port

This event is generated when a user deletes a server port or group of server ports from the system.

Host attribute add

This event is generated when a user creates a new host attribute.

Host attribute delete

This event is generated when a user deletes a user-defined host attribute.

Host attribute delete value

This event is generated when a user deletes a value assigned to a host attribute.

Host attribute set value

This event is generated when a user sets a host attribute value for a host.

Host attribute update

This event is generated when a user changes the definition of a user-defined host attribute.

Set host criticality

This event is generated when a user sets or modifies the host criticality value for a host.

Set operating system definition

This event is generated when a user sets the operating system for a host.

Set server definition

This event is generated when a user sets the vendor and version definitions for a server.

Set vulnerability impact qualification

This event is generated when a vulnerability impact qualification is set.

When a vulnerability is disabled at a global level from being used for impact qualifications, or when a vulnerability is enabled at a global level, this event is generated.

Vulnerability set invalid

This event is generated when a user invalidates (or reviews) a vulnerability or vulnerabilities.

Vulnerability set valid

This event is generated when a user validates a vulnerability that was previously marked as invalid.

Related Topics

Discovery event types, on page 907

View discovery and host input events

Discovery events workflows allow you to view data from both discovery events and host input events. You can manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access events differs depending on the workflow you use. You can use the predefined workflow, which includes the table view of discovery events and a terminating host view page. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1 Choose **Analysis** > **Hosts heading** > **Discovery Events**.

Step 2 Configure your event view using the following options:

- Adjust the time range as described in Change the time window, on page 695.

  Note
  Events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; refer to Access discovery and identity workflows, on page 905.
- Learn more about the contents of the columns in the table; refer to Discovery event fields, on page 913.

Related Topics

Access discovery and identity workflows, on page 905

Discovery event fields

Descriptions of the fields that can be viewed and searched in the discovery events table.

Time

The time that the system generated the event.

Event

The discovery event type or host input event type.

IP Address

The IP address associated with the host involved in the event.

User

The last user to log into the host involved in the event before the event was generated. If only non-authoritative users log in after an authoritative user, the authoritative user remains the current user for the host unless another authoritative user logs in.

MAC Address

The MAC address of the NIC used by the network traffic that triggered the discovery event. This MAC address can be either the actual MAC address of the host involved in the event, or the MAC address of a network device that the traffic passed through.

MAC Vendor

The MAC hardware vendor of the NIC used by the network traffic that triggered the discovery event.

When searching this field, enter **`virtual_mac_vendor`** to match events that involve virtual hosts.

Port

The port used by the traffic that triggered the event, if applicable.

Description

The text description of the event.

Domain

The domain of the device that discovered the host. This field is only present if you have ever configured the Firewall Management Center for multitenancy.

Device

The name of the managed device that generated the event. For new host and new server events based on NetFlow data, this is the managed device that processed the data.

Related Topics

Event searches, on page 707

Host data

Host data is information that

- is generated when the system detects a host and collects information about it to build the host profile, and
- can be viewed, searched, and deleted using the Firewall Management Center web interface.

Host data management capabilities

While viewing hosts, you can create traffic profiles and compliance allow lists based on selected hosts. You can also assign host attributes, including host criticality values (which designate business criticality) to groups of hosts. You can then use these criticality values, allow lists, and traffic profiles within correlation rules and policies.

The system can add hosts to the network map from exported NetFlow records, but the available information for these hosts is limited; see Differences between NetFlow and managed device data.

View host data

You can use the Firewall Management Center to view a table of hosts that the system has detected. Then, you can manipulate the view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access hosts differs depending on the workflow you use. Both predefined workflows terminate in a host view, which contains a host profile for every host that meets your constraints. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1 Access the host data:

- If you are using the predefined workflow, choose **Analysis** > **Hosts heading** > **Hosts**.
- If you are using a custom workflow that does not include the table view of hosts, click **(switch workflow)**, then choose **Hosts**.

Step 2 You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; refer to Access discovery and identity workflows, on page 905.
- Learn more about the contents of the columns in the table; refer to Host data fields, on page 915.
- Right-click an item in the table to see options. (Not every column offers options.)
- Assign a host attribute to specific hosts; refer to Set host attributes for selected hosts, on page 922.
- Create traffic profiles for specific hosts, refer to Create a traffic profile for selected hosts, on page 919.
- Create a compliance allow list based on specific hosts, refer to Creating a compliance allow list based on selected hosts , on page 920.

Host data fields

This reference explains the searchable and viewable fields used to analyze discovered host data in the hosts table.

When the system discovers a host, it collects data about that host. That data can include the host's IP addresses, the operating system it is running, and more. You can view some of that information in the table view of hosts.

Descriptions of the fields that can be viewed and searched in the hosts table follow below.

Last Seen

The date and time any of the host's IP addresses was last detected by the system. The Last Seen value is updated at least as often as the update interval you configured in the network discovery policy, as well as when the system generates a new host event for any of the host's IP addresses.

For hosts with operating system data updated using the host input feature, the Last Seen value indicates the date and time when the data was originally added.

IP Address

The IP addresses associated with the host.

MAC Address

The host's detected MAC address of the NIC.

The MAC Address field appears in the Table View of Hosts, which you can find in the Hosts workflow. You can also add the MAC Address field to:

- custom tables that include fields from the Hosts table
- drill-down pages in custom workflows based on the Hosts table

MAC Vendor

The host's detected MAC hardware vendor of the NIC.

The MAC Vendor field appears in the Table View of Hosts, which you can find in the Hosts workflow. You can also add the MAC Vendor field to:

- custom tables that include fields from the Hosts table
- drill-down pages in custom workflows based on the Hosts table

When searching this field, enter `virtual_mac_vendor` to match events that involve virtual hosts.

Current User

The user identity (username) of the currently logged in user on the host.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Host Criticality

The user-specified criticality value assigned to the host.

NetBIOS Name

The NetBIOS name of the host. Only hosts running the NetBIOS protocol will have a NetBIOS name.

VLAN ID

VLAN ID used by the host.

Hops

The number of network hops from the device that detected the host to the host.

Host Type

The type of host. Can be any of the following: host, mobile device, jailbroken mobile device, router, bridge, NAT device, and load balancer.

The methods the system uses to distinguish network devices include:

- the analysis of Cisco Discovery Protocol (CDP) messages, which can identify network devices and their type (Cisco devices only)
- the detection of the Spanning Tree Protocol (STP), which identifies a device as a switch or bridge
- the detection of multiple hosts using the same MAC address, which identifies the MAC address as belonging to a router
- the detection of TTL value changes from the client side, or TTL values that change more frequently than a typical boot time, which identify NAT devices and load balancers

If a device is not identified as a network device, it is categorized as a host.

When searching this field, enter `!host` to search for all network devices.

Hardware

The hardware platform for a mobile device.

OS

One of the following:

- The operating system (name, vendor, and version) either detected on the host or updated using Nmap or the host input feature
- `unknown` if the operating system does not match any known fingerprint
- `pending` if the system has not yet gathered enough information to identify the operating system

If the system detects multiple identities, it displays those identities in a comma-separated list.

This field appears when you invoke the hosts event view from the Custom Analysis widget on the dashboard. It is also a field option in custom tables based on the Hosts table.

When searching this field, enter `n/a` to include hosts where the operating system has not yet been identified.

OS Conflict

This field is search only.

OS Vendor

One of the following:

- The vendor of the operating system detected on the host or updated using Nmap or the host input feature
- `unknown` if the operating system does not match any known fingerprint
- `pending` if the system has not yet gathered enough information to identify the operating system

If the system detects multiple vendors, it displays those vendors in a comma-separated list.

When searching this field, enter `n/a` to include hosts where the operating system has not yet been identified.

OS Name

One of the following:

- The operating system detected on the host or updated using Nmap or the host input feature
- `unknown` if the operating system does not match any known fingerprint
- `pending` if the system has not yet gathered enough information to identify the operating system

If the system detects multiple names, it displays those names in a comma-separated list.

When searching this field, enter `n/a` to include hosts where the operating system has not yet been identified.

OS Version

One of the following:

- The version of the operating system detected on the host or updated using Nmap or the host input feature
- `unknown` if the operating system does not match any known fingerprint
- `pending` if the system has not yet gathered enough information to identify the operating system

If the system detects multiple versions, it displays those versions in a comma-separated list.

When searching this field, enter `n/a` to include hosts where the operating system has not yet been identified.

Source Type

The type of source used to establish the host's operating system identity:

- User: user_name
- Application: app_name
- Scanner: scanner_type (Nmap or scanner added through network discovery configuration)
- `Firepower` for operating systems detected by the system

The system may reconcile data from multiple sources to determine the identity of an operating system.

Confidence

One of the following:

- the percentage of confidence that the system has in the identity of the operating system running on the host, for hosts detected by the system
- 100%, for operating systems identified by an active source, such as the host input feature or Nmap scanner
- `unknown`, for hosts for which the system cannot determine an operating system identity, and for hosts added to the network map based on NetFlow data

When searching this field, enter `n/a` to include hosts added to the network map based on NetFlow data.

Notes

The user-defined content of the Notes host attribute.

Domain

The domain associated with the host. This field is only present if you have ever configured the Firewall Management Center for multitenancy.

Device

Either the managed device that detected the traffic or the device that processed NetFlow or host input data.

If this field is blank, either of the following conditions is true:

- The host was added to the network map by a device that is not explicitly monitoring the network where the host resides, as defined in the network discovery policy.
- The host was added using the host input feature and has not also been detected by the system.

Count

The number of events that match the information that appears in each row. This field appears only after you apply a constraint that creates two or more identical rows.

Related Topics

Event searches, on page 707
Operating system identity conflicts, on page 879

Create a traffic profile for selected hosts

A traffic profile is a profile of the traffic on your network, based on connection data collected over a timespan that you specify. After you create a traffic profile, you can detect abnormal network traffic by evaluating new traffic against your profile, which presumably represents normal network traffic.

You can use the Hosts page to create a traffic profile for a group of hosts that you specify. The traffic profile will be based on connections detected where one of the hosts you specify is the initiating host. Use the sort and search features to isolate the hosts for which you want to create a profile.

Before you begin

You must be an Admin user to perform this task.

Follow these steps to create a traffic profile for selected hosts:

Procedure

Step 1 On a table view in the hosts workflow, check the check boxes next to the hosts for which you want to create a traffic profile.

Step 2 At the bottom of the page, click **Create Traffic Profile**.

Step 3 Modify and save the traffic profile according to your specific needs.

Related Topics

Introduction to traffic profiles, on page 1027

Creating a compliance allow list based on selected hosts

Compliance allow lists allow you to specify which operating systems, clients, and network, transport, or application protocols are allowed on your network.

You can use the hosts page to create a compliance allow list based on the host profiles of a group of hosts that you specify. Use the sort and search features to isolate hosts you want to use when creating an allow list.

Before you begin

You must be an Admin user to perform this task.

Procedure

Step 1 In the hosts workflow table view, check the check boxes next to the hosts for which you want to create an allow list.

Step 2 At the bottom of the page, click **Create Allow List**.

Step 3 Modify and save the allow list according to your specific needs.

Related Topics

Introduction to compliance allow lists, on page 973

Host attribute data

A host attribute is additional information about hosts on your network that

- supplements the system-collected host detection information used to build host profiles,
- can include notes, business criticality settings, or any other information you choose to provide, and
- enhances host profiles with analyst-specific data beyond automated detection capabilities.

Host attribute usage

You can use host attributes in host profile qualifications, which constrain the data you collect while building a traffic profile, and also can limit the conditions under which you want to trigger a correlation rule. You can also set attribute values in response to a correlation rule.

Related Topics

View host attributes, on page 920
Configure set attribute remediations, on page 1048

View host attributes

You can use the Firewall Management Center to view a table of hosts detected by the system, along with their host attributes. Then, you can manipulate the view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access host attributes differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of host attributes that lists all detected hosts and their attributes, and terminates in a host view page, which contains a host profile for every host that meets your constraints.

You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1 Access the host attributes data:

- If you are using the predefined workflow, choose **Analysis** > **Hosts heading** > **Host Attributes**.
- If you are using a custom workflow that does not include the table view of host attributes, click **(switch workflow)**, then choose **Attributes**.

Step 2 You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; refer to Access discovery and identity workflows, on page 905.
- Learn more about the contents of the columns in the table; refer to Host attribute data fields, on page 921.
- Assign a host attribute to specific hosts; refer to Set host attributes for selected hosts, on page 922.

Host attribute data fields

Host attributes table does not display hosts identified only by MAC addresses.

The fields available in the host attributes table are:

IP address

The IP addresses associated with a host.

Current user

The user identity (username) of the currently logged in user on the host.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Host criticality

The user-assigned importance of a host to your enterprise. You can use the host criticality in correlation rules and policies to tailor policy violations and their responses to the importance of a host involved in an event. You can assign a host criticality of low, medium, high, or none.

Notes

Information about the host that you want other analysts to view.

Any user-defined host attribute, including those for compliance allow lists

The value of the user-defined host attribute. The host attributes table contains a field for each user-defined host attribute.

Domain

The domain associated with the host. This field is only present if you have ever configured the Firewall Management Center for multitenancy.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Related Topics

Event searches, on page 707

Set host attributes for selected hosts

You can configure predefined and user-defined host attributes from a host workflow.

Procedure

Step 1 In a host workflow, check the check boxes next to the hosts to which you want to add a host attribute.

Tip
Use the sort and search features to isolate the hosts to which you want to assign particular attributes.

Step 2 At the bottom of the page, click **Set Attributes**.

Step 3 Optionally, set the host criticality for the hosts you selected. You can choose **None**, **Low**, **Medium**, or **High**.

Step 4 Optionally, add notes to the host profiles of the hosts you selected in the text box.

Step 5 Optionally, set any user-defined host attributes you have configured.

Step 6 Click **Save**.

Indications of compromise data

A indication of compromise is an event type that

- marks hosts whose IP addresses are flagged for suspicious activity,
- identifies IP addresses of compromised hosts with a **Red Compromised Host icon** in event views, and
- tags users associated with compromised hosts, displaying them with a **Red User icon** in event views.

IOC generation timing

If a file containing malware is seen again within 300 seconds of being tagged as an IOC, another IOC is not generated. If the same file is seen more than 300 seconds later, a new IOC will be generated.

To configure the system to tag events as indications of compromise, refer to *Enabling Indications of Compromise Rules* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Related Topics

Editing server identities, on page 883

View and work with indications of compromise data

You can use the Firewall Management Center to view tables showing Indications of Compromise (IOC). Manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see depends on the workflow you use. The predefined IOC workflows terminate in a profile view, which contains a host or user profile for every host or user that meets your constraints. You can also create a custom workflow that displays only the information that matches your specific needs.

Before you begin

- For your system to detect and tag indications of compromise (IOC), you must activate the IOC feature in the network discovery policy and enable at least one IOC rule. See *Enabling Indications of Compromise Rules* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.
- Users must be identified in an active Identity policy.

Follow these steps to view and work with Indications of Compromise data:

Procedure

Step 1 Determine which location in the web interface presents information that meets your needs.

You can use the following locations to view or work with Indication of Compromise data:

- Event Viewer (under the Analysis menu): Connection, Security Intelligence, intrusion, malware, and IOC discovery event views indicate whether an event triggered an IOC. Note that malware events generated by Secure Endpoint that trigger IOC rules have the event type `AMP IOC` and appear with an event subtype that specifies the compromise.
- Dashboard: In the dashboard, Threats of the Summary Dashboard displays, by default, IOC tags by host and by user. The Custom Analysis widget offers presets based on IOC data.
- Context Explorer: The Indications of Compromise section of the Context Explorer displays graphs of hosts by IOC category and IOC categories by host.
- Network Map page: The Indications of Compromise under **Analysis** > **Hosts heading** > **Network Map** groups potentially compromised hosts on your network by type of compromise and IP address.
- Network File Trajectory details page: The details pages for files listed under **Analysis** > **Files** > **Network File Trajectory** let you track indications of compromise on your network.

- Host Indications of Compromise page: The Host Indications of Compromise page under the **Analysis** > **Hosts** menu lists monitored hosts, grouped by IOC tag. Use the workflows on this page to drill down into your data.
- User Indications of Compromise page: The User Indications of Compromise page under the **Analysis** > **Users** menu lists users associated with potential IOC events, grouped by IOC tag. Use the workflows on this page to drill down into your data.
- Host Profile page: The host profile for a potentially compromised host displays all IOC tags associated with that host, and lets you resolve IOC tags and configure IOC rule states.
- User Profile page: The user profile for a user associated with a potential IOC event displays all IOC tags associated with that user, and lets you resolve IOC tags and configure IOC rule states. (The user profile is labeled "User Identity" in the Firewall Management Center web interface.)

Step 2 If applicable, do one of the following and use the rest of the steps in this procedure:

Option	Description
To research IOCs on hosts:	• If you are using the predefined workflow, choose **Analysis** > **Hosts heading** > **Indications of Compromise**. • If you are using a custom workflow that does not include the Host IOC table view, click **(switch workflow)**, then choose **Host Indications of Compromise**.
To research IOCs associated with users:	• If you are using the predefined workflow, choose **Analysis** > **Hosts heading** > **Indications of Compromise**. • If you are using a custom workflow that does not include the User IOC table view, click **(switch workflow)**, then choose **User Indications of Compromise**.

Step 3 You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; refer to Access discovery and identity workflows, on page 905.
- Learn more about the contents of the columns in the table; refer to Indications of compromise data fields, on page 925.
- On a Host Indications of Compromise page: View the host profile for a compromised host by clicking **Compromised Host** in the **IP Address** column.
- On a User Indications of Compromise page: View the user profile associated with a compromise by clicking **Red User** in the **User** column.
- Mark IOC events resolved so they no longer appear in the list. To do so, check the check boxes next to the IOC events you want to modify, then click **Mark Resolved**.
- View details of events that triggered the IOC by clicking **View** ( ) in the **First Seen** or **Last Seen** columns.
- See more options: Right-click a value in the table.

Indications of compromise data fields

These are the fields in Host or User IOC (indication of compromise) tables. Not every IOC-related table includes all fields.

IOC data fields include the following components:

- **IP Address (When viewing Host IOC data)**: The IP address associated with the host that triggered the IOC.
- **User (When viewing User IOC data)**: The username, realm, and authentication source of the user associated with the event that triggered the IOC.
- **Category**: Brief description of the type of compromise indicated, such as `Malware Executed` or `Impact 1 Attack`.
- **Event Type**: Identifier associated with a specific IOC, referring to the event that triggered it.
- **Description**: Description of the impact on the potentially compromised host, such as `This host may be under remote control` or `Malware has been executed on this host`.
- **First Seen/Last Seen**: The first/most recent date and time that events triggering the IOC occurred.
- **Domain**: The domain of the host that triggered the IOC. This field is only present if you have ever configured the Firewall Management Center for multitenancy.

Related Topics

Event searches, on page 707

Edit indication of compromise rule states for a single host or user

When enabled in a network discovery policy, indication of compromise rules apply to all hosts in the monitored network and to authoritative users that are associated with IOC events on that network. You can disable a rule for an individual host or user to avoid unhelpful IOC tags (for example, you may not want to see IOC tags for a DNS server.) If a rule is disabled in the applicable network discovery policy, it cannot be enabled for a specific host or user. Disabling a rule for a particular host does not affect tagging for the user involved in the same event, and vice-versa.

Procedure

Step 1 Navigate to the **Indications of Compromise** section of a host or user profile.

Step 2 Click **Edit Rule States**.

Step 3 In the **Enabled** column for a rule, click the slider to enable or disable it.

Step 4 Click **Save**.

View source events for indication of compromise tags

You can use the Indications of Compromise section of the host profile and the user profile to navigate quickly to the events that triggered the IOC tags. Analyzing these events can give you the information you need to determine what, and whether, action is required to address threats of compromise.

Clicking **View** ( ) next to the timestamp of an IOC tag navigates to the table view of events for the relevant event type, constrained to show only the event that triggered the IOC tag.

Only the first instance of a User IOC is displayed in the Firewall Management Center. Subsequent instances are caught by the DNS Server.

Procedure

Step 1 In a host or user profile, navigate to the **Indications of Compromise** section.

Step 2 Click **View** ( ) in the **First Seen** or **Last Seen** column for the IOC tag you want to investigate.

Resolving indication of compromise tags

After you have analyzed and addressed the threats indicated by an indication of compromise (IOC) tag, or if you determine that an IOC tag represents a false positive, you can mark an event resolved. Marking an event resolved removes it from the host profile and the user profile; when all active IOC tags on a profile are resolved, the **Compromised Host** or a user is associated with an indication of compromise **Red User icon** no longer appears. You can still view the IOC-triggering events for the resolved IOC.

If the events that triggered the IOC tag recur, the tag is set again unless you have disabled the IOC rule for the host or user.

Procedure

Step 1 In a host or user profile, navigate to the **Indications of Compromise** section.

Step 2 You have two choices:

- To mark an individual IOC tag resolved, click **Delete** ( ) to the right of the tag you want to resolve.
- To mark all IOC tags on the profile resolved, click **Mark All Resolved**.

Server data

The system collects information about all servers running on hosts on monitored network segments. This information includes:

- the name of the server
- the application and network protocols used by the server
- the vendor and version of the server
- the IP address associated with the host running a server
- the port on which the server communicates

Server detection and event generation

When the system detects a server, it generates a discovery event unless the associated host has already reached its maximum number of servers. You can use the Firewall Management Center web interface to view, search, and delete server events.

You can also base correlation rules on server events. For example, you could trigger a correlation rule when the system detects a chat server, such as ircd, running on one of your hosts.

The system can add hosts to the network map from exported NetFlow records, but the available information for these hosts is limited; see Differences between NetFlow and managed device data.

View server data

You can use the Firewall Management Center to view a table of detected servers. Then, you can manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access servers differs depending on the workflow you use. All the predefined workflows terminate in a host view, which contains a host profile for every host that meets your constraints. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1 Access the server data:

- If you are using the predefined workflow, choose **Analysis** > **Hosts heading** > **Servers**.
- If you are using a custom workflow that does not include the table view of servers, click **(switch workflow)**, then choose **Servers**.

Step 2 You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; refer to Access discovery and identity workflows, on page 905.
- Learn more about the contents of the columns in the table; refer to Server data fields, on page 927.
- Edit server identities by checking the check boxes next to the events for servers you want to edit, then clicking **Set Server Identity**.
- Right-click an item in the table to see options. (Not every column offers options.)

Server data fields

This reference describes the fields that can be viewed and searched in the servers table.

The server data table provides these fields for each server entry, allowing users to view, search, and filter server information.

Last Used

The date and time the server was last used on the network or the date and time that the server was originally updated using the host input feature. The Last Used value is updated at least as often as the update interval

you configured in the network discovery policy, as well as when the system detects a server information update.

IP Address

The IP address associated with the host running the server.

Port

The port where the server is running.

Protocol

The network or transport protocol used by the server.

Application Protocol

One of the following:

- the name of the application protocol for the server
- `pending` if the system cannot positively or negatively identify the server for one of several reasons
- `unknown` if the system cannot identify the server based on known server fingerprints or if the server was added through host input and did not include the application protocol

Category, Tags, Risk, or Business Relevance for Application Protocols

The categories, tags, risk level, and business relevance assigned to the application protocol. These filters can be used to focus on a specific set of data.

Vendor

One of the following:

- the server vendor as identified by the system, Nmap or another active source, or that you specified using the host input feature
- blank, if the system cannot identify its vendor based on known server fingerprints, or if the server was added to the network map using NetFlow data

Version

One of the following:

- the server version as identified by the system, Nmap or another active source, or that you specified using the host input feature
- blank, if the system cannot identify its version based on known server fingerprints, or if the server was added to the network map using NetFlow data

Web Application

The web application based on the payload content detected by the system in the HTTP traffic. Note that if the system detects an application protocol of `HTTP` but cannot detect a specific web application, the system supplies a generic web browsing designation.

Category, Tags, Risk, or Business Relevance for Web Applications

The categories, tags, risk level, and business relevance assigned to the web application. These filters can be used to focus on a specific set of data.

Hits

The number of times the server was accessed. For servers added using the host input feature, this value is always 0.

Source Type

One of the following values:

- User: user_name
- Application: app_name
- Scanner: scanner_type (Nmap or scanner added through network discovery configuration)
- `Firepower`, `Firepower Port Match`, or `Firepower Pattern Match` for servers detected by the system
- `NetFlow` for servers added using NetFlow data

Domain

The domain of the host running the server. This field is only present if you have ever configured the Firewall Management Center for multitenancy.

Device

Either the managed device that detected the traffic or the device that processed NetFlow or host input data.

Current User

The user identity (username) of the currently logged in user on the host.

When a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Count

The number of events that match the information that appears in each row. This field appears only after you apply a constraint that creates two or more identical rows.

Related Topics

Event searches, on page 707

Application and application details data

A monitored application is a software entity that

- is identified by the system when a host communicates with another host, and
- includes email, instant messaging, peer-to-peer, and web applications, as well as other types of applications.

Application data usage and management

For each detected application, the system logs the IP address that used the application, the product, the version, and the number of times its use was detected. You can use the web interface to view, search, and delete application events. You can also update application data on a host or hosts using the host input feature.

Application detection knowledge supports multiple system functions. You can use this knowledge to create host profile qualifications, which constrain the data you collect while building a traffic profile, and also can limit the conditions under which you want to trigger a correlation rule. You can also base correlation rules on the detection of application. For example, if you want your employees to use a specific mail client, you could trigger a correlation rule when the system detects a different mail client running on one of your hosts.

You can obtain the latest information about application detectors by carefully reading both the release notes for each system update and advisories for each VDB update.

To collect and store application data for analysis, make sure that you enable application detection in your network discovery policy.

View application data

You can use the Firewall Management Center to view a table of detected applications. Then, you can manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access applications differs depending on the workflow you use. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1 Access the application data:

- If you are using the predefined workflow, choose **Analysis** > **Hosts heading** > **Application Details**.
- If you are using a custom workflow that does not include the table view of application details, click **(switch workflow)**, then choose **Clients**.

Step 2 You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; refer to Access discovery and identity workflows, on page 905.
- Learn more about the contents of the columns in the table; refer to Application data fields, on page 931.
- Open the Application Detail View for a specific application by clicking **Application Detail View** next to a client, application protocol, or web application.

- View data in sources external to your system, by right-clicking an event value. The options you see depend on the data type and include public sources; other sources depend on the resources you have configured. For information, refer to Event investigation using web-based resources, on page 638
- Gather intelligence about an event by right-clicking an event value in the table and choosing from a Cisco or third-party intelligence source. For example, you can get details about a suspicious IP address from Cisco Talos. The options you see depend on the data type and the integrations that are configured on your system. For more information, refer to Event investigation using web-based resources, on page 638.

Application data fields

When the system detects traffic for a known client, application protocol, or web application, it logs information about the application and the host running it.

The following fields can be viewed and searched in the applications table:

Application

The name of the detected application.

IP address

The IP address associated with the host using the application.

Type

The type of application:

Application Protocols

Represents communications between hosts.

Client Applications

Represents software running on a host.

Web Applications

Represents the content or requested URL for HTTP traffic.

Category

A general classification for the application that describes its most essential function. Each application belongs to at least one category.

Tag

Additional information about the application. Applications can have any number of tags, including none.

Risk

How likely the application is to be used for purposes that might be against your organization's security policy. An application's risk can range from Very Low to Very High.

Of Application Protocol Risk, Client Risk, and Web Application Risk, the highest of the three detected, when available, in the traffic that triggered the intrusion event.

Business relevance

The likelihood that the application is used within the context of your organization's business operations, as opposed to recreationally. An application's business relevance can range from Very Low to Very High.

Of Application Protocol Business Relevance, Client Business Relevance, and Web Application Business Relevance, the lowest of the three detected, when available, in the traffic that triggered the intrusion event.

Current user

The user identity (username) of the currently logged in user on the host.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Domain

The domain of the host using the application. This field is only present if you have ever configured the Firewall Management Center for multitenancy.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Related Topics

Event searches, on page 707

View application detail data

You can use the Firewall Management Center to view a table of detected application details. Then, you can manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access application details differs depending on the workflow you use. There are two predefined workflows. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1 Access the application details data:

- If you are using the predefined workflow, choose **Analysis** > **Hosts heading** > **Application Details**.
- If you are using a custom workflow that does not include the table view of application details, click **(switch workflow)**, then select **Clients**.

Step 2 You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.

- Perform basic workflow actions; see Access discovery and identity workflows, on page 905.
- Learn more about the contents of the columns in the table; refer to Application detail data fields, on page 933.
- Open the Application Detail View for a specific application by clicking **Application Detail View** next to a client.
- View data in available sources external to your system, by right-clicking an event value. The options you see depend on the data type and include public sources; other sources depend on the resources you have configured. For information, refer to Event investigation using web-based resources, on page 638
- Gather intelligence about an event by right-clicking an event value in the table and choosing from a Cisco or third-party intelligence source. For example, you can get details about a suspicious IP address from Cisco Talos. The options you see depend on the data type and the integrations that are configured on your system. For more information, refer to Event investigation using web-based resources, on page 638.

Application detail data fields

When the system detects traffic for a known client, application protocol, or web application, it logs information about the application and the host running it.

Descriptions of the fields that can be viewed and searched in the application details table follow.

Last used

The time that the application was last used or the time that the application data was updated using the host input feature. The Last Used value is updated at least as often as the update interval you configured in the network discovery policy, as well as when the system detects an application information update.

IP address

The IP address associated with the host using the application.

Client

The name of the application. Note that if the system detected an application protocol but could not detect a specific client, `client` is appended to the application protocol name to provide a generic name.

Version

The version of the application.

Category, tags, risk, or business relevance for clients, application protocols, and web applications

The categories, tags, risk level, and business relevance assigned to the application. These filters can be used to focus on a specific set of data.

Application protocol

The application protocol used by the application. Note that if the system detected an application protocol but could not detect a specific client, `client` is appended to the application protocol name to provide a generic name.

Web application

The web application based on the payload content or URL detected by the system in the HTTP traffic. Note that if the system detects an application protocol of HTTP but cannot detect a specific web application, the system supplies a generic web browsing designation here.

Hits

The number of times the system detected the application in use. For applications added using the host input feature, this value is always `0`.

Domain

The domain of the host using the application. This field is only present if you have ever configured the Firewall Management Center for multitenancy.

Device

The device that generated the discovery event containing the application detail.

Current user

The user identity (username) of the currently logged in user on the host.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Related Topics

Event searches, on page 707

Vulnerability data

Vulnerability data is a tracking database that

- is recorded in the system's dedicated vulnerability tracking database,
- is identified in conjunction with fingerprinting capabilities for each host on your network, and
- is mapped to operating systems, servers, and clients according to their associated risks.

Vulnerability data capabilities

The system includes its own vulnerability tracking database which is used, in conjunction with the system’s fingerprinting capability, to identify the vulnerabilities associated with the hosts on your network. The operating systems, servers, and clients running on your hosts have different sets of associated vulnerabilities.

You can use the Firewall Management Center to:

- Track and review the vulnerabilities for each host.
- Deactivate vulnerabilities for a host after you patch the host or otherwise judge it immune to a vulnerability.

Vulnerabilities for vendorless and versionless servers are not mapped unless the application protocols used by the servers are mapped in the Firewall Management Center configuration. Vulnerabilities for vendorless and versionless clients cannot be mapped.

Related Topics

Mapping vulnerabilities for servers, on page 133

Vulnerability data fields

This reference explains the vulnerability data fields that appear on various pages under the Vulnerabilities menu in Firewall Management Center

Except as noted, these fields appear on all pages under **Analysis** > **Hosts heading** > **Vulnerabilities**.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

CVE ID

The identification number associated with the vulnerability in MITRE's Common Vulnerabilities and Exposures (CVE) database (https://cve.mitre.org/).

To view details about this vulnerability in the National Vulnerability Database (NVD), right-click the CVE ID and choose **View description in NVD**.

Date published

The date the vulnerability was published.

Description

A brief description of the vulnerability, from the National Vulnerability Database (NVD).

For the complete description, right-click the CVE ID and choose **View description in NVD** to view details in the National Vulnerability Database (NVD).

Impact

Refer to "Vulnerability Impact" (below.)

Impact qualification

This field is available only on the Vulnerability Details page.

Use the drop-down list to enable or disable a vulnerability. The Firewall Management Center ignores disabled vulnerabilities in its impact correlations.

The setting you specify here determines how the vulnerability is treated on a system-wide basis and is not limited to the host profile where you select the value.

Remote

Indicates whether the vulnerability is remotely exploitable (`TRUE`/`FALSE`).

Severity

The base score and Common Vulnerability Scoring System score (CVSS) from the National Vulnerability Database (NVD).

Snort ID

The identification number associated with the vulnerability in the Snort ID (SID) database. That is, if an intrusion rule can detect network traffic that exploits a particular vulnerability, that vulnerability is associated with the intrusion rule's SID.

Note that a vulnerability can be associated with more than one SID (or no SIDs at all). If a vulnerability is associated with more than one SID, the vulnerabilities table includes a row for each SID.

SVID

The vulnerability identification number that the system uses to track vulnerabilities.

To view details for this vulnerability, click **View** ( ).

Vulnerability Impact/Impact

The severity of the vulnerability on a scale of 0 to 10, with 10 being the most severe.

Related Topics

Event searches, on page 707

Vulnerability deactivation

Vulnerability deactivation is a workflow feature that

- stops the system from using a specific vulnerability during intrusion impact correlation,
- enables administrators to control vulnerability status after patching or confirming immunity, and
- applies scope and conditions based on workflow constraints, including host and domain relationships.

Vulnerability deactivation characteristics

You can deactivate a vulnerability after you patch the hosts on your network or otherwise judge them immune. Note that if the system discovers a new host that is affected by that vulnerability, the vulnerability is considered valid (and is *not* automatically deactivated) for that host.

Deactivating a vulnerability within a vulnerabilities workflow that is **not** constrained by IP addresses deactivates the vulnerability for *all* detected hosts on your network. You can deactivate vulnerabilities within the vulnerabilities workflow only on:

- the second page of the default vulnerabilities workflow, **Vulnerabilities on the Network**, which shows only the vulnerabilities that apply to the hosts on your network
- a page in a vulnerabilities workflow, custom or predefined, that you constrained based on IP address using a search.

You can deactivate a vulnerability for a single host using the network map, using the host's host profile, or by constraining the vulnerabilities workflow based on the IP addresses of the host or hosts for which you want to deactivate vulnerabilities. For hosts with multiple associated IP addresses, this function applies only to the single, selected IP address of that host.

In a multidomain deployment, deactivating a vulnerability in an ancestor domain deactivates it in all descendant domains. Leaf domains can activate or deactivate a vulnerability for their devices if the vulnerability remains active in the ancestor domain.

Related Topics

Deactivate vulnerabilities for individual hosts, on page 895
Deactivate individual vulnerabilities, on page 895
Deactivate multiple vulnerabilities, on page 938

View vulnerability data

You can use the Firewall Management Center to view a table of vulnerabilities. Then, you can manipulate the event view depending on the information you are looking for.

The page you see when you access vulnerabilities differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of vulnerabilities. The table view contains a row for each vulnerability in the database, regardless of whether any of your detected hosts exhibit the vulnerabilities. The second page of the predefined workflow contains a row for each vulnerability (that you have not deactivated) that applies to detected hosts on your network. The predefined workflow terminates in a vulnerability detail view, which contains a detailed description for every vulnerability that meets your constraints.

Tip If you want to see the vulnerabilities that apply to a single host or set of hosts, perform a search for vulnerabilities, specifying an IP address or range of IP addresses for the hosts.

You can also create a custom workflow that displays only the information that matches your specific needs.

The table of vulnerabilities is not restricted by domain in a multidomain deployment.

Procedure

Step 1 Access the table of vulnerabilities:

- If you are using the predefined vulnerabilities workflow, choose **Analysis** > **Hosts heading** > **Vulnerabilities**.
- If you are using a custom workflow that does not include the table view of vulnerabilities, click **(switch workflow)**, then choose **Vulnerabilities**.

Step 2 You have these options:

- Perform basic workflow actions; refer to Access discovery and identity workflows, on page 905.
- Deactivate vulnerabilities so they are no longer used for intrusion impact correlation for currently vulnerable hosts; refer to Deactivate multiple vulnerabilities, on page 938.

- View the details for a vulnerability by clicking **View** ( ) in the SVID column. Alternatively, constrain on the vulnerability ID and drill down to the vulnerability details page. Refer to options for viewing additional details at View vulnerability details, on page 938.
- View the full text of a vulnerability title by right-clicking the title and choosing **Show Full Text**.

View vulnerability details

Vulnerability details provide comprehensive information about security issues affecting your network, including severity ratings, affected systems, and remediation guidance.

Procedure

You can view vulnerability details in any of the following ways:

- Choose **Analysis** > **Hosts heading** > **Vulnerabilities**, and click **View** ( ) next to the SVID.
- Choose **Analysis** > **Hosts heading** > **Third-Party Vulnerabilities** and click **View** ( ) next to the SVID.
- Choose **Analysis** > **Hosts heading** > **Network Map**, and click **Vulnerabilities**.
- View the profile of a host affected by the vulnerability (**Analysis** > **Hosts heading** > **Network Map**, click **Hosts**, then drill down and click the host you are investigating), and expand the **Vulnerabilities** section of the profile.
- In any table under **Analysis** > **Hosts heading** > **Network Map**, right-click the value in the **CVE ID** column and choose **View description in NVD** to view that CVE on the NVD (National Vulnerabilities Database) web site.

Deactivate multiple vulnerabilities

Deactivating a vulnerability within a vulnerabilities workflow that is **not** constrained by IP addresses deactivates the vulnerability for *all* detected hosts on your network.

In a multidomain deployment, deactivating a vulnerability in an ancestor domain deactivates it in all descendant domains. Leaf domains can activate or deactivate a vulnerability for their devices so long as the vulnerability is activated in the ancestor domain.

Procedure

Step 1 Access the table of vulnerabilities:

- If you are using the predefined vulnerabilities workflow, choose **Analysis** > **Hosts heading** > **Vulnerabilities**.
- If you are using a custom workflow that does not include the table view of vulnerabilities, click **(switch workflow)**, then choose **Vulnerabilities**.

Step 2 Click **Vulnerabilities on the Network**.

Step 3 Check the check boxes next to vulnerabilities you want to deactivate.

Step 4 Click **Review** at the bottom of the page.

Related Topics

Deactivate vulnerabilities for individual hosts, on page 895
Deactivate individual vulnerabilities, on page 895

Third-party vulnerability data

A third-party vulnerability data source is a vulnerability database that

- augments the system's own vulnerability tracking capabilities,
- allows importing data from external applications using scripts or command line import files, and
- can be mapped to operating system and application definitions in the system database for impact correlation.

System vulnerability tracking database

The system includes its own vulnerability tracking database which is used, in conjunction with the system's fingerprinting capability, to identify the vulnerabilities associated with the hosts on your network.

You can augment the system's vulnerability data with imported network map data from third-party applications. To do so, your organization must be able to write scripts or create command line import files to import the data. For more information, refer to the *Firepower System Host Input API Guide*.

To include imported data in impact correlations, you must map third-party vulnerability information to the operating system and application definitions in the database. You cannot map third-party vulnerability information to client definitions.

View third-party vulnerability data

After you use the host input feature to import third-party vulnerability data, you can use the Firewall Management Center to view a table of third-party vulnerabilities. Then, you can manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access third-party vulnerabilities differs depending on the workflow you use. There are two predefined workflows. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1 Access the third-party vulnerabilities data:

- If you are using the predefined workflow, choose **Analysis** > **Hosts heading** > **Third-Party Vulnerabilities**.

- If you are using a custom workflow that does not include the table view of third-party vulnerabilities, click **(switch workflow)**, then choose **Vulnerabilities by Source** or **Vulnerabilities by IP Address**.

Step 2 You have these options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; refer to Access discovery and identity workflows, on page 905.
- Learn more about the contents of the columns in the table; refer to Third-party vulnerability data fields, on page 940.
- View the vulnerability details for a third-party vulnerability by clicking **View** ( ) in the SVID column. Alternatively, constrain on the vulnerability ID and drill down to the vulnerability details page.

Third-party vulnerability data fields

This reference explains the searchable and viewable fields used to analyze third-party vulnerability data.

These fields provide information about third-party vulnerabilities and their associated metadata:

Vulnerability source

The source of the third-party vulnerabilities, for example, QualysGuard or NeXpose.

Vulnerability ID

The ID number associated with the vulnerability for its source.

IP address

The IP address associated with the host affected by the vulnerability.

Port

A port number, if the vulnerability is associated with a server running on a specific port.

Bugtraq ID

The identification number associated with the vulnerability in the Bugtraq database.

CVE ID

The identification number associated with the vulnerability in MITRE's Common Vulnerabilities and Exposures (CVE) database (https://cve.mitre.org/).

SVID

The legacy vulnerability identification number that the system uses to track vulnerabilities

Click **View** ( ) to access the vulnerability details for the SVID.

Snort ID

The identification number associated with the vulnerability in the Snort ID (SID) database. That is, if an intrusion rule can detect network traffic that exploits a particular vulnerability, that vulnerability is associated with the intrusion rule's SID.

Note that a vulnerability can be associated with more than one SID (or no SIDs at all). If a vulnerability is associated with more than one SID, the vulnerabilities table includes a row for each SID.

Title

The title of the vulnerability.

Description

A brief description of the vulnerability.

Domain

The domain of the host with the vulnerability. This field is only present if you have ever configured the Firewall Management Center for multitenancy.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Related Topics

Event searches, on page 707

Active sessions, users, and user activity data

A identity source is a user data management system that

- collects active session data, user data, and user activity data from network users,
- organizes this data into distinct workflows for individual user analysis, and
- enables administrators to monitor and manage user behavior across the network.

User data workflows

Identity sources collect active session data, user data, and user activity data. The data is displayed in individual user-related workflows:

- Active Sessions: This workflow displays all current user sessions on your network. A single user running several simultaneous active sessions would occupy several rows in this table. For more information about the types of user data displayed in this workflow, refer to Active sessions data, on page 948.
- Users: This workflow displays all users seen on your network. A single user occupies a single row in this table. For more information about the types of user data displayed in this workflow, refer to User data, on page 949.

- User Activity: This workflow displays all user activity seen on your network. A single user with more than one instance of user activity would occupy several rows in this table. For more information about the types of user activity displayed in this workflow, refer to User activity data, on page 952.

For more information about the user identity sources that populate these workflows, refer to the *Cisco Secure Firewall Management Center Device Configuration Guide*.

User-related fields

User-related data is displayed in the active sessions, users, and user activity tables.

Note Active sessions for Azure AD realm users are displayed only in the **Active Sessions** new UI layout and not in the legacy UI.

Table 125: Active sessions, users, and user activity field descriptions

Field	Description	Active Sessions Table	Users Table	User Activity Table
Active Session Count	The number of active sessions associated with the user.	No	Yes	No
Authentication Type	The type of authentication: `No Authentication`, `Passive Authentication`, `Active Authentication`, `Guest Authentication`, `Failed Authentication`, or `VPN Authentication`. For more information about the supported identity sources for each Authentication Type, refer to the *Cisco Secure Firewall Management Center Device Configuration Guide* .	Yes	No	Yes
Available for Policy	A value of **Yes** means the user was retrieved from the user store (for example, Active Directory).) A value of **No** means the Firewall Management Center received a report of a login for that user but the user is not in the user store. One way this can happen is if a user in an excluded group logs in to the user store. You can exclude groups from being downloaded when you configure a realm. Users not available for policy are recorded in the Firewall Management Center but are not sent to managed devices.	No	Yes	No

Field	Description	Active Sessions Table	Users Table	User Activity Table
Count	**Note** The **Count** field is displayed only after you apply a constraint that creates two or more identical rows. Depending on the table, the number of sessions, users, or activity events that match the information that appears in a particular row.	Yes	Yes	Yes
Current IP	(See also Current IP/Domain and IP address.) The IP address associated with the host that the user is logged into. This field is blank in the Users table if there are no active sessions for a user.	Yes	No	No
Department	The user's department, as obtained by a realm. If there is no department explicitly associated with the user on your servers, the department is listed as whatever default group the server assigns. For example, on Active Directory, this is `Users (ad)`. This field is blank if: • You have not configured a realm. • The Firewall Management Center cannot correlate the user in the Firewall Management Center database with an LDAP record (for example, for users added to the database via an AIM, Oracle, or SIP login).	Yes	Yes	No
Description	More information, if available, about the session, user, or user activity.	No	No	Yes
Device	For user activity detected by traffic-based detection or an active authentication identity source, the name of the device that identified the user. For other types of user activity, the managing Firewall Management Center. **Note** If you have configured your VPN in a high-availability deployment, the device name displayed against active VPN sessions can be the primary or secondary device that identified the user session.	Yes	No	Yes

Field	Description	Active Sessions Table	Users Table	User Activity Table
Discovery Application	The application or protocol used to detect the user. • For user activity detected by traffic-based detection, one of the following: ldap, pop3, imap, oracle, sip, http, ftp, mdns, or aim. **Note** Users are not added to the database based on SMTP logins. • For all other user activity: ldap.	Yes	Yes	Yes
Current IP Domain / **Domain**	In the Active Sessions table, the multitenancy domain where the user activity was detected. In the Users table, the multitenancy domain associated with the user's realm. In the User Activity table, the multitenancy domain where the user activity was detected. This field is only present if you have ever configured the Firewall Management Center for multitenancy.	Yes	Yes	Yes
Email	The user's email address. This field is blank if: • The user was added to the database via an AIM login. • The user was added to the database via an LDAP login and there is no email address associated with the user on your LDAP servers.	Yes	Yes (as E-Mail)	No
End Port	If the user was reported by the TS Agent and their session is currently active, this field identifies the end value for the port range assigned to the user. This field is blank if the user's TS Agent session is inactive or if the user was reported by another identity source.	No	No	Yes
Endpoint Location	The IP address of the network device that used ISE to authenticate the user, as identified by ISE. If you do not configure ISE, this field is blank.	No	No	Yes
Endpoint Profile	The user's endpoint device type, as identified by Cisco ISE. If you do not configure ISE, this field is blank.	No	No	Yes
Event	The user activity event type.	No	No	Yes

Field	Description	Active Sessions Table	Users Table	User Activity Table
First Name	The user's first name, as obtained by a realm. This field is blank if: • You have not configured a realm. • The Firewall Management Center cannot correlate the user in the Firewall Management Center database with an LDAP record (for example, for users added to the database via an AIM, Oracle, or SIP login). • There is no first name associated with the user on your servers.	Yes	Yes	No
IP Address	For **User Login** user activity, the IP address or internal IP address involved in the login: • LDAP, POP3, IMAP, FTP, HTTP, MDNS, and AIM logins — the address of the user's host • SMTP and Oracle logins — the address of the server • SIP logins — the address of the session originator (See also Current IP and Current IP/Domain.) An associated IP address does not mean the user is the current user for that IP address; when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. For other types of user activity, this field is blank.	No	No	Yes
Last Name	The user's last name, as obtained by a realm. This field is blank if: • You have not configured a realm. • The Firewall Management Center cannot correlate the user in the Firewall Management Center database with an LDAP record (for example, for users added to the database via an AIM, Oracle, or SIP login). • There is no last name associated with the user on your servers.	Yes	Yes	No

Field	Description	Active Sessions Table	Users Table	User Activity Table
Last Seen	The date and time that a session was last initiated (or user data was updated) for the user.	Yes	Yes	No
Login Time	The date and time that the session was initiated for the user.	Yes	No	No
Phone Number	The user's telephone number, as obtained by a realm. This field is blank if: • You have not configured a realm. • The Firewall Management Center cannot correlate the user in the Firewall Management Center database with an LDAP record (for example, for users added to the database via an AIM, Oracle, or SIP login). • There is no telephone number associated with the user on your servers.	Yes (as Phone)	Yes	No
Realm	The identity realm associated with the user.	Yes	Yes	Yes
Security Group Tag	The Security Group Tag (SGT) attribute applied by Cisco TrustSec as the packet entered a trusted TrustSec network. If you do not configure ISE, this field is blank.	No	No	Yes
Session Duration	The duration of the user session, calculated from the **Login Time** and the current time.	Yes	No	No
Start Port	If the user was reported by the TS Agent and their session is currently active, this field identifies the start value for the port range assigned to the user. This field is blank if the user's TS Agent session is inactive or if the user was reported by another identity source.	No	No	Yes
Time	The time that the system detected the user activity.	No	No	Yes

Field	Description	Active Sessions Table	Users Table	User Activity Table
User	At minimum, this field displays the user's realm and username. For example, `Lobby\jsmith`, where Lobby is the realm and jsmith is the username. If a realm downloads additional user data from an LDAP server and the system associates it with a user, this field also displays the user's first name, last name, and type. For example, `John Smith (Lobby\jsmith, LDAP)`, where John Smith is the user's name and LDAP is the type. **Note** Because traffic-based detection can record unsuccessful AIM logins, the Firewall Management Center may store invalid AIM users (for example, if the username is misspelled by the user).	Yes	Yes	No
Username	The username associated with the user.	Yes	Yes	Yes
VPN Bytes In	For Remote Access VPN-reported user activity, the total number of bytes received from the remote peer or client by the Firewall Threat Defense. **Note** You can view the total number of bytes received once the user's VPN session is terminated. For ongoing VPN sessions, this is not a dynamic counter. For other types of user activity, this field is blank.	Yes	No	Yes
VPN Bytes Out	For Remote Access VPN-reported user activity, the total number of bytes transmitted to the remote peer or client by the Firewall Threat Defense. **Note** You can view the total number of bytes transmitted once the user's VPN session is terminated. For ongoing VPN sessions, this is not a dynamic counter. For other types of user activity, this field is blank.	No	No	Yes
VPN Client Application	For Remote Access VPN-reported user activity, the remote user's AnyConnect Remote Access VPN application. For other types of user activity, this field is blank.	Yes	No	Yes
VPN Client Country	For Remote Access VPN-reported user activity, the country name as reported by the AnyConnect Client VPN. For other types of user activity, this field is blank.	No	No	Yes

Field	Description	Active Sessions Table	Users Table	User Activity Table
VPN Client OS	For Remote Access VPN-reported user activity, the remote user's endpoint operating system as reported by the AnyConnect Client VPN. For other types of user activity, this field is blank.	Yes	No	Yes
VPN Client Public IP	For Remote Access VPN-reported user activity, the publicly routable IP address of the AnyConnect Client VPN device. For other types of user activity, this field is blank.	Yes	No	Yes
VPN Connection Duration	For Remote Access VPN-reported user activity, the total time (HH:MM:SS) that the session was active. For other types of user activity, this field is blank.	No	No	Yes
VPN Connection Profile	For Remote Access VPN-reported user activity, the name of the connection profile (tunnel group) used by the VPN session. Connection profiles are part of a Remote Access VPN Policy. For other types of user activity, this field is blank.	Yes	No	Yes
VPN Group Policy	For Remote Access VPN-reported user activity, the name of the group policy assigned to the client when the VPN session is established; either the statically-assigned group policy associated with the VPN Connection Profile, or the dynamically-assigned group policy if RADIUS is used for authentication. If assigned by the RADIUS server, this group policy overrides the static policy configured for the VPN Connection Profile. Group policies configure common attributes for groups of users in Remote Access VPN policies. For other types of user activity, this field is blank.	Yes	No	Yes
VPN Session Type	For Remote Access VPN-reported user activity, the type of session: LAN-to-LAN or Remote. For other types of user activity, this field is blank.	Yes	No	Yes

Active sessions data

The **Analysis** > **Users heading** > **Active Sessions** workflow displays select information about current user sessions.

When a user on your network runs several sessions simultaneously, the system can uniquely identify the sessions if:

- they have unique **IP Address** values.

- they have unique **Start Port** and **End Port** values, as provided by the Cisco Terminal Services (TS) Agent.
- they have unique **Current IP Domain** values.
- they were authenticated by different identity sources.
- they were associated with different identity realms.

For more information about the user and user activity data stored by the system, refer to User data, on page 949 and User activity data, on page 952.

For information about general user-related event troubleshooting and Remote Access VPN Troubleshooting, refer to the *Troubleshoot Realms and User Downloads* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

View active session data

You can view a table of active sessions, and then manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access users differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of users that lists all detected users, and terminates in a user details page. The user details page provides information on every user that meets your constraints.

Procedure

Step 1 Access the users data:

- If you are using the predefined workflow, click **Analysis** > **Users heading** > **Active Sessions**.
- If you are using a custom workflow that does not include the table view of active sessions, click **(switch workflow)**, then choose **Active Sessions**.

Step 2 Choose **Analysis** > **Users heading** > **Active Sessions**.

Step 3 You have these options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; refer to Access discovery and identity workflows, on page 905.
- Learn more about the contents of the columns in the table; refer to Active sessions data, on page 948 and User-related fields, on page 942.

User data

This reference describes how user records are created and updated in the database and what identity-related information the system stores.

When an identity source reports a user login for a user who is not already in the database, the user is added to the database, unless you have specifically restricted that login type.

The system updates the users database when one of the following occurs:

- A user on the Firewall Management Center manually deletes a non-authoritative user from the Users table.
- An identity source reports a logoff by that user.
- A realm ends the user session as specified by the realm's **User Session Timeout: Authenticated Users**, **User Session Timeout: Failed Authentication Users**, or **User Session Timeout: Guest Users** setting.



Note If you have ISE/ISE-PIC configured, you may see host data in the users table. Because host detection by ISE/ISE-PIC is not fully supported, you cannot perform user control using ISE-reported host data.

The type of user login that the system detected determines what information is stored about the new user.

Identity Source	Login Type	User Data Stored
ISE/ISE-PIC	Active Directory LDAP RADIUS RSA	• username • current IP address • Security Group Tag (SGT) — not supported with ISE-PIC • endpoint profile/device type — not supported with ISE-PIC • endpoint location/location IP — not supported with ISE-PIC • type (LDAP)
TS Agent	Active Directory	• username • current IP address • start port • end port • type (LDAP)
captive portal	Active Directory LDAP	• username • current IP address • type (LDAP)

Identity Source	Login Type	User Data Stored
traffic-based detection	LDAP AIM Oracle SIP HTTP FTP MDNS	• username • current IP address • type (AD)
	POP3 IMAP	• username • current IP address • email address • type (`pop3` or `imap`)

If you configure a realm to automatically download users, the Firewall Management Center queries the servers based on the interval you specified. It may take five to ten minutes for the Firewall Management Center database to update with user metadata after the system detects a new user login. The Firewall Management Center obtains the following information and metadata about each user:

- username
- first and last names
- email address
- department
- telephone number
- current IP address
- Security Group Tag (SGT), if available
- endpoint profile, if available
- endpoint location, if available
- start port, if available
- end port, if available

The number of users the Firewall Management Center can store in its database depends on your Firewall Management Center model. When a non-authoritative user login is detected on a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user login is detected for that host, only another authoritative user login changes the current user.

Note that traffic-based detection of AIM, Oracle, and SIP logins create duplicate user records because they are not associated with any of the user metadata that the system obtains from LDAP servers. To prevent

overuse of user count because of duplicate user records from these protocols, configure traffic-based detection to ignore those protocols.

You can search, view, and delete users from the database; you can also purge all users from the database.

For information about general user-related event troubleshooting, refer to *Cisco Secure Firewall Management Center Device Configuration Guide*.

View user data

You can view a table of users, and then manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access users differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of users that lists all detected users, and terminates in a user details page. The user details page provides information on every user that meets your constraints.

Procedure

Step 1 Access the users data:

- If you are using the predefined workflow, choose **Analysis** > **Users heading** > **Users**.
- If you are using a custom workflow that does not include the table view of users, click **(switch workflow)**, then choose **Users**.

Step 2 Choose **Analysis** > **Users heading** > **Users**.

Step 3 Choose from the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; refer to Access discovery and identity workflows, on page 905.
- Learn more about the contents of the columns in the table; refer to User-related fields, on page 942.

User activity data

A user activity data record is a security event type that

- logs the appearance and login of users on a network,
- enables correlation between user actions and other network events such as intrusions, and
- supports automated remediation and alerting through customizable rules.

The system generates events that communicate the details of user activity on your network. When the system detects user activity, the user activity data is logged to the database. You can view, search, and delete user activity; you can also purge all user activity from the database.

The system logs a user activity event when a user is seen on your network for the first time. Subsequent appearances by that user do not log new user activity events. However, if the user's IP address changes, the system logs a new user activity event.

The system also correlates user activity with other types of events. For example, intrusion events can tell you the users who were logged into the source and destination hosts at the time of the event. This correlation can tell you who was logged into the host that was targeted by an attack, or who initiated an internal attack or portscan.

You can also use user activity in correlation rules. Based on the type of user activity as well as other criteria that you specify, you can build correlation rules that, when used in a correlation policy, launch remediations and alert responses when network traffic meets your criteria.



Note If you have ISE/ISE-PIC configured, you may see host data in the users table. Because host detection by ISE/ISE-PIC is not fully supported, you cannot perform user control using ISE-reported host data.

Descriptions of the four types of user activity data follow.

New User Identity

This type of event is generated when the system detects a login by an unknown user that is not in the database.

The system logs a user activity event when a user is seen on your network for the first time. Subsequent appearances by that user do not log new user activity events. However, if the user's IP address changes, the system logs a new user activity event.

User Login

This type of event is generated when any of the following occur:

- Captive portal performs a successful or failed user authentication.
- Traffic-based detection detects a successful or failed user login.



Note SMTP logins detected by traffic-based detection are not recorded unless there is already a user with a matching email address in the database.

When a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user.

If you are using captive portal or traffic-based detection, note the following about failed user login and failed user authentication data:

- Failed logins reported by traffic-based detection (LDAP, IMAP, FTP, and POP3 traffic) are displayed in the table view of user activity, but not in the table view of users. If a known user failed to log in, the system identifies them by their username. If an unknown user failed to log in, the system uses **Failed Authentication** as their username.
- Failed authentications reported by captive portal are displayed in both the table view of user activity and the table view of users. If a known user failed to authenticate, the system identifies them by their username. If an unknown user failed to authenticate, the system identifies them by the username they entered.

Delete User Identity

This type of event is generated when you manually delete a user from the database.

User Identity Dropped: User Limit Reached

This type of event is generated when the system detects a user that is not in the database, but cannot add the user because you have reached the maximum number of users in the database as determined by your Firewall Management Center model.

After you reach the user limit, in most cases the system stops adding new users to the database. To add new users, you must either manually delete old or inactive users from the database, or purge all users from the database.

However, the system favors authoritative users. If you have reached the limit and the system detects a login for a previously undetected authoritative user, the system deletes the non-authoritative user who has remained inactive for the longest time, and replaces it with the new authoritative user.

User Indications of Compromise Events

The following user IOC changes are logged in the user activity database:

- When indications of compromise are resolved.
- When indication of compromise rules are enabled or disabled for users.

For information about general user-related event troubleshooting, refer to the *Cisco Secure Firewall Management Center Device Configuration Guide*.

View user activity data

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

You can view a table of user activity, and then manipulate the event view depending on the information you are looking for. The page you see when you access user activity differs depending on the workflow you use. You can use the predefined workflow, which includes the table view of user activity and terminates in a user details page, which contains user details for every user that meets your constraints. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1 Access the user activity data:

- If you are using the predefined workflow, choose **Analysis** > **Users heading** > **User Activity**.
- If you are using a custom workflow that does not include the table view of user activity, click **(switch workflow)**, then choose **User Activity**.

Tip
If no events appear, you may need to adjust the time range; refer to Change the time window, on page 695.

Step 2 Choose **Analysis** > **Users heading** > **User Activity**.

Step 3 You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; refer to Access discovery and identity workflows, on page 905.
- Learn more about the contents of the columns in the table; refer to User-related fields, on page 942.

User profile and host history

A user profile is a web interface display that

- provides detailed information about a specific user through the User pop-up window (titled "User Identity" in the web interface)
- shows the same user data as you would see in the table view of users, and
- includes host history that provides a graphic representation of the last twenty-four hours of the user's activity.

User profile access points

You can learn more about a specific user by viewing the User pop-up window. The page that appears, called the "User Profile" in this document, is titled "User Identity" in the web interface.

You can display the window from:

- any event view that associates user data with other kinds of events
- the table view of active sessions
- the table view of users

User information also appears in the terminating page for users workflows.

The user data you see is the same as you would see in the table view of users.

Indications of compromise section

For information about this section, refer to:

- *Indications of Compromise* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- Indications of compromise data fields, on page 925
- Edit indication of compromise rule states for a single host or user, on page 925
- Resolving indication of compromise tags, on page 926
- View source events for indication of compromise tags, on page 925

Host history section

The host history provides a graphic representation of the last twenty-four hours of the user’s activity. A list of IP addresses of the hosts that the user logged into and logged off of approximates login and logout times with bar graphs. A typical user might log on to and off of multiple hosts in the course of a day. For example,

periodic automated logins to a mail server would display as multiple short sessions, while longer logins (such as during working hours) display longer sessions.

If you use traffic-based detection or captive portal to capture failed logins, the host history also includes hosts where the user failed to log in.

The data used to generate the host history is stored in the user history database, which by default stores 10 million user login events. If you do not see any data in the host history for a particular user, either that user is inactive, or you may need to increase the database limit.

Related Topics

User data fields

View user details and host history

User details and host history provide valuable insights for security analysis and incident investigation.

Procedure

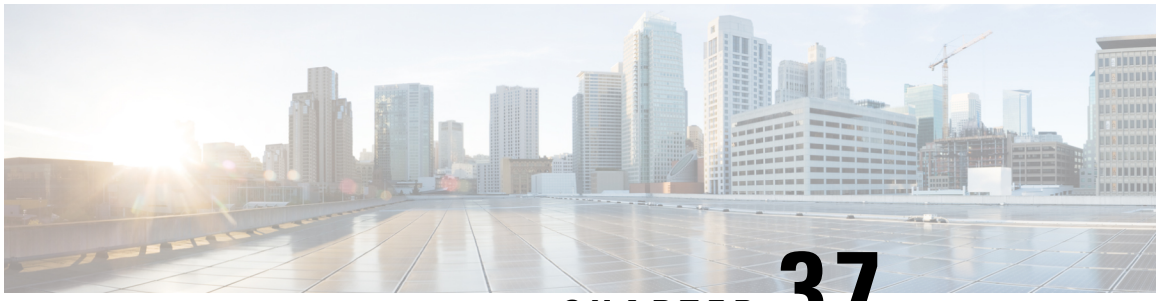
You have two options:

- In any event view that lists users, click user that appears next to a user identity **User icon**, or, for users associated with an indication of compromise, **Red User icon**.
- In any users workflow, click the Users terminating page.

History for working with discovery events

This reference documents the historical changes and feature updates related to discovery events, specifically focusing on vulnerability data processing changes across different system versions.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Vulnerabilities pages changes	6.7	Any	Bugtraq and its vulnerability data are no longer available. The following changes have been made: • Most vulnerability data now comes from the National Vulnerability Database (NVD). • Obsolete and redundant fields have been removed. • A new CVE ID column has been added to table views, and a new Severity field has been added to tables and details pages. • You can now right-click the CVE ID in tables to view details about that vulnerability in the NVD. • The Vulnerability Impact column in tables has been renamed to Impact. (No change to the field name in Detail views.) • When viewing vulnerabilities in host profiles under Analysis > Hosts > Network Map > Hosts, details for vulnerabilities (excluding third-party vulnerabilities) use the new set of fields. • The Bugtraq option has been removed from the Vulnerabilities options on the Analysis > Hosts > Network Map > Vulnerabilities page. Modified screens: • All pages under Analysis > Hosts > Vulnerabilities • Hosts and Vulnerabilities tabs on Analysis > Hosts > Network Map pages Supported Platforms: Firewall Management Center



CHAPTER 37

Correlation and Compliance Events

The following topics describe how to view correlation and compliance events.

- View correlation events, on page 959
- Analyze compliance events using compliance allow list workflows, on page 962
- Remediation status events, on page 967

View correlation events

When a correlation rule within an active correlation policy triggers, the system generates a correlation event and logs it for further analysis. You can customize the event view to focus on relevant security incidents and supporting data.

Note When a compliance allow list within an active correlation policy triggers, the system generates an allow list event.

You can view a table of correlation events and then manipulate the event view based on the information you need.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access correlation events differs based on the workflow you select. The predefined workflow includes a table view of correlation events. You can also create a custom workflow to display only the information that matches your specific requirements.

Before you begin

You must be an Admin or Security Analyst user to perform this task.

Procedure

Step 1 Choose **Analysis** > **Correlation** > **Correlation Events** .

Optionally, to use a different workflow, including a custom workflow, click **(switch workflow)** by the workflow title.

Note
If you are using a custom workflow that does not include the table view of correlation events, click **(switch workflow)**, then choose **Correlation Events**.

Step 2 Optionally, adjust the time range as described in Change the time window, on page 695.

Step 3 Perform any of the following actions as needed:

- To learn more about the columns that appear, refer to Correlation event fields, on page 960.
- To view the host profile for an IP address, click the host profile link next to the IP address.
- To view user identity information, click the user icon that appears next to the **User Identity**, or for users associated with IOCs, **Red User**.
- To sort and constrain events, or to navigate within the current workflow page, refer to Use workflows, on page 676.
- To navigate between pages in the current workflow, keeping the current constraints, click the appropriate page link at the top left of the workflow page.
- To drill down to the next page in the Workflows, constraining on a specific value, refer to Use drill-down pages, on page 684.
- To delete some or all correlation events, check the checkboxes next to the events you want to delete and click **Delete**, or click **Delete All** and confirm you want to delete all the events in the current constrained view.
- To navigate to other event views to view associated events, refer to Navigate between workflows, on page 701.
- To view data in available sources external to your system, right-click an event value. The options you see depend on the data type and include public sources; other sources depend on the resources you have configured. For more information, refer to Event investigation using web-based resources, on page 638.
- To gather intelligence about an event, right-click an event value in the table and choose from a Cisco or third-party intelligence source. For example, you can get details about a suspicious IP address from Cisco Talos. The options you see depend on the data type and the integrations that are configured on your system. For more information, refer to Event investigation using web-based resources, on page 638.

Related Topics

Database event limits, on page 68
Workflow pages, on page 680

Correlation event fields

When a correlation rule triggers, the system generates a correlation event. The table describes the fields available in the correlation events table when a correlation rule triggers.

Table 126: Correlation event fields

Field	Description
Description	The description of the correlation event. The information in the description depends on how the rule was triggered. For example, if the rule was triggered by an operating system information update event, the new operating system name and confidence level appear.

Field	Description
Device	The name of the device that generated the event that triggered the policy violation.
Domain	The domain of the device whose monitored traffic triggered the policy violation. This field is only present if you have ever configured the Firewall Management Center for multitenancy.
Impact	The impact level assigned to the correlation event based on the correlation between intrusion data, discovery data, and vulnerability information. When searching this field, valid case-insensitive values are `Impact 0`, `Impact Level 0`, `Impact 1`, `Impact Level 1`, `Impact 2`, `Impact Level 2`, `Impact 3`, `Impact Level 3`, `Impact 4`, and `Impact Level 4`. Do not use impact icon colors or partial strings (for example, do not use `blue`, `level 1`, or `0`).
Ingress Interface or Egress Interface	The ingress or egress interface in the intrusion or connection event that triggered the policy violation.
Ingress Security Zone or Egress Security Zone	The ingress or egress security zone in the intrusion or connection event that triggered the policy violation.
Inline Result	One of: • a black down arrow, indicating that the system dropped the packet that triggered the intrusion rule • a gray down arrow, indicating that the system would have dropped the packet in an inline, switched, or routed deployment if you enabled the **Drop when Inline** intrusion policy option • blank, indicating that the triggered intrusion rule was not set to Drop and Generate Events When using this field to search for policy violations triggered by intrusion events, type either: • `dropped`, to specify whether the packet was dropped in an inline, switched, or routed deployment • `would have dropped`, to specify whether the packet would have been dropped if the intrusion policy had been set to drop packets in an inline, switched, or routed deployment Note that the system does not drop packets in a passive deployment, including when an inline set is in tap mode, regardless of the rule state or the drop behavior of the intrusion policy.
Policy	The name of the policy that was violated.
Priority	The priority of the correlation event, which is determined by the priority of either the triggered rule or the violated correlation policy. When searching this field, enter `none` for no priority.
Rule	The name of the rule that triggered the policy violation.
Security Intelligence Category	The name of the object that represents or contains the blocked IP address in the event that triggered the policy violation. When searching this field, specify the Security Intelligence category associated with the correlation event that triggered the policy violation. The Security Intelligence category can be the name of a Security Intelligence object, the global Block list, a custom Security Intelligence list or feed, or one of the categories in the Intelligence Feed.

Field	Description
Source Continent or Destination Continent	The continent associated with the source or destination host IP addresses in the event that triggered the policy violation.
Source Country or Destination Country	The country associated with the source or destination IP address in the event that triggered the policy violation.
Source Host Criticality or Destination Host Criticality	The user-assigned host criticality of the source or destination host involved in the correlation event: `None`, `Low`, `Medium`, or `High`. Note that only correlation events generated by rules based on discovery events, host input events, or connection events contain a source host criticality.
Source IP or Destination IP	The IP address of the source or destination host in the event that triggered the policy violation.
Source Port/ICMP Type or Destination Port/ICMP Code	The source port or ICMP type for the source traffic or the destination port or ICMP code for destination traffic associated with the event that triggered the policy violation.
Source User or Destination User	The name of the user logged in to the source or destination host in the event that triggered the policy violation.
Time	The date and time that the correlation event was generated. This field is not searchable.
Count	The number of events that match the information that appears in each row. Note that the **Count** field appears only after you apply a constraint that creates two or more identical rows. This field is not searchable.

Related Topics

Event searches, on page 707

Analyze compliance events using compliance allow list workflows

Analyze compliance events and violations generated for your network, determining the compliance of your network assets, and manage workflow-based event analysis efficiently.

The Firewall Management Center provides a set of workflows that you can use to analyze the allow list events and violations that are generated for your network. The workflows are, along with the network map and dashboard, a key source of information about the compliance of your network assets.

The system provides predefined workflows for allow list events and violations. You can also create custom workflows. When you are using a compliance allow list workflow, you can perform many common actions.

Before you begin

You must be an Admin, Security Analyst, or Discovery Admin user to perform this task.

Follow these steps to analyze compliance events using compliance allow list workflows:

Procedure

Step 1 Access an allow list workflow using the **Analysis** > **Correlation** menu.

Step 2 Use the available workflow options as needed:

- Switch Workflow: To use a different workflow, including a custom workflow, click **(switch workflow)**.
- Time Range: To adjust the time range, which is useful if no events appear, refer to Change the time window, on page 695.
- Host Profile: To view the host profile for an IP address, click **Host Profile**() or, for hosts with active indications of compromise (IOC) tags, the **Compromised Host** that appears next to the IP address.
- User Profile (events only): To view user identity information, click the user icon that appears next to the **User Identity**, or for users associated with IOCs, **Red User**.
- Constrain: To constrain the columns that appear, click **Close** (×) in the column heading that you want to hide. In the pop-up window that appears, click **Apply**.

  Tip
  To hide or show other columns, select or clear the appropriate check boxes before you click **Apply**. To add a disabled column back to the view, expand the search constraints, then click the column name under Disabled Columns.

- Drill Down: Refer to Use drill-down pages, on page 684.
- Sort: To sort data in a workflow, click the column title. Click the column title again to reverse the sort order.
- Navigate This Page: Refer to Workflow page traversal tools, on page 681.
- Navigate Between Pages: To navigate between pages in the current workflow, keeping the current constraints, click the appropriate page link at the top left of the workflow page.
- Navigate Between Event Views: To navigate to other event views to view associated events, click **Jump to** and select the event view from the drop-down list.
- Delete Events (events only): To delete some or all items in the current constrained view, select the check boxes next to items you want to delete and click **Delete** or click **Delete All**.

Related Topics

Workflow pages, on page 680
Configure event view settings, on page 234

View allow list events

After its initial evaluation, the system generates an *allow list event* whenever a monitored host goes out of compliance with an active allow list. Allow list events are a special kind of correlation event and are logged to the Firewall Management Center correlation event database.

You can use the Firewall Management Center to view a table of compliance allow list events. You can then adjust the event view to display the information you need.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page that appears when you access allow list events differs depending on the workflow you use. Use a predefined workflow to display a table view of events. You can also create a custom workflow that displays only the information that matches your specific needs.

Before you begin

You must be an Admin, Security Analyst, or Discovery Admin user to perform this task.

Complete these steps to view compliance allow list events:

Procedure

Step 1 Choose **Analysis** > **Correlation** > **Allow List** > **Events**.

Step 2 You have these options:

- To perform basic workflow actions, refer to Analyze compliance events using compliance allow list workflows, on page 962.
- To learn more about the contents of the columns in the table, refer to Allow list event fields, on page 964.
- To see more options, right-click values in the table.

Allow list event fields

Allow list events record instances of compliance violations, which can be reviewed and searched using workflows. Each event contains these fields:

Device

Name of the managed device that detected the allow list violation.

Description

Details of how the allow list was violated.

For example:

```
Client "AOL Instant Messenger" is not allowed.
```

Violations that involve an application protocol indicate the application protocol name and version, as well as the port and protocol (TCP or UDP) it is using. If you restrict prohibitions to a particular operating system, the description includes the operating system name. For example:

```
Server "ssh / 22 TCP (OpenSSH 3.6.1p2)" is not allowed on Operating System
"Linux Linux 2.4 or 2.6".
```

Domain

The domain of the host that has become non-compliant with the allow list. This field is only present if you have ever configured the Firewall Management Center for multitenancy.

Host criticality

The user-assigned criticality for the source host that is out of compliance with the allow list: None, Low, Medium, or High.

IP address

The IP address of the host that has become non-compliant with the allow list.

Policy

The name of the correlation policy violated; that is, the correlation policy that includes the allow list.

Port

The port, if any, associated with the discovery event that triggered an application protocol allow list violation (a violation that occurred as a result of a non-compliant application protocol). For other types of allow list violations, this field is blank.

Priority

The priority specified by the policy or allow list that triggered the policy violation. This is determined either by the priority of the allow list in a correlation policy or by the priority of the correlation policy itself. Note that the allow list priority overrides the priority of its policy. When searching this field, enter `none` for no priority.

Time

The date and time that the allow list event was generated. This field is not searchable.

User

The identity of any known user logged in to the host that has become non-compliant with the allow list.

Allow List

The name of the allow list.

Count

The number of events that match the information that appears in each row. The Count field appears only after you apply a constraint that creates two or more identical rows. This field is not searchable.

View allow list violations

The system keeps a record of the current *allow list violations* on your network. Each violation represents something disallowed running on one of your hosts. If a host becomes compliant, the system removes the now-corrected violation from the database.

The page you see when you access allow list violations differs depending on the workflow you use. The predefined workflows terminate in a host view, which contains a host profile for every host that meets your constraints. You can also create a custom workflow that displays only the information that matches your specific needs.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

Follow these steps to view compliance allow list violations:

Procedure

Step 1 Choose **Analysis** > **Correlation** > **Allow List** > **Violations**.

Step 2 You have these options:

- To perform basic workflow actions, refer to Analyze compliance events using compliance allow list workflows, on page 962.
- To learn more about the contents of the columns in the table, refer to Allow list violation fields, on page 966.
- To see more options, right-click values in the table.

You can view, filter, and manipulate the compliance allow list violations table according to your needs using the Firewall Management Center. Violations update in real-time as hosts become compliant.

Allow list violation fields

Allow list violations, which you can view and search using workflows, contain these fields.

Domain

The domain where the non-compliant host resides. This field is only present if you have ever configured the Firewall Management Center for multitenancy.

Information

Any available vendor, product, or version information associated with the allow list violation. For protocols that violate an allow list, this field also indicates whether the violation is due to a network or transport protocol.

IP address

The IP address of the non-compliant host.

Port

The port, if any, associated with the event that triggered an application protocol allow list violation (a violation that occurred as a result of a non-compliant application protocol). For other types of allow list violations, this field is blank.

Protocol

The protocol, if any, associated with the event that triggered an application protocol allow list violation (a violation that occurred as a result of a non-compliant application protocol). For other types of allow list violations, this field is blank.

Time

The date and time that the allow list violation was detected.

Type

The type of allow list violation, that is, whether the violation occurred as a result of a non-compliant:

- operating system (os) (When searching this field, enter **os** or **operating system**.)
- application protocol (server)
- client
- protocol
- web application (web) (When searching this field, enter **web application**.)

Allow List

The name of the allow list that was violated.

Count

The number of events that match the information that appears in each row. The Count field appears only after you apply a constraint that creates two or more identical rows. This field is not searchable.

Remediation status events

A remediation status event is a database log entry that records when a remediation is triggered in the system.

Remediation status event management

You can view remediation status events on the Remediation Status page. You can search, view, and delete remediation status events.

Related Topics

Remediation status table fields, on page 968

View remediation status events

Monitor remediation status events to track the progress and status of remediation activities. You can use predefined or custom workflows to visualize remediation data in a table view or other formats. The table view contains a row for each remediation status event. You can also create a custom workflow that displays only the information that matches your specific needs.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

Ensure you have Admin privileges to perform this task.

Procedure

Step 1 Choose **Analysis** > **Correlation** > **Status**.

Step 2 Optionally, adjust the time range as described in Change the time window, on page 695.

Step 3 Optionally, to use a different workflow, including a custom workflow, click **(switch workflow)** by the workflow title.

Tip
If you are using a custom workflow that does not include the table view of remediations, click **(switch workflow)** menu by the workflow title, then choose **Remediation Status**.

Step 4 You have these options to work with remediation status events:

- To learn more about the columns that appear, refer to Remediation status table fields, on page 968.
- To sort and constrain the events, refer to Use workflows, on page 676.
- To navigate to the correlation events view to see associated events, click **Correlation Events**.
- To bookmark the current page to quickly return to it later, click **Bookmark This Page**. To manage bookmarks, click **View Bookmarks**.
- To generate a report based on the data in the table view, click **Report Designer** as described in Create a report template from an event view, on page 536.
- To drill down to the next page in the workflow, refer to Use drill-down pages, on page 684.
- To delete remediation status events from the system, check the check boxes next to events you want to delete and click **Delete** or click **Delete All** and confirm you want to delete all the events in the current constrained view.
- To search for remediation status events, click **Search**.

Related Topics

Use workflows, on page 676

Remediation status table fields

The table describes the fields in the remediation status table that can be viewed and searched.

Table 127: Remediation status fields

Field	Description
Domain	The domain of the device whose monitored traffic triggered the policy violation, that in turn triggered the remediation. This field is only present if you have ever configured the Firewall Management Center for multitenancy.
Policy	The name of the correlation policy that was violated and triggered the remediation.
Remediation Name	The name of the remediation that was launched.
Result Message	A message that describes what happened when the remediation was launched. Status messages include: `Successful completion of remediation`, `Error in the input provided to the remediation module`, `Error in the remediation module configuration`, `Error logging into the remote device or server`, `Unable to gain required privileges on remote device or server`, `Timeout logging into remote device or server`, `Timeout executing remote commands or servers`, `The remote device or server was unreachable`, `The remediation was attempted but failed`, `Failed to execute remediation program`, and `Unknown/unexpected error`. If custom remediation modules are installed, you may see additional status messages that are implemented by the custom module.
Rule	The name of the correlation rule that triggered the remediation.
Time	The date and time when the Firewall Management Center launched the remediation.
Count	The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows. This field is not searchable.

Related Topics

Event searches, on page 707

Configure the remediation status events table

Adjust the layout of the remediation status events table, or filter the events by specific field values. If you disable columns, your view changes during your session. If you remove the first column, the table adds the Count column.

When you click a value within a row, the table view is constrained and does not open another page.

Tip Every page with a table view includes 'Table View' in its name.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

Ensure you have Admin privileges to perform this task.

Follow these steps to configure and constrain the remediation status events table.

Procedure

Step 1 Choose **Analysis** > **Correlation** > **Status**.

Tip
If you are using a custom workflow that does not include the table view of remediations, click **(switch workflow)** menu next to the workflow title, and select **Remediation Status**.

Step 2 You have these options:

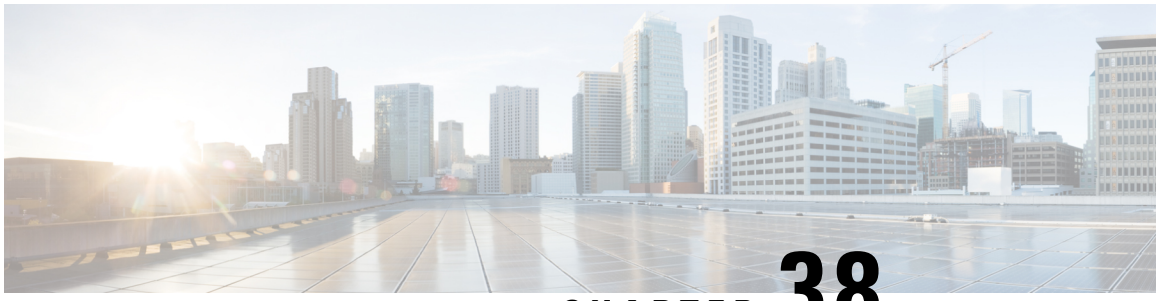
- To learn more about the columns that appear, refer to Remediation status table fields, on page 968.
- To sort and constrain the events, refer to Use workflows, on page 676.



PART IX

Correlation and Compliance

- Compliance Lists, on page 973
- Correlation Policies, on page 989
- Traffic profiling, on page 1027
- Remediations, on page 1039



CHAPTER 38

Compliance Lists

The following topics describe how to configure compliance allow lists before you add them to correlation policies.

- Introduction to compliance allow lists, on page 973
- Requirements and prerequisites for compliance, on page 978
- Create a compliance allow list, on page 979
- Manage compliance allow lists, on page 984
- Manage shared host profiles, on page 986

Introduction to compliance allow lists

A compliance allow list, sometimes abbreviated as an allow list, is a set of criteria that specifies which operating systems, applications (web and client), and protocols are allowed on hosts on your network. The system generates an event (violation) if a host is not on this list.

Components of compliance allow lists

A compliance allow list has two main components:

- **Targets**: Targets are the hosts you select for compliance evaluation. You can evaluate all monitored hosts or select specific hosts. You may constrain the selection by subnet, VLAN, or host attribute. In a multidomain deployment, you can target domains and subnets within a domain or across domains.
- **Host profiles**: Host profiles specify the compliance criteria for the targets. The global host profile is operating system agnostic. You can also configure operating-system specific host profiles, either unique to one allow list or shared across multiple allow lists.

The Talos Intelligence Group provides a default allow list with recommended settings. You can also create custom allow lists. A simple custom list might allow only hosts running a certain operating system. A more complex list might allow all operating systems, but specify which operating system a host must use to run a certain application protocol on a specific port.

Note The system can add hosts to the network map from exported NetFlow records, but the available information for these hosts is limited; see Differences between NetFlow and managed device data. This limitation may affect the way you build compliance allow lists.

Implementing compliance allow lists

To implement allow lists, add the list to an active correlation policy. The system evaluates the targets and assigns every host a corresponding attribute:

- Compliant: The host does not violate the list.
- Non-Compliant: The host violates the list.
- Not Evaluated: The host is not a target of the list, the host is currently being evaluated, or the system has insufficient information to determine whether the host is in compliance.



Note To delete the host attribute, delete its corresponding allow list. Deactivating, deleting, or removing an allow list from a correlation policy does **not** delete the host attribute, nor does it change the attribute's value for each host.

After its initial evaluation, the system generates an allow list event whenever a monitored host goes out of compliance with an active allow list; it also records an allow list violation.

You can use workflows, dashboards, and network maps to monitor system-wide compliance activity and determine when and how an individual host violates your allow lists. You can also automatically respond to such violations with remediations and alerts.

Restricting HTTP to web servers

Your security policy states that only web servers may run HTTP. You create an allow list that evaluates your entire network, excluding your web farm, to determine which hosts are running HTTP.

Using the network map and the dashboard, you can obtain an at-a-glance summary of the compliance of your network. In just a few seconds, you can determine exactly which hosts in your organization are running HTTP in violation of your policy, and take appropriate action.

Then, using the correlation feature, you can configure the system to alert you whenever a host that is not in your web farm starts running HTTP.

Related Topics

Configure correlation policies, on page 991

Compliance allow list target networks

A target network specifies the hosts you want to evaluate for compliance. An allow list can have more than one target network, and it evaluates hosts that meet the criteria of any of its targets.

Target network configuration and behavior

Initially, you constrain a target network by IP address or range. In multidomain deployments, the initial constraints also include a domain.

The system-provided default allow list targets all monitored hosts: 0.0.0.0/0 and ::/0. In a multidomain deployment, the default allow list is constrained to (and only available in) the Global domain

If you modify a target network or a host so that the host is no longer a valid target for the allow list, the host is no longer evaluated by the list and is considered neither compliant nor non-compliant.

Surveying and refining target networks

When you add a target network to an allow list, the system prompts you to survey the network map to help you characterize compliant hosts. The survey adds a target to the allow list that represents the hosts you surveyed.

You can survey a subnet or individual host. In a multidomain deployment, you can survey an entire domain, or you can survey across domains. Surveying an ancestor domain causes the system to survey that domain's descendants.

In addition to the added target, the survey also populates the allow list with one host profile for each operating system detected in the survey. These host profiles allow all the clients, application protocols, web applications, and protocols that the system has detected on the applicable operating systems.

After you survey a target network (or skip the survey), refine the target. You can exclude hosts by IP address, or constrain target networks by host attribute or VLAN.

Targeting domains with compliance allow lists

In a multidomain deployment, domains and target networks are closely linked:

- Leaf-domain administrators can create allow lists that evaluate hosts within their leaf domains.
- Higher-level domain administrators can create allow lists that evaluate hosts across domains. You can target different subnets in different domains in the same allow list.



Note In addition to targeting IP addresses and ranges in leaf domains, you can also constrain a target network using a higher-level domain. Targeting a subnet in a higher-level domain targets the **same** subnet in **each** of the descendant leaf domains. The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.

Global domain compliance scenario

Consider a scenario where you are a Global domain administrator, and you want to apply the same compliance criteria to web servers across the entire deployment. You can create one allow list in the Global domain that defines the compliance criteria. Then, constrain the allow list with target networks that specify the IP space (or individual IP addresses) of the web servers in each leaf domain.

Compliance allow list host profiles

In a compliance allow list, host profiles specify which operating systems, clients, application protocols, web applications, and protocols are allowed to run on the target hosts. There are three types of host profile you can use in a compliance allow list; each type appears differently in the compliance list editor.

Host profile types

Table 128: Compliance allow list host profile types

Host Profile Type	Appearance	Description
global	Any Operating System	Specifies what is allowed to run on target hosts, regardless of operating system.
operating-system specific	is listed in plain text	Specifies what is allowed to run on target hosts of a particular operating system.
shared	is listed in italics	Specifies operating-system criteria that can be used in multiple allow lists.

Operating system-specific host profiles

In a compliance allow list, operating-system specific host profiles indicate the allowed operating systems on your network. These profiles also specify the application protocols, clients, web applications, and protocols permitted to run on those operating systems.

You could require that compliant hosts run a particular version of Microsoft Windows. As another example, you could allow SSH to run on Linux hosts on port 22, and further restrict the vendor and version of the SSH client.

Host profile configuration guidelines

Create one host profile for each operating system you want to allow on your network. Create host profiles only for operating systems you want to allow on your network. For example, to make sure that all the hosts on your network are running Windows, configure the allow list so it contains host profiles for Windows.

Note Unidentified hosts remain in compliance with all allow lists until they are identified. You can, however, create an allow list host profile for unknown hosts. **Unidentified** hosts are hosts about which the system has not yet gathered enough information to identify their operating systems. **Unknown** hosts are hosts whose operating systems do not match known fingerprints.

Shared host profiles

In a compliance allow list, shared host profiles are tied to specific operating systems, but you can use each shared host profile in more than one allow list.

Global office deployment scenario

You might have offices worldwide with a separate allow list for each location, but you want to use the same profile for all hosts running Apple Mac OS X. You can create a shared profile for that operating system and use it in all your allow lists.

Built-in host profiles

The default allow list uses a special category of shared host profiles, called **built-in host profiles**. These profiles use built-in application protocols, web applications, protocols, and clients. In the compliance allow list editor, the system marks these profiles with the **Built-In Host Profile icon**.

In a multidomain deployment, the system displays shared host profiles created in the current domain; you can edit these profiles. It also displays shared host profiles from ancestor domains; you cannot edit these profiles. To view and edit shared host profiles created in a lower domain, switch to that domain.



Note If you modify a shared host profile (including built-in profiles), or modify a built-in application protocol, protocol, or client, your change affects every allow list that uses the modified element. If you unintentionally change or delete a built-in element, you can restore the default settings by resetting to factory defaults.

Allow violation triggers

This section describes the conditions that cause the system to trigger allow list violations and change host compliance status.

Automatic violation triggers

The allow list compliance of a host can change when the system:

- Detects a change in a host's operating system.
- Detects an identity conflict for a host's operating system or an application protocol on the host.
- Detects a new TCP server port active on a host, such as a port used by SMTP or web servers, or detects a new UDP server running on a host.
- Detects a change in a discovered TCP or UDP server running on a host, for example, a version change due to an upgrade.
- Detects a new client or web application running on a host.
- Drops a client or web application from its database due to inactivity.
- Detects that a host is communicating with a new network protocol or a new transport protocol.
- Detects a new jailbroken mobile device.
- Detects that a TCP or UDP port has closed or timed out on a host.

In addition, you can trigger a compliance change for a host by using the host input feature or the host profile to:

- Add a client, protocol, or server to a host.
- Delete a client, protocol, or server from a host.
- Set the operating system definition for a host.
- Change a host attribute for a host so that the host is no longer a valid target.

Note To avoid overwhelming you with events, the system does not generate allow list events for non-compliant hosts on its initial evaluation, nor hosts made non-compliant as a result of you modifying an active allow list or shared host profile. The violations, however, are still recorded. If you want to generate allow list events for all non-compliant targets, purge discovery data. Rediscovering network assets may trigger allow list events.

Operating system compliance

If your allow list specifies that only Microsoft Windows hosts are allowed on your network, and the system detects a host running Mac OS X, the system generates an allow list event. In addition, the host attribute associated with the allow list changes from Compliant to Non-Compliant for that host.

For the host in this example to come back into compliance, one of these must occur:

- You edit the allow list so that the Mac OS X operating system is allowed.
- You manually change the operating system definition of the host to Microsoft Windows.
- The system detects that the operating system has changed back to Microsoft Windows.

Deleting a non-compliant asset from the network map

If your allow list disallows the use of FTP, and you then delete FTP from the application protocols network map or from an event view, hosts running FTP become compliant. However, if the system detects the application protocol again, the system generates an allow list event and the hosts become non-compliant.

Triggering on complete information only

If your allow list allows only TCP FTP traffic on port 21, and the system detects indeterminate activity on port 21/TCP, the allow list does not trigger. The allow list triggers only when the system identifies the traffic as something other than FTP, or you use the host input feature to designate the traffic as non-FTP traffic. The system does not record a violation with only partial information.

Requirements and prerequisites for compliance

This table summarizes the requirements and prerequisites for compliance:

Requirement	Supported value
Model support	Any
Supported domains	Any
User roles	Admin

Create a compliance allow list

When you create a compliance allow list, the system prompts you to survey your network to create an initial target and to help you characterize compliant hosts.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Allow List**.

Step 2 Click **New Allow List**.

Step 3 Enter the **IP Address** and **Netmask** for an initial target network. In a multidomain deployment, choose the **Domain** where the target network resides.

Tip
To survey the entire monitored network, use the default values of 0.0.0.0/0 and ::/0.

Note
After you choose a domain for the target network, you cannot change it. Targeting a subnet in a higher-level domain targets the **same** subnet in **each** of the descendant leaf domains. The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.

Step 4 Add the target network:

- Add—To add the target network without a survey, click **Add**.
- Add and Survey Network—To add and survey the target network, click **Add and Survey Network**.
- Skip—To create an allow list without surveying your network, click **Skip**.

Step 5 Enter a new **Name** and **Description** for the allow list.

Step 6 **Allow Jailbroken Mobile Devices** on your network. Disabling this option causes jailbroken devices to generate allow list violations.

Step 7 Add at least one **Target Network** to the allow list, as described in Set target networks for a compliance allow list, on page 980.

Step 8 Characterize compliant hosts using **Allowed Host Profiles**:

- Global Host Profile—To edit the allow list's global host profile, click **Any Operating System** and proceed as described in Configure allow list host profiles, on page 981.
- Edit Surveyed Profiles—To edit an existing operating system-specific host profile created by a network survey, click its name and proceed as described in Configure allow list host profiles, on page 981.
- Create New Profiles—To create a new operating system-specific host profile for this allow list, click **Add** (+) next to **Allowed Host Profiles**, and proceed as described in Configure allow list host profiles, on page 981.
- Add Shared Host Profile—To add an existing shared host profile to the allow list, click **Add Shared Host Profile**, select the shared host profile you want to add, then click **OK**. Shared host profiles appear in italics.

Step 9 Click **Save Allow List**.

What to do next

- Add the allow list to an active correlation policy as described in Configure correlation policies, on page 991. The system immediately starts evaluating the allow list and generating violations.

Related Topics

Compliance allow list target networks, on page 974
Creating a compliance allow list based on selected hosts, on page 920
IP address conventions, on page 27

Set target networks for a compliance allow list

When you add a target network, you can survey it to characterize compliant hosts. This survey populates the allow list with one host profile for each operating system detected in the survey. These host profiles allow all the clients, application protocols, web applications, and protocols that the system has detected on the applicable operating systems.

Procedure

Step 1 In the compliance allow list editor, click **Add Target Network**.

Step 2 Enter the **IP Address** and **Netmask** for the target network.

Step 3 In a multidomain deployment, choose the **Domain** where the target network resides.

Note
After you choose a domain for the target network, you cannot change it. If you target a subnet in a higher-level domain, the **same** subnet is targeted in **each** descendant leaf domain. The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.

Step 4 Add the target network:

- Add — To add the target network without a survey, click **Add**.
- Add and Survey Network — To add and survey the target network, click **Add and Survey Network**.

Step 5 Optionally, click the new target to configure it further:

- Name — Enter a new **Name**.
- Add Networks — To target additional hosts, click **Add** (+), then enter the **IP Address** and **Netmask**. To exclude the network from allow list compliance, select **Exclude**.
- Add Host Attributes — To target hosts with a specific host attribute, click **Add** (+), then specify the **Attribute** and its **Value**.
- Add VLANs — To target a VLAN, click **Add** (+), then type a VLAN number (for 802.1q VLANs).

- Delete — To remove a target restriction, click **Delete** ( ).

Step 6 To immediately implement all changes made since the last time you saved, click **Save Allow List**.

Related Topics

Compliance allow list target networks, on page 974
IP address conventions, on page 27

Configure allow list host profiles

Host profiles specify the allow list's compliance criteria, that is, which operating systems, clients, application protocols, web applications, and protocols are allowed to run on the target hosts.

Every allow list has a global host profile which is operating-system agnostic. For example, instead of editing multiple Microsoft Windows and Linux host profiles to allow Mozilla Firefox, you can configure the global host profile to allow Firefox regardless of the operating system where it was detected.

You can also configure operating-system specific host profiles, either unique to one allow list or shared across allow lists.

Note If you modify a shared host profile (including built-in profiles), or modify a built-in application protocol, protocol, or client, your modification affects every allow list that uses it. To undo unintended changes to these built-in elements, reset to factory defaults.

Before you begin

- Create or edit a host profile within an allow list as described in Edit a compliance allow list, on page 985, or create or edit a shared host profile as described in Manage shared host profiles, on page 986.

Procedure

Step 1 In the compliance allow list host profile editor, configure a host profile:

- Name: Type a **Name**.
- Operating System: To restrict the host profile to a specific operating system, use the **OS Vendor**, **OS Name**, and **Version** drop-down lists. Because its purpose is to apply to hosts running any operating system, you cannot restrict a global host profile.
- Application Protocol: To allow an application protocol, click **Add** ( ) and proceed as described in Add an application protocol to a compliance allow list, on page 982.
- Client: To allow a client, click **Add** ( ) and proceed as described in Add a client to a compliance allow list, on page 983.
- Web Application: To allow a web application, click **Add** ( ) and proceed as described in Add a web application to a compliance allow list, on page 983.

- Protocol: To allow a protocol, click **Add** (+) and proceed as described in Add a protocol to a compliance allow list, on page 983.
- Delete: To disallow an item you previously allowed, click **Delete** ( ).
- Edit Properties: To edit the properties of an allowed application protocol, client, or protocol, click its name. The changes you make are reflected in every host profile that uses that element.

Tip
Select the appropriate **Allow all...** checkbox to allow all application protocols, clients, or web applications for hosts matching this profile.

Step 2 To immediately implement all changes made since the last time you saved, click **Save Allow List** (or **Save All Profiles** if you are editing a shared host profile).

Add an application protocol to a compliance allow list

Using allow list host profiles, you can allow application protocols either globally or on specific operating systems. Optionally, you can restrict the application protocol by port, vendor, or version. For example, you could allow a particular version of Open SSH to run on Linux hosts on port 22/TCP.

Procedure

Step 1 While you are creating or modifying a compliance allow list host profile, click **Add** (+) next to **Allowed Application Protocols** (or next to **Globally Allowed Application Protocols** if you are modifying the global host profile).

Step 2 You have two options:

- If the application protocols you want to allow are listed, select them. The web interface displays these application protocols as allowed or currently allowed by the compliance allow list.
- To allow an application protocol that is not in the list, select **<New Application Protocol>** and click **OK** to display the application protocol editor. Select the application protocol **Type** and **Protocol** you want to allow. Optionally, restrict the application protocol by **port**, **Vendor**, and **Version**.

  Note
  You must type the vendor and version exactly as they would appear in a table view of applications. If you do not specify a vendor or version, the allow list allows all vendors and versions as long as the type and protocol match.

Step 3 Click **OK**.

Step 4 To immediately implement all changes made since the last time you saved, click **Save Allow List**.

Add a client to a compliance allow list

Using allow list host profiles, you can allow clients either globally or on specific operating systems. Optionally, you can require that the client be a specific version. For example, you could allow only Microsoft Internet Explorer 10 to run on Microsoft Windows hosts.

Procedure

Step 1 To create or modify a compliance allow list host profile, click **Add** (+) next to **Allowed Clients** (or next to **Globally Allowed Clients** if you are modifying the global host profile).

Step 2 Choose one of these options:

- If the clients you want to allow are listed, select them. The web interface lists clients that have been allowed or are currently allowed by the allow list.
- To allow a client not in the list, select **<New Client>** and click **OK** to display the client editor. Select the **Client** you want to allow from the drop-down list, and, optionally, restrict the client to an allowed **Version**.

  Note
  You must type the version exactly as it would appear in a table view of clients. If you do not specify a version, all versions are allowed.

Step 3 Click **OK**.

Step 4 To immediately implement all changes made since the last time you saved, click **Save Allow List**.

Add a web application to a compliance allow list

Use allow list host profiles to allow web applications either globally or on specific operating systems.

Procedure

Step 1 To create or modify a compliance allow list host profile, click **Add** (+) next to **Allowed Web Applications** (or next to **Globally Allowed Web Applications** if you are modifying the global host profile).

Step 2 Select the web applications you want to allow.

Step 3 Click **OK**.

Step 4 To immediately implement all changes made since the last time you saved, click **Save Allow List**.

Add a protocol to a compliance allow list

Use allow list host profiles to allow protocols either globally or on specific operating systems. ARP, IP, TCP, and UDP are always allowed to run on any host; you cannot disallow them.

Procedure

Step 1 To create or modify a compliance allow list host profile, click the **Add** (+) icon next to **Allowed Protocols** (or next to **Globally Allowed Protocols** if you are modifying the global host profile).

Step 2 Choose one of the following options:

- If the protocols you want to allow are listed, select them. The web interface lists protocols that have been allowed or are currently allowed by the allow list.
- To allow a protocol not in the list, select **<New Protocol>** and click **OK** to display the protocol editor. From the **Type** drop-down list, select the protocol type (**Network** or **Transport**), then select the **Protocol** from the drop-down list.

Tip
Select **Other (manual entry)** to specify a protocol that is not in the list. For network protocols, type the appropriate number as listed in http://www.iana.org/assignments/ethernet-numbers/. For transport protocols, type the appropriate number as listed in http://www.iana.org/assignments/protocol-numbers/.

Step 3 Click **OK**

Step 4 To immediately implement all changes made since the last time you saved, click **Save Allow List**.

Manage compliance allow lists

You can use the Allow List page to manage compliance allow lists and shared host profiles. The default allow list represents recommended settings and uses a special category of shared host profiles, called **built-in host profiles**.

In a multidomain deployment, the system displays compliance allow lists created in the current domain, which you can edit. It also displays selected allow lists from ancestor domains, which you cannot edit. To view and edit allow lists created in a lower domain, switch to that domain.

Note The system does not display configurations from ancestor domains if the configurations expose information about unrelated domains, including names, managed devices, and so on. The default allow list is only available in the Global domain.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Allow List**.

Step 2 Manage your compliance allow lists:

- Create — To create a new allow list, click **New Allow List** and proceed as described in Create a compliance allow list, on page 979.

- Delete — To delete an allow list that is not in use, click **Delete** ( ), then confirm you want to delete the allow list. Deleting an allow list also removes its associated host attribute from all hosts on your network. If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Edit — To modify an existing allow list, click **Edit** ( ) and proceed as described in Edit a compliance allow list, on page 985. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Shared Host Profiles — To manage your allow lists' shared host profiles, click **Edit Shared Profiles** and proceed as described in Manage shared host profiles, on page 986.

Edit a compliance allow list

When you modify and save a compliance allow list that is included in an active correlation policy, the system immediately re-evaluates the compliance of the hosts in the allow list's target networks. Although this re-evaluation may bring some hosts into or out of compliance, the system does not generate any allow list events.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Allow List**.

Step 2 Next to the allow list you want to modify, click **Edit** ( ).

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Edit your compliance allow list:

Field	Description
Name and Description	To change the name or description, click the allow list name in the left panel to display basic allow list information, then type the new information.
Allow Jailbroken Devices	To allow jailbroken mobile devices on your network, click the allow list name in the left panel to display basic allow list information, then enable **Allow Jailbroken Mobile Devices**. Disabling this option causes jailbroken devices to generate allow list violations.
Add Allowed Host Profile	To create an operating system-specific host profile for this allow list, click **Add** ( ) next to Allowed Host Profiles and proceed as described in Configure allow list host profiles, on page 981.
Add Shared Host Profile	To add an existing shared host profile to the allow list, click **Add Shared Host Profile**, select the shared host profile you want to add, then click **OK**. Shared host profiles appear in italics.

Field	Description
Add Target Network	To add a new target network without surveying its hosts, click **Add** (+) next to Target Networks and proceed as described in Set target networks for a compliance allow list, on page 980.
Delete Host Profile	To delete a shared or operating-system specific host profile from the allow list, click **Delete** ( ) next to the host profile, then confirm your choice. Deleting a shared host profile removes it from the allow list, but does not delete the profile or remove it from any other allow lists that use it. You cannot delete an allow list's global host profile.
Delete Target Network	To remove a target network from the allow list, click **Delete** ( ) next to the network, then confirm your choice.
Edit Global Host Profile	To edit the allow list's global host profile, click **Any Operating System** and proceed as described in Configure allow list host profiles, on page 981.
Edit Other Host Profile	To edit a shared or operating-system specific host profile, click the host profile's name and proceed as described in Configure allow list host profiles, on page 981.
Edit Target Network	To edit a target network, click the network's name and proceed as directed in Set target networks for a compliance allow list, on page 980.

Step 4 To immediately implement all changes made since the last time you saved, click **Save Allow List**.

Manage shared host profiles

In a compliance allow list, **shared host profiles** are tied to specific operating systems, but you can use each shared host profile in more than one allow list. If you create multiple allow lists but want to use the same host profile to evaluate hosts running a particular operating system across the allow lists, use a shared host profile.

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

Note If you modify a shared host profile (including built-in profiles), or modify a built-in application protocol, protocol, or client, your change affects every allow list that uses it. If you make unintended changes to or delete these built-in elements, you can reset to factory defaults.

Procedure

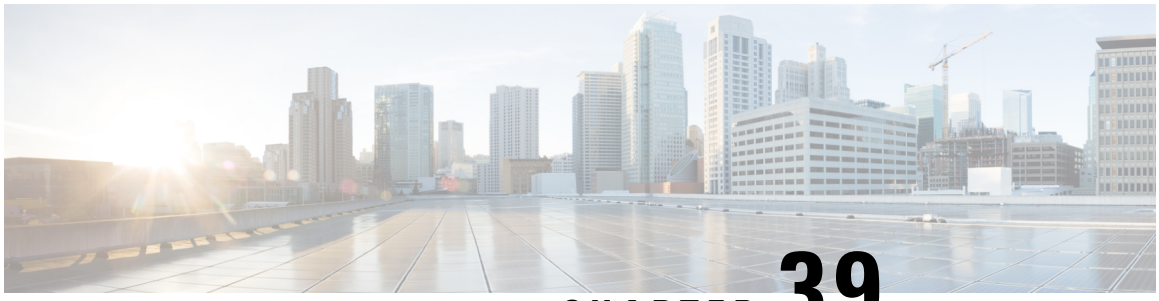
Step 1 Choose **Policies** > **Correlation**, then click **Allow List**.

Step 2 Click **Edit Shared Profiles**.

Step 3 Manage your shared host profiles:

- Create Shared Host Profile — To create a new shared host profile without surveying hosts, click **Add** (+) next to Shared Host Profiles and proceed as described in Configure allow list host profiles, on page 981.
- Create Shared Host Profile by Survey — To create multiple new shared host profiles by surveying a network, click **Add Target Network** and proceed as described in Set target networks for a compliance allow list, on page 980.
- Delete — To delete a shared host profile, click **Delete** (🗑), then confirm your choice.
- Edit — To modify an existing shared host profile (including a built-in shared host profile), click its name and proceed as described in Configure allow list host profiles, on page 981.
- Reset Built-In Host Profiles — To reset all built-in host profiles to factory defaults, click **Built-in Host Profiles**, then click **Reset to Factory Defaults** and confirm your choice.

Step 4 To immediately implement all changes made since the last time you saved, click **Save All Profiles**.



CHAPTER 39

Correlation Policies

The following topics describe how to configure correlation policies and rules.

- Correlation policies and rules, on page 989
- Requirements and prerequisites for compliance, on page 990
- Configure correlation policies, on page 991
- Configure correlation rules, on page 993
- Configure correlation response groups, on page 1024

Correlation policies and rules

A correlation policy is a security mechanism that

- responds in real time to threats to your network using correlation rules and compliance allow lists,
- generates correlation policy violations when network activity triggers either a correlation rule or compliance allow list within an active policy, and
- launches responses such as alerts and remediations when violations occur.

Correlation rules

When a correlation rule in an active correlation policy triggers, the system generates a *correlation event*. Correlation rules can trigger when:

- The system generates a specific type of event (connection, intrusion, malware, discovery, user activity, and so on).
- Your network traffic deviates from its normal profile.

You can constrain correlation rules in these ways:

- Add a *host profile qualification* to constrain the rule using information from the host profile of a host involved in the triggering event.
- Add a *connection tracker* to a correlation rule so that after the rule's initial criteria are met, the system begins tracking certain connections. Then, a correlation event is generated only if the tracked connections meet additional criteria.

- Add a *user qualification* to a correlation rule to track certain users or groups of users. For example, you can constrain a correlation rule so that it triggers only for a particular user's traffic, or traffic from a specific department.
- Add *snooze periods*. When a correlation rule triggers, a snooze period causes that rule not to trigger again for a specified interval. After the snooze period elapses, the rule can trigger again and start a new snooze period.
- Add *inactive periods*. During inactive periods, correlation rules do not trigger.

Although you can configure correlation rules without licensing your deployment, rules that use unlicensed components do not trigger.

A compliance allow list specifies which operating systems, applications (web and client), and protocols are allowed on hosts on your network. When a host violates an allow list used in an active correlation policy, the system generates an *allow list event*.

Responses to correlation policy violations include simple alerts and various remediations (such as scanning a host). You can associate each correlation rule or allow list with a single response or group of responses.

If network traffic triggers multiple rules or allow lists, the system launches all the responses associated with each rule and allow list.

In a multidomain deployment, you can create correlation policies at any domain level, using whatever rules, allow lists, and responses are available at that level. Higher-level domain administrators can perform correlation within or across domains:

- Constraining a correlation rule by domain matches events reported by that domain's descendants.
- Higher-level domain administrators can create compliance allow lists that evaluate hosts across domains. You can target different subnets in different domains in the same allow list.

Note The system builds a separate network map for each leaf domain. Using literal configurations (such as IP addresses, VLAN tags, and usernames) to constrain cross-domain correlation rules can have unexpected results.

Requirements and prerequisites for compliance

This table summarizes the requirements and prerequisites for compliance:

Requirement	Supported value
Model support	Any
Supported domains	Any
User roles	Admin

Configure correlation policies

Correlation policies help you organize and prioritize security event monitoring by combining correlation rules, compliance allow lists, alert responses, and remediations into a single policy framework.

Use correlation rules, compliance allow lists, alert responses, and remediations to build correlation policies.

In a multidomain deployment, you can create correlation policies at any domain level, using whatever constituent configurations are available at that level.

You can assign a priority to each correlation policy, as well as to each rule and allow list used in that policy. Rule and allow list priorities override correlation policy priorities. If network traffic violates the correlation policy, the resultant correlation events display the policy priority value, unless the violated rule or allow list has its own priority.

Procedure

Step 1 Choose **Policies** > **Correlation**.

Step 2 Click **Create Policy**.

Step 3 Enter a **Policy Name** and **Policy Description**.

Step 4 From the **Default Priority** drop-down list, choose a priority for the policy. Choose **None** to use rule priorities only.

Step 5 Click **Add Rules**, check the rules and allow lists that you want to use in the policy, then click **Add**.

Step 6 From the **Priority** list for each rule or allow list, choose a priority:

- A priority value from 1 to 5
- **None**
- **Default** to use the policy's default priority

Step 7 Add responses to rules and allow lists as described in Add responses to rules and allow lists, on page 991.

Step 8 Click **Save**.

What to do next

Activate the policy by clicking the slider.

Add responses to rules and allow lists

You can associate each correlation rule or allow list with a single response or group of responses. If network traffic triggers multiple rules or allow lists, the system launches all the responses associated with each rule and allow list. In a multidomain deployment, you can use responses created in the current domain or in ancestor domains.



Note An Nmap remediation does not launch when used as a response to a traffic profile change.

Procedure

Step 1 In the correlation policy editor, next to a rule or allow list where you want to add responses, click **Responses** ( ).

Step 2 Under Unassigned Responses, choose the responses you want to launch when the rule or allow list triggers, and click the ^.

Step 3 Click **Update**.

Manage correlation policies

Manage correlation policies to control how the system processes events and triggers responses in your security environment.

Changes made to active correlation policies take effect immediately.

When you activate a correlation policy, the system immediately begins processing events and triggering responses. Note that the system does not generate allow list events for non-compliant hosts on its initial, post-activation evaluation.

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

Note The system does not display configurations from ancestor domains if the configurations expose information about unrelated domains, including names, managed devices, and so on.

Procedure

Step 1 Choose **Policies** > **Correlation**.

Step 2 Manage your correlation policies:

- Activate or Deactivate — Click the slider. If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Create — Click **Create Policy**; see Configure correlation policies, on page 991.
- Edit — Click **Edit** ( ); see Configure correlation policies, on page 991. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Delete — Click **Delete** ( ). If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

You have successfully managed your correlation policies. Active policies will immediately begin processing events and triggering responses according to their configuration.

Configure correlation rules

Correlation rules enable automated detection and response to specific network events or combinations of events. You can create simple rules that trigger on single event types or complex rules with multiple conditions and constraints.

A simple correlation rule requires only that an event of a certain type occurs. You do not need to provide more specific conditions. For example, correlation rules based on traffic profile changes do not require conditions. You can also create complex correlation rules, with multiple conditions and added constraints.

When you create correlation rule trigger criteria, host profile qualifications, user qualifications, or connection trackers, the syntax varies but the mechanics remain consistent.



Note In a multidomain deployment, constraining a correlation rule by an ancestor domain matches events reported by that domain's descendants.

Before you begin

- Confirm that your deployment is collecting the type of information you want to use to trigger correlation events. For example, the information available for any individual connection or connection summary event depends on several factors, including the detection method, the logging method, and event type. The system can add hosts to the network map from exported NetFlow records, but the available information for these hosts is limited; see Differences between NetFlow and managed device data.

Follow these steps to configure correlation rules:

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Rule Management**.

Step 2 Click **Create Rule**.

Step 3 Enter a **Rule Name** and **Rule Description**.

Step 4 Optionally, choose a **Rule Group** for the rule.

Step 5 Choose a base event type and, optionally, specify additional trigger criteria for the correlation rule. You can choose from these base event types:

- **a VPN troubleshooting event occurs**—See Syntax for VPN troubleshoot event trigger criteria, on page 994.
- **an intrusion event occurs**—See Syntax for intrusion event trigger criteria, on page 995.
- **a malware event occurs**—See Syntax for malware event trigger criteria, on page 997.
- **a discovery event occurs**—See Discovery event trigger criteria and event types, on page 999.
- **user activity is detected**—See Syntax for user activity event trigger criteria, on page 1000.
- **a host input event occurs**—See Syntax for host input event trigger criteria, on page 1001.
- **a connection event occurs**—See Syntax for Connection Event Trigger Criteria, on page 1002.

- **a traffic profile changes**—See Syntax for traffic profile changes, on page 1005.

Step 6 Optionally, further constrain the correlation rule by adding any or all of these options:

- Host Profile Qualification—Click **Add Host Profile Qualification**; see Syntax for correlation host profile qualifications, on page 1007.
- Connection Tracker—Click **Add Connection Tracker**; see Connection trackers, on page 1010.
- User Qualification—Click **Add User Qualification**; see Syntax for user qualifications, on page 1009.
- Snooze Period—Under Rule Options, use the **Snooze** text field and drop-down list to specify the interval that the system should wait to trigger a correlation rule again, after the rule triggers.
- Inactive Period—Under Rule Options, click **Add Inactive Period**. Using the text field and drop-down lists, specify when and how often you want the system to refrain from evaluating network traffic against the correlation rule.

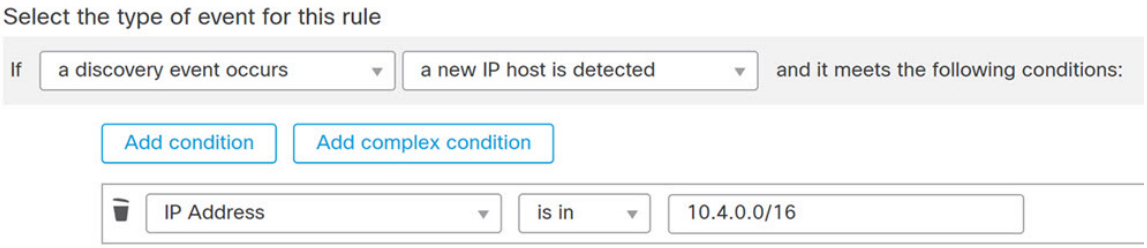
Tip
To remove a snooze period, specify an interval of `0` (seconds, minutes, or hours).

Step 7 Click **Save Rule**.

The correlation rule is saved and can be used in correlation policies to monitor network events and trigger automated responses.

Example simple correlation rule

This simple correlation rule triggers if a new host is detected in a specific subnet. Note that when the category represents an IP address, choosing **is in** or **is not in** as the operator allows you to specify whether the IP address *is in* or *is not in* a block of IP addresses, as expressed in special notation such as CIDR.



What to do next

- Use the rule in correlation policies as described in Configure correlation policies, on page 991.

Syntax for VPN troubleshoot event trigger criteria

This table describes how to build a correlation rule condition when you choose a VPN troubleshooting event as the base event.

Table 129: Syntax for VPN troubleshoot events

If you specify...	Choose an operator, then enter...
Device	Choose one or more devices with VPN troubleshoot syslog enabled.
Syslog Message Class	Choose the VPN syslog message class. When syslog with the selected message class is generated, it fulfils the correlation rule criteria and generates a correlation event.
Syslog Message ID	Specify the VPN syslog message IDs for the correlation rule.
Syslog Message Text	Specify the VPN syslog message text for the correlation rule.
Syslog Severity	Specify the VPN syslog severity. VPN troubleshoot syslog generated for the selected severity triggers the correlation event.
Username	Mention the VPN user name for whose traffic a correlation event need to be generated.

Syntax for intrusion event trigger criteria

This table describes how to build a correlation rule condition when you choose an intrusion event as the base event.

Table 130: Syntax for intrusion events

If you specify...	Choose an operator, then...
Access Control Policy	Choose one or more access control policies that use the intrusion policy that generated the intrusion event.
Access Control Rule Name	Enter all or part of the name of the access control rule that uses the intrusion policy that generated the intrusion event.
Application Protocol	Choose one or more application protocols associated with the intrusion event.
Application Protocol Category	Choose one or more category of application protocol.
Classification	Choose one or more classifications.
Client	Choose one or more clients associated with the intrusion event.
Client Category	Choose one or more category of client.
Destination Country or Source Country	Choose one or more countries associated with the source or destination IP address in the intrusion event.
Destination IP, Source IP, Both Source IP and Destination IP, or Either Source IP or Destination IP	Enter a single IP address or address block.

If you specify...	Choose an operator, then...
Destination Port/ICMP Code or Source Port/ICMP Type	Enter the port number or ICMP type for source traffic or the port number or ICMP code for destination traffic.
Device	Choose one or more devices that may have generated the event.
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants. This field is only present if you have ever configured the Firewall Management Center for multitenancy.
Egress Interface or Ingress Interface	Choose one or more interfaces.
Egress Security Zone or Ingress Security Zone	Choose one or more security zones or tunnel zones.
Generator ID	Choose one or more preprocessors.
Impact Flag	Choose the impact level assigned to the intrusion event. Because no operating system information is available for hosts added to the network map from NetFlow data, the system cannot assign Vulnerable (impact level 1: red) impact levels for intrusion events involving those hosts. In such cases, use the host input feature to manually set the operating system identity for the hosts.
Inline Result	Choose whether the system **dropped** or **would have dropped** packets as a result of the intrusion policy violation. The system can drop packets in an inline, switched, or routed deployment. It does not drop packets in a passive deployment, including when an inline set is in tap mode, regardless of intrusion rule state or the drop behavior of the intrusion policy.
Intrusion Policy	Choose one or more intrusion policies that generated the intrusion event.
IOC Tag	Choose whether an indication of compromise tag was set as a result of the intrusion event.
Priority	Choose the rule priority. For rule-based intrusion events, the priority corresponds to either the value of the `priority` keyword or the value for the `classtype` keyword. For other intrusion events, the priority is determined by the decoder or preprocessor.
Protocol	Enter the name or number of the transport protocol as listed in http://www.iana.org/assignments/protocol-numbers.
Rule Message	Enter all or part of the rule message.
Rule SID	Enter a single Snort ID (SID) or multiple SIDs separated by commas. If you choose **is in** or **is not in** as the operator, you cannot use the multi-selection pop-up window. You must enter a comma-separated list of SIDs.

If you specify...	Choose an operator, then...
Rule Type	Specify whether the rule is local. Local rules include custom standard text intrusion rules, standard text rules that you modified, and any new instances of shared object rules created when you saved the rule with modified header information.
SSL Actual Action	Choose the SSL rule action that indicates how the system handled an encrypted connection.
SSL Certificate Fingerprint	Enter the fingerprint of the certificate used to encrypt the traffic, or choose a subject common name associated with the fingerprint.
SSL Certificate Subject Common Name (CN)	Enter all or part of the subject common name of the certificate used to encrypt the session.
SSL Certificate Subject Country (C)	Choose one or more subject country codes of the certificate used to encrypt the session.
SSL Certificate Subject Organization (O)	Enter all or part of the subject organization name of the certificate used to encrypt the session.
SSL Certificate Subject Organizational Unit (OU)	Enter all or part of the subject organizational unit name of the certificate used to encrypt the session.
SSL Flow Status	Choose one or more statuses based on the result of the system's attempt to decrypt the traffic.
Username	Enter the username of the user logged into the source host in the intrusion event.
VLAN ID	Enter the innermost VLAN ID associated with the packet that triggered the intrusion event
Web Application	Choose one or more web applications associated with the intrusion event.
Web Application Category	Choose one or more category of web application.

Syntax for malware event trigger criteria

To base a correlation rule on a malware event, first specify the type of malware event you want to use. Your choice determines the set of trigger criteria you can use. You can choose:

- **by endpoint-based malware detection** (detection by Secure Endpoint)
- **by network-based malware detection** (detection by Malware Defense)
- **by retrospective network-based malware detection** (retroactive detection by Malware Defense)

This table describes how to build a correlation rule condition when you choose a malware event as the base event.

Table 131: Syntax for malware events

If you specify...	Choose an operator, then...
Application Protocol	Choose one or more application protocols associated with the malware event.

If you specify...	Choose an operator, then...
Application Protocol Category	Choose one or more category of application protocol.
Client	Choose one or more clients associated with the malware event.
Client Category	Choose one or more category of client.
Destination Country or Source Country	Choose one or more countries associated with the source or destination IP address in the malware event.
Destination IP, Host IP, or Source IP	Enter a single IP address or address block.
Destination Port/ICMP Code	Enter the port number or ICMP code for destination traffic.
Disposition	Choose either or both **Malware** or **Custom Detection**.
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants. This field is only present if you have ever configured the Firewall Management Center for multitenancy.
Event Type	Choose one or more event types associated with the malware event detected by Secure Endpoint.
File Name	Enter the name of the file.
File Type	Choose the file type.
File Type Category	Choose one or more file type categories.
IOC Tag	Choose whether an indication of compromise tag **is** or **is not** set as a result of the malware event.
SHA-256	Enter or paste the SHA-256 hash value of the file.
SSL Actual Action	Choose the SSL rule action that indicates how the system handled an encrypted connection.
SSL Certificate Fingerprint	Enter the fingerprint of the certificate used to encrypt the traffic, or choose a subject common name associated with the fingerprint.
SSL Certificate Subject Common Name (CN)	Enter all or part of the subject common name of the certificate used to encrypt the session.
SSL Certificate Subject Country (C)	Choose one or more subject country codes of the certificate used to encrypt the session.
SSL Certificate Subject Organization (O)	Enter all or part of the subject organization name of the certificate used to encrypt the session.
SSL Certificate Subject Organizational Unit (OU)	Enter all or part of the subject organizational unit name of the certificate used to encrypt the session.
SSL Flow Status	Choose one or more statuses based on the result of the system's attempt to decrypt the traffic.
Source Port/ICMP Type	Enter the port number or ICMP type for source traffic.

If you specify...	Choose an operator, then...
Web Application	Choose one or more web applications associated with the malware event.
Web Application Category	Choose one or more category of web application.

Discovery event trigger criteria and event types

You cannot trigger a correlation rule on hops changes, or when the system drops a new host due to reaching the host limit. You can, however, choose **there is any type of event** to trigger the rule when any type of discovery event occurs.

To base a correlation rule on a discovery event, first specify the type of discovery event you want to use. Your choice determines the set of trigger criteria you can use. The following table lists the discovery event types you can choose.

Table 132: Correlation Rule Trigger Criteria vs Discovery Event Types

Choose this option...	To use this discovery event type...
a client has changed	Client Update
a client timed out	Client Timeout
a host IP address is reused	DHCP: IP Address Reassigned
a host is deleted because the host limit was reached	Host Deleted: Host Limit Reached
a host is identified as a network device	Host Type Changed to Network Device
a host timed out	Host Timeout
a host's IP address has changed	DHCP: IP Address Changed
a NETBIOS name change is detected	NETBIOS Name Change
a new client is detected	New Client
a new IP host is detected	New Host
a new MAC address is detected	Additional MAC Detected for Host
a new MAC host is detected	New Host
a new network protocol is detected	New Network Protocol
a new transport protocol is detected	New Transport Protocol
a TCP port closed	TCP Port Closed
a TCP port timed out	TCP Port Timeout
a UDP port closed	UDP Port Closed

Choose this option...	To use this discovery event type...
a UDP port timed out	UDP Port Timeout
a VLAN tag was updated	VLAN Tag Information Update
an IOC was set	Indication of Compromise
an open TCP port is detected	New TCP Port
an open UDP port is detected	New UDP Port
the OS information for a host has changed	New OS
the OS or server identity for a host has a conflict	Identity Conflict
the OS or server identity for a host has timed out	Identity Timeout
there is any kind of event	any event type
there is new information about a MAC address	MAC Information Change
there is new information about a TCP server	TCP Server Information Update
there is new information about a UDP server	UDP Server Information Update

Syntax for user activity event trigger criteria

To base a correlation rule on user activity, first choose the type of user activity you want to use. Your choice determines the set of trigger criteria you can use. You can choose:

- **a new user identity is detected**
- **a user logs into a host**

This table describes how to build a correlation rule condition when you choose user activity as the base event.

Table 133: Syntax for user activity

If you specify...	Choose an operator, then...
Device	Choose one or more devices that may have detected the user activity.
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants. This field is only present if you have ever configured the Firewall Management Center for multitenancy.
IP Address	Enter a single IP address or address block.
Username	Enter a username.

Syntax for host input event trigger criteria

To base a correlation rule on a host input event, first specify the type of host input event you want to use. Your choice determines the set of trigger criteria you can use. The host input event types you can choose are listed in correlation rule trigger criteria.

You cannot trigger a correlation rule when you add, delete, or change the definition of a user-defined host attribute, or set a vulnerability impact qualification.

Table 134: Correlation rule trigger criteria vs host input event types

Choose this option...	To trigger the rule on this event type...
a client is added	Add Client
a client is deleted	Delete Client
a host is added	Add Host
a protocol is added	Add Protocol
a protocol is deleted	Delete Protocol
a scan result is added	Add Scan Result
a server definition is set	Set Server Definition
a server is added	Add Port
a server is deleted	Delete Port
a vulnerability is marked invalid	Vulnerability Set Invalid
a vulnerability is marked valid	Vulnerability Set Valid
an address is deleted	Delete Host/Network
an attribute value is deleted	Host Attribute Delete Value
an attribute value is set	Host Attribute Set Value
an OS definition is set	Set Operating System Definition
host criticality is set	Set Host Criticality

This table describes how to build a correlation rule condition when you choose a host input event as the base event.

Table 135: Syntax for host input events

If you specify...	Choose an operator, then...
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants. This field is only present if you have ever configured the Firewall Management Center for multitenancy.

If you specify...	Choose an operator, then...
IP Address	Enter a single IP address or address block.
Source	Choose the source for the host input data.
Source Type	Choose the type of the source for the host input data.

Syntax for Connection Event Trigger Criteria

To base a correlation rule on a connection event, first specify the type of connection event you want to use. Note that the information available for a connection event can vary depending on how, why, and when the system logged the connection. You can choose:

- **at either the beginning or the end of the connection**
- **at the beginning of the connection**
- **at the end of the connection**

The following table describes how to build a correlation rule condition when you choose a connection event as the base event.

Table 136: Syntax for Connection Events

If you specify...	Choose an operator, then...
Access Control Policy	Choose one or more access control policies that logged the connection.
Access Control Rule Action	Choose one or more actions associated with the access control rule that logged the connection. Choose **Monitor** to trigger correlation events when network traffic matches the conditions of any Monitor rule, regardless of the rule or default action that later handles the connection.
Access Control Rule	Enter all or part of the name of the access control rule that logged the connection. You can enter the name of any Monitor rule whose conditions were matched by a connection, regardless of the rule or default action that later handled the connection.
Application Protocol	Choose one or more application protocols associated with the connection.
Application Protocol Category	Choose one or more categories of application protocol.
Client	Choose one or more clients.
Client Category	Choose one or more categories of client.
Client Version	Enter the version number of the client.
Connection Duration	Enter the duration of the connection event, in seconds.

If you specify...	Choose an operator, then...
Connection Type	Specify whether you want to trigger the correlation rule based on how the connection information was obtained: • Choose **is** and **Netflow** for connection events generated from exported NetFlow data. • Choose **is not** and **Netflow** for connection events detected by a managed device.
Destination Country or Source Country	Choose one or more countries associated with the source or destination IP address in the connection event.
Device	Choose one or more devices that either detected the connection, or that processed the connection (for connection data from exported NetFlow records).
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants. This field is only present if you have ever configured the Firewall Management Center for multitenancy.
Egress Interface or Ingress Interface	Choose one or more interfaces.
Egress Security Zone or Ingress Security Zone	Choose one or more security zones or tunnel zones .
Initiator Bytes, Responder Bytes, or Total Bytes	Enter one of: • The number of bytes sent (**Initiator Bytes**). • The number of bytes received (**Responder Bytes**). • The number of bytes both sent and received (**Total Bytes**).
Initiator IP, Responder IP, Both Initiator and Responder IP, or Either Initiator IP or Responder IP	Specify a single IP address or address block.
Initiator Packets, Responder Packets, or Total Packets	Enter one of: • The number of packets sent (**Initiator Packets**). • The number of packets received (**Responder Packets**). • The number of packets both sent and received (**Total Packets**)
Initiator Port/ICMP Type or Responder Port/ICMP Code	Enter the port number or ICMP type for initiator traffic or the port number or ICMP code for responder traffic.
IOC Tag	Specify whether an indication of compromise tag **is** or **is not** set due to the connection event.
NetBIOS Name	Enter the NetBIOS name of the monitored host in the connection.
NetFlow Device	Choose the IP address of the NetFlow exporter you want to use to trigger the correlation rule. If you did not add any NetFlow exporters to the network discovery policy, the **NetFlow Device** drop-down list is blank.

If you specify...	Choose an operator, then...
Prefilter Policy	Choose one or more prefilter policies that handled the connection.
Reason	Choose one or more reasons associated with the connection event.
Security Intelligence Category	Choose one or more Security Intelligence categories associated with the connection event. To use Security Intelligence Category as a condition for end-of-connection events, set that category to **Monitor** instead of **Block** in your access control policy.
SSL Actual Action	Specify the SSL rule action that indicates how the system handled an encrypted connection.
SSL Certificate Fingerprint	Enter the fingerprint of the certificate used to encrypt the traffic, or choose a subject common name associated with the fingerprint.
SSL Certificate Status	Choose one or more statuses associated with the certificate used to encrypt the session.
SSL Certificate Subject Common Name (CN)	Enter all or part of the subject common name of the certificate used to encrypt the session.
SSL Certificate Subject Country (C)	Choose one or more subject country codes of the certificate used to encrypt the session.
SSL Certificate Subject Organization (O)	Enter all or part of the subject organization name of the certificate used to encrypt the session.
SSL Certificate Subject Organizational Unit (OU)	Enter all or part of the subject organizational unit name of the certificate used to encrypt the session.
SSL Cipher Suite	Choose one or more cipher suites used to encrypt the session.
SSL Encrypted Session	Choose **Successfully Decrypted** .
SSL Flow Status	Choose one or more statuses based on the result of the system's attempt to decrypt the traffic.
SSL Policy	Choose one or more SSL policies that logged the encrypted connection.
SSL Rule Name	Enter all or part of the name of the SSL rule that logged the encrypted connection.
SSL Server Name	Enter all or part of the name of the server with which the client established an encrypted connection.
SSL URL Category	Choose one or more URL categories for the URL visited in the encrypted connection.
SSL Version	Choose one or more SSL or TLS versions used to encrypt the session.
TCP Flags	Choose a TCP flag that a connection event must contain in order to trigger the correlation rule. Only connection data generated from NetFlow records contains TCP flags.
Transport Protocol	Enter the transport protocol used by the connection: **`TCP`** or **`UDP`** .
Tunnel/Prefilter Rule	Enter all or part of the name of the tunnel or prefilter rule that handled the connection.
URL	Enter all or part of the URL visited in the connection.
URL Category	Choose one or more URL categories for the URL visited in the connection.

If you specify...	Choose an operator, then...
URL Reputation	Choose one or more URL reputation values for the URL visited in the connection.
Username	Enter the username of the user logged in to either host in the connection.
Web Application	Choose one or more web applications associated with the connection.
Web Application Category	Choose one or more categories of web application.

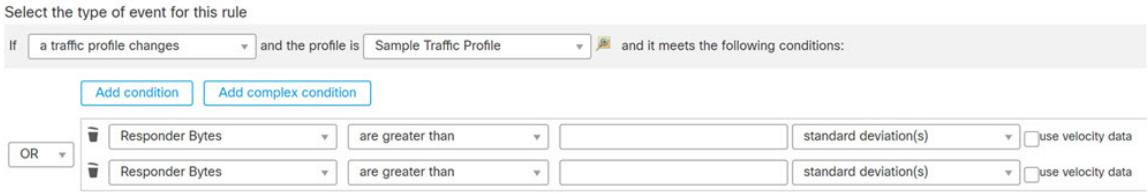
Syntax for traffic profile changes

To base a correlation rule on a traffic profile change, first choose the traffic profile you want to use. The rule triggers when network traffic deviates from the pattern characterized by the profile you choose.

You can trigger the rule based on either raw data or on the statistics calculated from the data. For example, you could write a rule that triggers if the amount of data traversing your network (measured in bytes) suddenly spikes, which could indicate an attack or other security policy violation. You could specify that the rule trigger if either:

- the number of bytes traversing your network spikes above a certain number of bytes
- the number of bytes traversing your network spikes above a certain number of standard deviations above or below the mean amount of traffic

  Note that to create a rule that triggers when the number of bytes traversing your network falls outside a certain number of standard deviations (whether above or below), you must specify upper and lower bounds, as shown in this graphic.



  To create a rule that triggers when the number of bytes traversing is greater than a certain number of standard deviations *above* the mean, use only the first condition shown in the graphic.

  To create a rule that triggers when the number of bytes traversing is greater than a certain number of standard deviations *below* the mean, use only the second condition.

Check the **use velocity data** check box to trigger the correlation rule based on rates of change between data points. If you wanted to use velocity data in the above example, you could specify that the rule triggers if either:

- the change in the number of bytes traversing your network spikes above or below a certain number of standard deviations above the mean rate of change
- the change in the number of bytes traversing your network spikes above a certain number of bytes

This table describes how to build a condition in a correlation rule when you choose a traffic profile change as the base event.

Table 137: Syntax for traffic profile changes

If you specify...	Choose an operator, then enter...	Then choose one of...
Number of Connections	the total number of connections detected **or** the number of standard deviations either above or below the mean that the number of connections detected must be in to trigger the rule	connections standard deviation(s)
Total Bytes, Initiator Bytes, or Responder Bytes	one of: • the total bytes transmitted (**Total Bytes**) • the number of bytes transmitted (**Initiator Bytes**) • the number of bytes received (**Responder Bytes**) **or** the number of standard deviations either above or below the mean that one of the above criteria must be in to trigger the rule	bytes standard deviation(s)
Total Packets, Initiator Packets, or Responder Packets	one of: • the total packets transmitted (**Total Packets**) • the number of packets transmitted (**Initiator Packets**) • the number of packets received (**Responder Packets**) **or** the number of standard deviations either above or below the mean that one of the above criteria must be in trigger the rule	packets standard deviation(s)
Unique Initiators	the number of unique hosts that initiated sessions **or** the number of standard deviations either above or below the mean that the number of unique initiators detected must be to trigger the rule	initiators standard deviation(s)
Unique Responders	the number of unique hosts that responded to sessions **or** the number of standard deviations either above or below the mean that the number of unique responders detected must be to trigger the rule	responders standard deviation(s)

Syntax for correlation host profile qualifications

To constrain a correlation rule based on the host profile of a host involved in the event, add a *host profile qualification*. You cannot add a host profile qualification to a correlation rule that triggers on a malware event, traffic profile change, or on the detection of a new IP host.

When you build a host profile qualification, first specify the host you want to use to constrain your correlation rule. The host you can choose depends on the rule's base event type:

- connection event — Choose **Responder Host** or **Initiator Host**.
- intrusion event — Choose **Destination Host** or **Source Host**.
- discovery event, host input event, or user activity — Choose **Host**.

This table describes how to build a host profile qualification for a correlation rule.

Table 138: Syntax for host profile qualifications

If you specify...	Choose an operator, then...
Application Protocol > Application Protocol	Choose an application protocol.
Application Protocol > Application Port	Enter the application protocol port number.
Application Protocol > Protocol	Choose a protocol.
Application Protocol Category	Choose a category.
Client > Client	Choose a client.
Client > Client Version	Enter the client version.
Client Category	Choose a category.
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants. This field is only present if you have ever configured the Firewall Management Center for multitenancy.
Hardware	Enter the hardware model for the mobile device. For example, to match all Apple iPhones, enter `iPhone`.
Host Criticality	Choose the host criticality.
Host Type	Choose one or more host types. You can choose between a normal host or one of several types of network device.
IOC Tag	Choose one or more indication of compromise tags.
Jailbroken	Choose **Yes** to indicate that the host in the event is a jailbroken mobile device or **No** to indicate that it is not.
MAC Address > MAC Address	Enter all or part of the MAC address of the host.

If you specify...	Choose an operator, then...
MAC Address > MAC Type	Choose whether the MAC type is ARP/DHCP detected: • the system positively identified the MAC address as belonging to the host (**is ARP/DHCP Detected**) • the system is seeing many hosts with that MAC address because, for example, there is a router between the device and the host (**is not ARP/DHCP Detected**) • the MAC type is irrelevant (**is any**)
MAC Vendor	Enter all or part of the MAC vendor of hardware used by the host.
Mobile	Choose **Yes** to indicate that the host in the event is a mobile device or **No** to indicate that it is not.
NetBIOS Name	Enter the NetBIOS name of the host.
Network Protocol	Enter the network protocol number as listed in http://www.iana.org/assignments/ethernet-numbers.
Operating System > OS Vendor	Choose one or more operating system vendor names.
Operating System > OS Name	Choose one or more operating system names.
Operating System > OS Version	Choose one or more operating system versions.
Transport Protocol	Enter the name or number of the transport protocol as listed in http://www.iana.org/assignments/protocol-numbers.
VLAN ID	Enter the VLAN ID number of the host.
Web Application	Choose a web application.
Web Application Category	Choose a category.
any available host attribute, including the default compliance allow list host attribute	Enter or choose the appropriate value, depending on the host attribute type.

Using implied or generic clients to build a host profile qualification

When system reports a detected client using an application protocol name followed by `client` (for example, `HTTPS client`), that client is an *implied* or *generic* client. In these cases, the system has not detected a particular client, but is inferring the existence of a client based on server response traffic.

To create a host profile qualification using an implied or generic client, constrain using the application protocol running on the responder host, not the client.

Using event data to build a host profile qualification

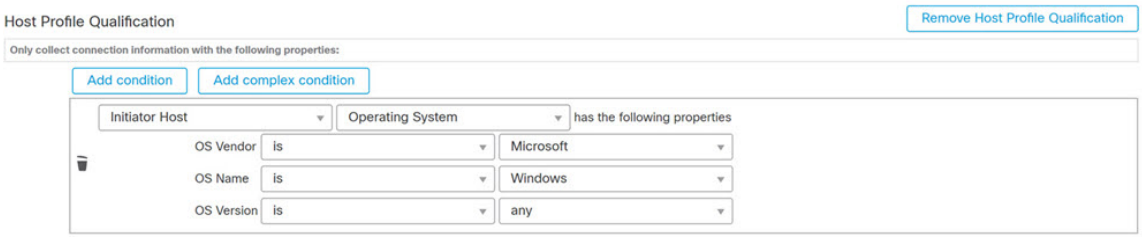
You can often use data from the correlation rule's base event when constructing a host profile qualification.

For example, assume your correlation rule triggers when the system detects the use of a particular browser on one of your monitored hosts. Further assume that when you detect this use, you want to generate an event if the browser version is not the latest.

You could add a host profile qualification to this correlation rule so that the rule triggers only if the **Client** is the **Event Client**, but the **Client Version** is not the latest version.

Example host profile qualification

This host profile qualification constrains a correlation rule so the rule triggers only if the host involved in the discovery event on which the rule is based is running a version of Microsoft Windows.



Syntax for user qualifications

If you are using a connection, intrusion, discovery, or host input event to trigger your correlation rule, you can constrain the rule based on the identity of a user involved in the event. This constraint is called a *user qualification*. For example, you could constrain a correlation rule so that it triggers only when the identity of the source or destination user is one from the sales department.

You cannot add a user qualification to a correlation rule that triggers on a traffic profile change or on the detection of user activity. Also, the system obtains user details through the Firewall Management Center-server connection established in an identity realm. This information may not be available for all users in the database.

When you build a user qualification, first specify the identity you want to use to constrain your correlation rule. The identity you can choose depends on the rule's base event type:

- connection event — Choose **Identity on Initiator** or **Identity on Responder**.
- intrusion event — Choose **Identity on Destination** or **Identity on Source**.
- discovery event — Choose **Identity on Host**.
- host input event — Choose **Identity on Host**.

This table describes how to build a user qualification for a correlation rule.

Table 139: Syntax for user qualifications

If you specify...	Choose an operator, then...
Authentication Protocol	Choose the authentication protocol (or user type) protocol used to detect the user.
Department	Enter a department.

If you specify...	Choose an operator, then...
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants. This field is only present if you have ever configured the Firewall Management Center for multitenancy.
Email	Enter an email address.
First Name	Enter a first name.
Last Name	Enter a last name.
Phone	Enter a telephone number.
Username	Enter a username.

Connection trackers

A *connection tracker* is a correlation rule constraint that

- constrains a correlation rule so that after the rule's initial criteria are met (including host profile and user qualifications), the system begins tracking certain connections,
- generates a correlation event for the rule if the tracked connections meet additional criteria gathered over a time period that you specify, and
- monitors very specific traffic and, when triggered, runs only for a finite, specified time.

Connection tracker event generation

Tip Connection trackers typically monitor very specific traffic and, when triggered, run only for a finite, specified time. Compare connection trackers with traffic profiles, which typically monitor a broad range of network traffic and run persistently.

There are two ways a connection tracker can generate an event.

Connection trackers that fire immediately when conditions are met:

- You can configure a connection tracker so that the correlation rule fires as soon as network traffic meets the tracker's conditions. When this happens, the system stops tracking connections for this connection tracker instance, even if the timeout period has not expired. If the same type of policy violation that triggered the correlation rule occurs again, the system creates a new connection tracker.
- However, if time expires before network traffic meets the conditions in the connection tracker, the system does not generate a correlation event, and also stops tracking connections for that rule instance.

Connection trackers that fire at the end of the timeout period:

- You can configure a connection tracker so that it relies on data collected over the entire timeout period, and therefore cannot fire until the end of the timeout period.

Connection tracker use cases

For example, a connection tracker can serve as a kind of event threshold by generating a correlation event only if a certain type of connection occurs more than a specific number of times within a specific time period. Or, you can generate a correlation event only if the system detects excessive data transfer after an initial connection.

For example, if you configure a connection tracker to fire if you detect fewer than a certain number of bytes being transferred during a certain time period, the system waits until that time period passes and then generates an event if network traffic met that condition.

Add a connection tracker

Connection trackers monitor specific network connections and generate correlation events when predefined conditions are met, enabling proactive network security monitoring and incident response.

Connection trackers work with correlation rules to analyze network traffic patterns and trigger alerts or responses when specific connection behaviors occur within defined time intervals.

Before you begin

- Create a correlation rule based on a connection, intrusion, discovery, user identity, or host input event. You cannot add a connection tracker to a rule based on a VPN troubleshoot event, malware event or traffic profile change.

Follow these steps to add a connection tracker:

Procedure

Step 1 Choose **Policies** > **Correlation**, then click the **Rule management** tab. In the correlation rule editor, click **Edit**.

Step 2 Specify the connections to track; see Syntax for connection trackers, on page 1011.

Step 3 Based on the tracked connections, specify when you want to generate a correlation event; see Syntax for connection tracker events, on page 1014.

Step 4 Specify the interval (in seconds, minutes, or hours) during which the tracker's conditions must be met.

The connection tracker is now configured and will monitor the specified connections according to your defined criteria and time interval, generating correlation events when the conditions are met.

Syntax for connection trackers

This table describes how to build a connection tracker condition that specifies the kind of connections you want to track.

Table 140: Syntax for connection trackers

If you specify...	Choose an operator, then...
Access Control Policy	Choose one or more access control policies that handled the connections you want to track.

If you specify...	Choose an operator, then...
Access Control Rule Action	Choose one or more access control rule actions associated with the access control rule that logged the connections you want to track. Choose **Monitor** to track connections that match the conditions of any Monitor rule, regardless of the rule or default action that later handles the connections.
Access Control Rule Name	Enter all or part of the name of the access control rule that logged the connections you want to track. To track connections that match a Monitor rule, enter the name of the Monitor rule. The system tracks the connections, regardless of the rule or default action that later handles them.
Application Protocol	Choose one or more application protocols.
Application Protocol Category	Choose one or more application protocol categories.
Client	Choose one or more clients.
Client Category	Choose one or more client categories.
Client Version	Enter the version of the client.
Connection Duration	Enter the connection duration, in seconds.
Connection Type	Specify whether you want to trigger the correlation rule based on how the connection information was obtained: • Choose **is** and **Netflow** for connection events generated from exported NetFlow records. • Choose **is not** and **Netflow** for connection events detected by a managed device.
Destination Country or Source Country	Choose one or more countries.
Device	Choose one or more devices whose detected connections you want to track. If you want to track NetFlow connections, choose the devices that process the connection data from exported NetFlow records.
Ingress Interface or Egress Interface	Choose one or more interfaces.
Ingress Security Zone or Egress Security Zone	Choose one or more security zones or tunnel zones.
Initiator IP, Responder IP, or Initiator/Responder IP	Enter a single IP address or address block.
Initiator Bytes, Responder Bytes, or Total Bytes	Enter one of: • the number of bytes transmitted (**Initiator Bytes**) • the number of bytes received (**Responder Bytes**) • the number of bytes both transmitted and received (**Total Bytes**)

If you specify...	Choose an operator, then...
Initiator Packets, Responder Packets, or Total Packets	Enter one of: • the number of packets transmitted (**Initiator Packets**) • the number of packets received (**Responder Packets**) • the number of packets both transmitted and received (**Total Packets**)
Initiator Port/ICMP Type or Responder Port/ICMP Code	Enter the port number or ICMP type for initiator traffic or the port number or ICMP code for responder traffic.
IOC Tag	Choose whether an indication of compromise tag **is** or **is not** set.
NETBIOS Name	Enter the NETBIOS name of the monitored host in the connection.
NetFlow Device	Choose the IP address of the NetFlow exporter you want to track. If you did not add any NetFlow exporters to the network discovery policy, the NetFlow Device drop-down list is blank.
Prefilter Policy	Choose one or more prefilter policies that handled the connections you want to track.
Reason	Choose one or more reasons associated with the connections you want to track.
Security Intelligence Category	Choose one or more Security Intelligence categories associated with the connections you want to track.
TCP Flags	Choose the TCP flag that connections must contain in order to track them. Only connections generated from exported NetFlow records contain TCP flag data.
Transport Protocol	Choose the transport protocol used by the connection.
URL	Enter all or part of the URL visited in the connections you want to track.
URL Category	Choose one or more URL categories for the URL visited in the connections you want to track.
URL Reputation	Choose one or more URL reputation values for the URL visited in the connections you want to track
Username	Enter the username of the user logged into either host in the connections you want to track.
Web Application	Choose one or more web applications.
Web Application Category	Choose one or more web application categories.

Using event data to build a connection tracker

You can often use data from the correlation rule's base event when constructing a connection tracker.

For example, assume your correlation rule triggers when the system detects a new client. When you add a connection tracker to this type of correlation rule, the system automatically populates the tracker with constraints that refer to the base event:

- The **Initiator/Responder IP** is set to the **Event IP Address**.
- The **Client** is set to the **Event Client**.

Tip To track connections for a specific IP address or block of IP addresses, click **switch to manual entry** to manually specify the IP. Click **switch to event fields** to go back to using the IP address in the event.

Syntax for connection tracker events

This table describes how to build a connection tracker condition that specifies when you want to generate a correlation event based on the connections you are tracking.

Table 141: Syntax for connection tracker events

If you specify...	Choose an operator, then enter...
Number of Connections	the total number of connections detected
Number of SSL Encrypted Sessions	the total number of SSL- or TLS-encrypted sessions detected
Total Bytes, Initiator Bytes, or Responder Bytes	one of: • the total bytes transmitted (**Total Bytes**) • the number of bytes transmitted (**Initiator Bytes**) • the number of bytes received (**Responder Bytes**)
Total Packets, Initiator Packets, or Responder Packets	one of: • the total packets transmitted (**Total Packets**) • the number of packets transmitted (**Initiator Packets**) • the number of packets received (**Responder Packets**)
Unique Initiators or Unique Responders	one of: • the number of unique hosts that initiated sessions that were detected (**Unique Initiators**) • the number of unique hosts that responded to connections that were detected (**Unique Responders**)

How excessive connection detection works

Consider a scenario where you archive sensitive files on network 10.1.0.0/16, and where hosts outside this network typically do not initiate connections to hosts inside the network. An occasional connection initiated from outside the network might occur, but you have determined that when four or more connections are initiated within two minutes, there is cause for concern.

The rule shown in the following graphic specifies that when a connection occurs from outside the 10.1.0.0/16 network to inside the network, the system begins tracking connections that meet that criterion. The system

then generates a correlation event if the system detects four connections (including the original connection) within two minutes that match that signature.

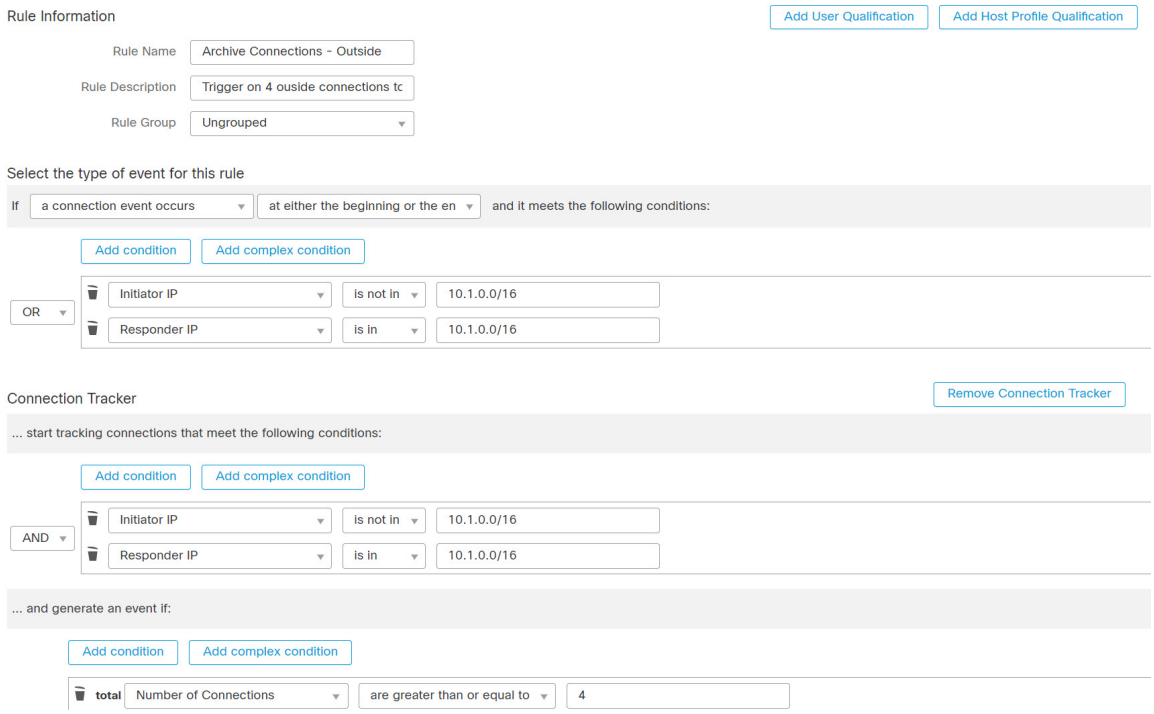
Summary

The key components involved in excessive connection detection are:

- Correlation rule: Defines the conditions and thresholds for connection monitoring
- Connection tracker: Monitors and counts qualifying connections within the specified time window
- External hosts: Sources of connections from outside the protected network
- Internal hosts: Targets of connections within the protected network

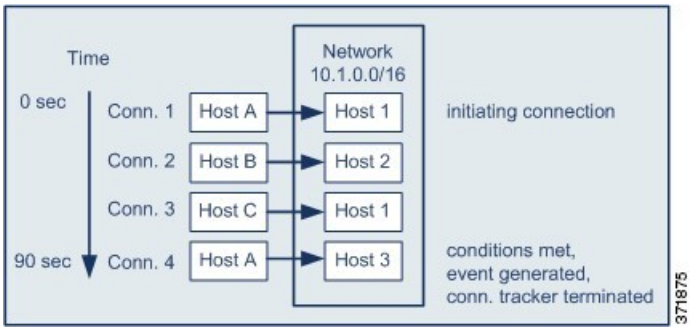
Workflow

Figure 29: Correlation rule configuration



The following diagram shows how network traffic can trigger the above correlation rule.

Figure 30: Network traffic triggering correlation rule



In this example, the system detected a connection that met the basic conditions of the correlation rule, that is, the system detected a connection from a host outside the 10.1.0.0/16 network to a host inside the network. This created a connection tracker.

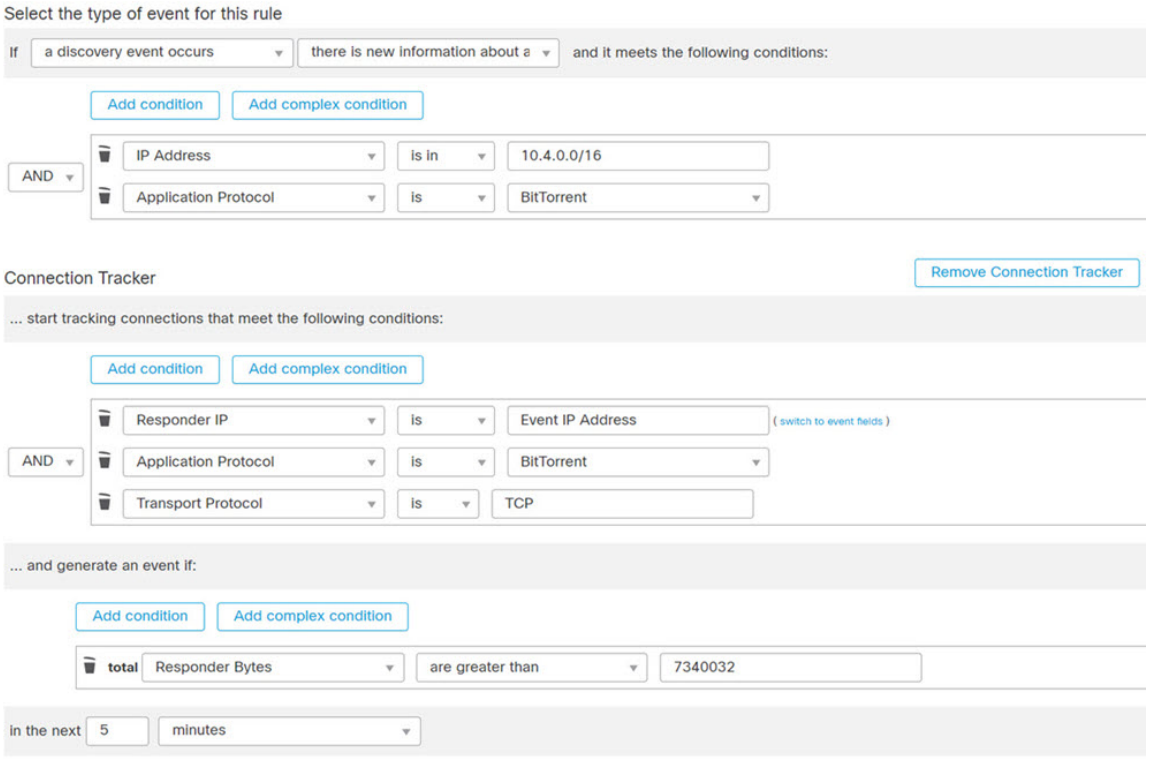
These stages describe how the connection tracker is processed:

1. First, the system starts tracking connections when it detects a connection from Host A outside the network to Host 1 inside the network.
2. The system detects two more connections that match the connection tracker signature: Host B to Host 2 and Host C to Host 1.
3. The system detects a fourth qualifying connection when Host A connects to Host 3 within the two-minute time limit. The rule conditions are met.
4. Finally, the system generates a correlation event and the system stops tracking connections.

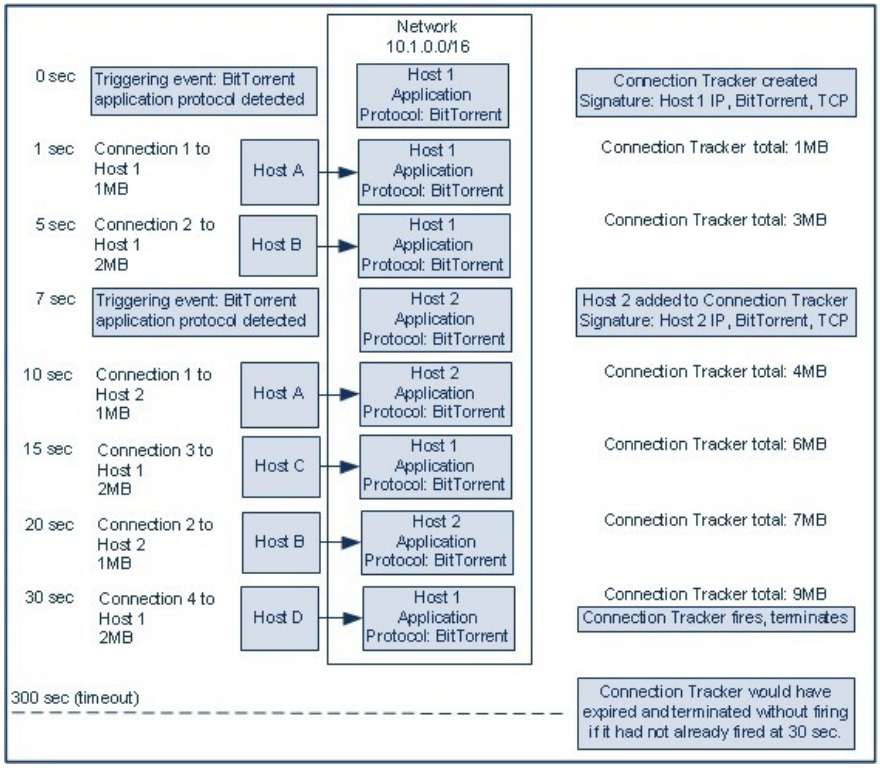
Sample configuration for excessive BitTorrent data transfers

Consider a scenario where you want to generate a correlation event if the system detects excessive BitTorrent data transfers after an initial connection to any host on your monitored network.

This graphic shows a correlation rule that triggers when the system detects the BitTorrent application protocol on your monitored network. The rule has a connection tracker that constrains the rule so that the rule triggers only if hosts on your monitored network (in this example, 10.1.0.0/16) collectively transfer more than 7MB of data (7340032 bytes) via BitTorrent in the five minutes after the initial policy violation.



This diagram shows how network traffic can trigger the correlation rule.



In this example, the system detected the BitTorrent TCP application protocol on two different hosts: Host 1 and Host 2. These two hosts transmitted data via BitTorrent to four other hosts: Host A, Host B, Host C, and Host D.

This connection tracker is processed in the stages:

- First, the system starts tracking connections at the 0-second marker when the system detects the BitTorrent application protocol on Host 1. Note that the connection tracker will expire if the system does not detect 7MB of BitTorrent TCP data being transmitted in the next 5 minutes (by the 300-second marker).
- At 5 seconds, Host 1 has transmitted 3MB of data that matches the signature:
  - 1MB from Host 1 to Host A, at the 1-second marker (1MB total BitTorrent traffic counted towards fulfilling the connection tracker)
  - 2MB from Host 1 to Host B, at the 5-second marker (3MB total)
- At 7 seconds, the system detects the BitTorrent application protocol on Host 2 and starts tracking BitTorrent connections for that host as well.
- At 20 seconds, the system has detected additional data matching the signature being transmitted from both Host 1 and Host 2:
  - 1MB from Host 2 to Host A, at the 10-second marker (4MB total)
  - 2MB from Host 1 to Host C, at the 15-second marker (6MB total)
  - 1MB from Host 2 to Host B, at the 20-second marker (7MB total)
- Although Host 1 and Host 2 have now transmitted a combined 7MB of BitTorrent data, the rule does not trigger because the total number of bytes transmitted must be **more** than 7MB (**Responder Bytes are greater than 7340032**). At this point, if the system were to detect no additional BitTorrent transfers for the remaining 280 seconds in the tracker's timeout period, the tracker would expire and the system would not generate a correlation event.
- However, at 30 seconds, the system detects another BitTorrent transfer, and the rule conditions are met:
  - 2MB from Host 1 to Host D at the 30-second marker (9MB total)
- Finally, the system generates a correlation event. The system also stops tracking connections for this connection tracker instance, even though the 5-minute period has not expired. If the system detects a new connection using the BitTorrent TCP application protocol at this point, it will create a new connection tracker. Note that the system generates the correlation event *after* Host 1 transmits the entire 2MB to Host D, because it does not tally connection data until the session terminates.

Snooze and inactive periods

Snooze and inactive periods are time-based configuration options that

- control when correlation rules can trigger and generate events
- prevent excessive correlation events during specified intervals, and
- allow administrators to suppress rule triggering during planned activities.

Snooze and inactive period characteristics

Snooze periods and inactive periods provide different types of rule control:

- *Snooze periods*: When a correlation rule triggers, a snooze period instructs the system to stop firing that rule for a specified interval, even if the rule is violated again during the interval. When the snooze period has elapsed, the rule can trigger again and start a new snooze period.
- Inactive periods: During inactive periods, the correlation rule will not trigger. You can set up inactive periods to recur daily, weekly, or monthly.

Snooze period usage

For example, you may have a host on your network that should never generate traffic. A simple correlation rule that triggers whenever the system detects a connection involving that host may create multiple correlation events in a short period of time, depending on the network traffic to and from the host. To limit the number of correlation events exposing your policy violation, you can add a snooze period so that the system generates a correlation event only for the first connection (within a time period that you specify) that the system detects involving that host.

Inactive period usage

For example, you might perform a nightly Nmap scan on your internal network to look for host operating system changes. In that case, you could set a daily inactive period on the affected correlation rules for the time and duration of your scan so that those rules do not trigger erroneously.

Correlation rule building mechanics

Correlation rule building mechanics is a process that

- uses three parts for most conditions: a *category*, an *operator*, and a *value*,
- allows categories to depend on whether you are building correlation rule triggers, host profile qualifications, connection trackers, or user qualifications, and
- requires linking multiple conditions with **AND** or **OR** operators.

Correlation rule condition components

The categories you can choose depend on whether you are building correlation rule triggers, a host profile qualification, a connection tracker, OR a user qualification. Within correlation rule triggers, the categories further depend on the base event type for the rule. Some conditions may contain several categories, each of which may have their own operators AND values.

A condition's available operators depend on the category.

The syntax you can use to specify a condition's value depends on the category AND operator. Sometimes you type the value in a text field. Other times, you can choose a value (OR multiple values) from a drop-down list.

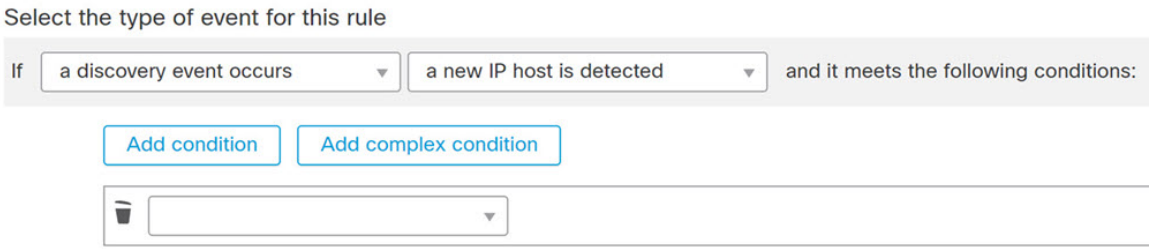
When your construct includes more than one condition, you must link them with an **AND** OR an **OR** operator. Conditions on the same level are evaluated together:

- The **AND** operator requires that all conditions on the level it controls must be met.
- The **OR** operator requires that at least one of the conditions on the level it controls must be met.

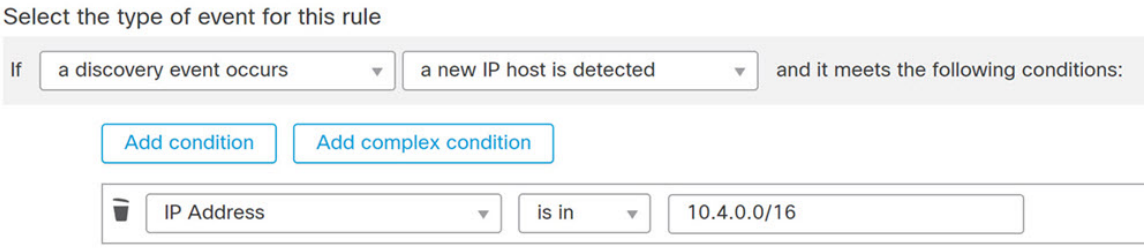
Caution Evaluating complex correlation rules that trigger on frequently occurring events can degrade system performance. For example, a multicondition rule that the system must evaluate against every logged connection can cause resource overload.

Simple correlation rule example

For example, if you want to generate a correlation event every time a new host is detected, you can create a simple rule with no conditions.



If you want to further constrain the rule AND generate an event only if that new host was detected on the 10.4.x.x network, you can add a single condition.



This rule, which detects SSH activity on a nonstandard port on the 10.4.x.x network AND the 192.168.x.x network, has four conditions, with the bottom two constituting a complex condition.

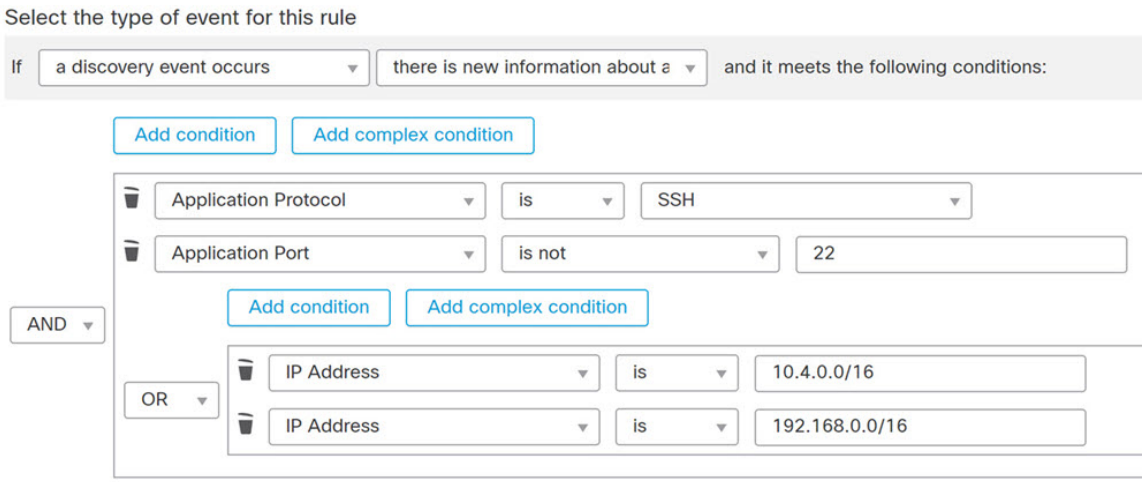


Table 142: Rule evaluation

Where...	Is the condition that states...
A	Application Protocol is SSH
B	Application Port is not 22
C	IP Address is in 10.0.0.0/8
D	IP Address is in 196.168.0.0/16

Logically, the rule is evaluated as follows:

```
(A and B and (C or D))
```

Add AND link conditions in correlation rules

Add conditions to correlation rules to specify the criteria that events must meet for the rule to trigger, AND link multiple conditions using logical operators to create complex matching logic.

Correlation rules use conditions to match events. Simple conditions evaluate a single criterion, while complex conditions combine multiple criteria. Conditions can be linked using AND OR OR operators to create sophisticated matching logic.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click the **Rule management** tab. In the correlation rule editor, add a simple OR complex condition:

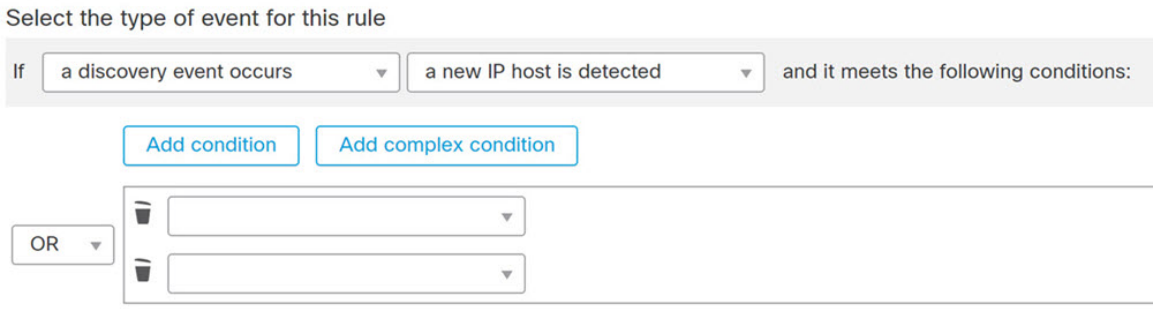
- Simple — Click **Add condition**.
- Complex — Click **Add complex condition**.

Step 2 Link conditions by choosing the **AND** OR **OR** operator from the drop-down list to the left of the conditions.

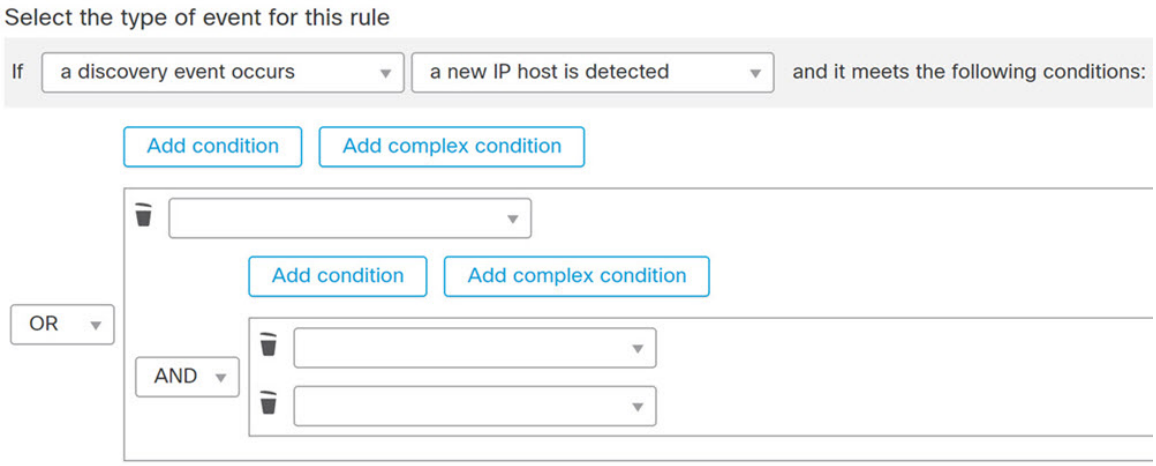
The correlation rule is configured with the specified conditions AND logical operators, ready to match events according to your defined criteria.

Example: simple vs complex conditions

This graphic shows a correlation rule with two simple conditions joined by the **OR** operator.



This graphic shows a correlation rule with one simple condition AND one complex condition, joined by the **OR** operator. The complex condition comprises two simple conditions joined by the **AND** operator.



Use multiple values in correlation rule conditions

This task enables you to select multiple values from drop-down lists when building correlation rule conditions, providing greater flexibility in defining rule criteria.

When you are building a correlation condition, and the condition syntax allows you to pick a value from a drop-down list, you can often use multiple values from the list.

Procedure

Step 1 In the correlation rule editor, build a condition, choosing **is in** or **is not in** as the operator.

Step 2 Click anywhere in the text field or on the **Edit** link.

Step 3 Under **Available**, choose multiple values. You can also click and drag to choose multiple adjacent values.

Step 4 Click the (Right Arrow >) to move the selected entries to **Selected**.

Step 5 Click **OK**.

The correlation rule condition is configured with multiple selected values, allowing the rule to match against any of the specified criteria.

Manage correlation rules

This task allows you to create, edit, and organize correlation rules and groups to effectively manage event correlation across your multidomain deployment.

In a multidomain deployment, the system displays correlation rules and groups created in the current domain, which you can edit. It also displays selected correlation rules and groups from ancestor domains, which you cannot edit. To view and edit correlation rules and groups created in a lower domain, switch to that domain.

Note The system does not display configurations from ancestor domains if the configurations expose information about unrelated domains, including names, managed devices, and so on.

Changes made to rules in active correlation policies take effect immediately.

Before you begin

- If you want to delete a rule, delete it from all correlation policies, as described in Manage correlation policies, on page 992.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Rule Management**.

Step 2 Manage your rules:

- Create — Click **Create Rule**; see Configure correlation rules, on page 993.
- Create Group — Click **Create Group**, enter a name for the group, and click **Save**. To add a rule to a group, edit the rule.
- Edit — Click **Edit** ( ); see Configure correlation rules, on page 993. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Delete Rule or Rule Group— Click **Delete** ( ). Deleting a rule group ungroups the rules. If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

You have successfully managed your correlation rules and groups. Any changes made to rules in active correlation policies will take effect immediately.

Configure correlation response groups

You can create a correlation response group of alerts and remediations, then activate and assign the group to a correlation rule within an active correlation policy. The system launches all the grouped responses when network traffic matches the correlation rule.

When used in an active correlation policy, changes to an active group or any of its grouped responses take affect immediately.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Groups**.

Step 2 Click **Create Group**.

Step 3 Enter a **Name**.

Step 4 Check the **Active** check box if you want to activate the group upon creation.

Deactivated groups do not launch responses.

Step 5 Choose the **Available Responses** to group. then click the Right Arrow > to move them to the **Responses in Group**. To move responses the other way, use the Left Arrow <.

Step 6 Click **Save**.

The correlation response group is created and can be assigned to correlation rules within active correlation policies.

What to do next

- If you did not activate the group upon creation and you want to activate it now, click the slider.

Manage correlation response groups

This task allows you to manage correlation response groups by performing various administrative actions including activation, deactivation, creation, editing, and deletion.

You can delete a response group if it is not used in a correlation policy. Deleting a response group ungroups its responses. You can also temporarily deactivate a response group without deleting it. This leaves the group on the system but does not launch it when policies are violated.

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

Changes made to active, in-use response groups take effect immediately.

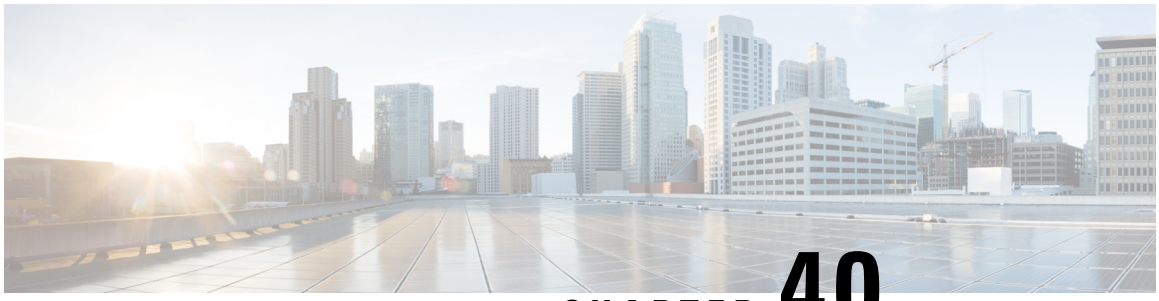
Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Groups**.

Step 2 Manage response groups:

- Activate or Deactivate — Click the slider. If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Create — Click **Create Group**; see Configure correlation response groups, on page 1024.
- Edit — Click **Edit** ( ); see Configure correlation response groups, on page 1024. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Delete — Click **Delete** ( ). If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

The correlation response group is managed according to the selected action, with changes to active groups taking effect immediately.



CHAPTER 40

Traffic profiling

The following topics describe how to configure traffic profiles:

- Introduction to traffic profiles, on page 1027
- Requirements and prerequisites for traffic profiles, on page 1031
- Manage traffic profiles, on page 1031
- Configure traffic profiles, on page 1032

Introduction to traffic profiles

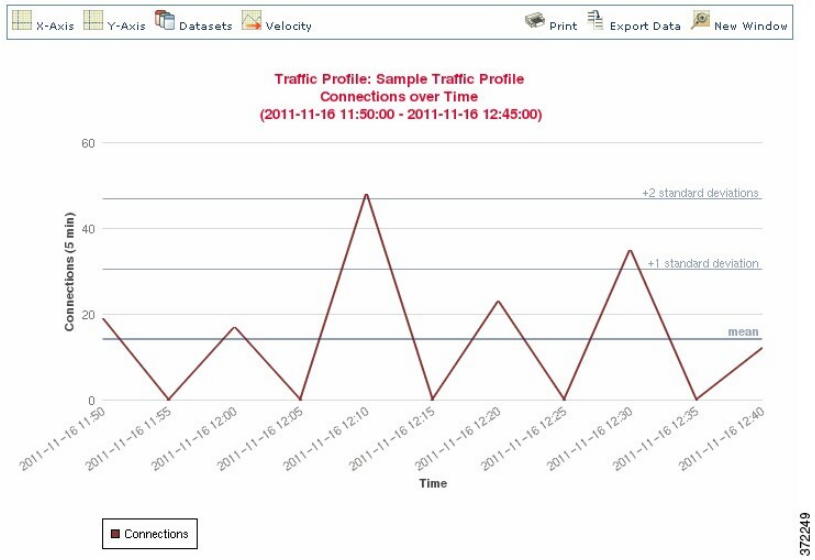
A *traffic profile* is a graph of network traffic based on connection data collected over a profiling time window (PTW). This measurement presumably represents normal network traffic. After the learning period, evaluate new network traffic against your profile to detect abnormalities.

Traffic profile configuration

The default PTW is one week, but you can change it to be as short as an hour or as long as several weeks. By default, traffic profiles generate statistics on connection events generated by the system over five-minute intervals. However, you can increase this sampling rate to as long as an hour.

Tip It is recommended that the PTW include at least 100 data points. Configure your PTW and sampling rate so that your traffic profiles contain enough data to be statistically meaningful.

This graphic shows a traffic profile with a PTW of one day and a sampling rate of five minutes.



You can also set up inactive periods in traffic profile. Traffic profiles collect data during inactive periods, but they do not use this data when calculating profile statistics. Traffic profile graphs plotted over time show inactive periods as a shaded region.

For example, in some network infrastructures, all workstations are backed up at midnight every night. The backup takes about 30 minutes and causes a spike in network traffic. You could configure a recurring inactive period for your traffic profile to coincide with the scheduled backups.

Note The system uses end-of-connection data to create connection graphs and traffic profiles. To use traffic profiles, make sure you log end-of-connection events to the Firewall Management Center database.

Implementing traffic profiles

When you activate a traffic profile, the system collects and evaluates connection data for the learning period (PTW) you configured. After the learning period, the system evaluates correlation rules written against the traffic profile.

For example, you could write a rule that triggers if the amount of data traversing your network (measured in packets, KBytes, or number of connections) suddenly spikes to three standard deviations above the mean amount of traffic, which could indicate an attack or other security policy violation. Then, you could include that rule in a correlation policy to alert you of the traffic spike or to perform a remediation in response.

Targeting traffic profiles

Profile conditions and *host profile qualifications* constrain traffic profiles.

Using profile conditions, you can profile all network traffic, or you can restrict the traffic profile to monitoring a domain, subnets within or across domains, or individual hosts. In a multidomain deployment:

- Leaf-domain administrators can profile network traffic within their leaf domains.
- Higher-level domain administrators can profile traffic within or across domains.

Profile conditions can also constrain traffic profiles based on connection data. For example, you could set the profile conditions so that the traffic profile only profile includes only sessions using a specific port, protocol, or application.

Finally, you can also constrain traffic profiles using information about the tracked hosts. This constraint is called a *host profile qualification*. For example, you could collect connection data only for hosts with high criticality.

Note

- Constraining a traffic profile to a higher-level domain aggregates and profiles the **same** type of traffic in **each** of the descendant leaf domains. The system builds a separate network map for each leaf domain. In a multidomain deployment, profiling traffic across domains can have unexpected results.
- The traffic profiling graph displays traffic segment data generated based on configured constraints and intervals. These traffic segments may be purged according to the database retention policies. As a result, the traffic profiling graph may display mismatching data against the metrics that are calculated from history of data stored.

Related Topics

Correlation policies and rules, on page 989

Traffic profile conditions

You can create simple traffic profile conditions and host profile qualifications. You can also create more elaborate constructs by combining and nesting conditions.

Condition structure and components

Conditions have three parts: a category, an operator, and a value:

- The available categories depend on whether you are building traffic profile conditions or a host profile qualification.
- The operators available for use depend on your category that you have chosen.
- The syntax that is used to specify a condition's value depends on the category and operator. In some cases, you must enter the value in a text field. In other cases, you can select one or more values from a drop-down list.

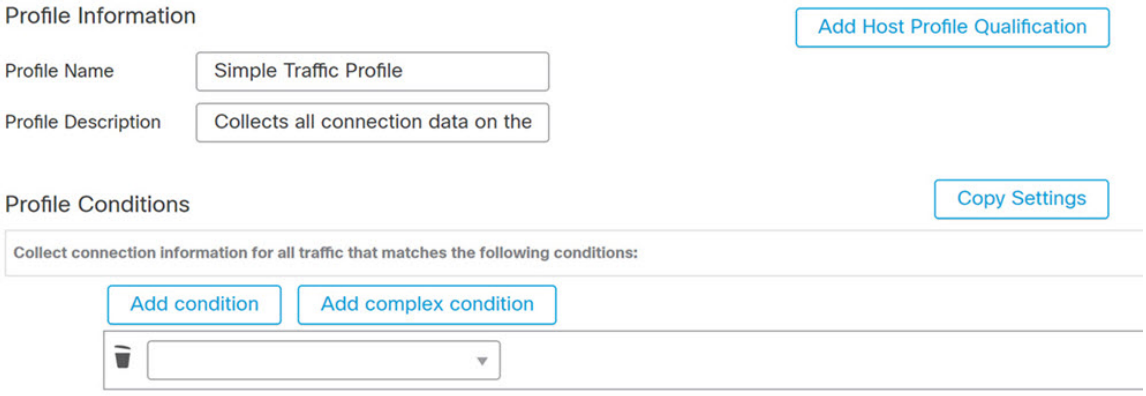
For a host profile qualification, you must also specify whether you are constraining the traffic profile using information data about the initiating OR responding hosts.

When your construct includes more than one condition, you must link them with an **AND** OR an **OR** operator. Conditions on the same level are evaluated together:

- The **AND** operator requires that all conditions on the level it controls must be met.
- The **OR** operator requires that at least one of the conditions on the level it controls must be met.

Unconstrained traffic profile

If you want to create a traffic profile that collects data for your entire monitored network segment, you can create a very simple profile with no conditions, as shown in this graphic.



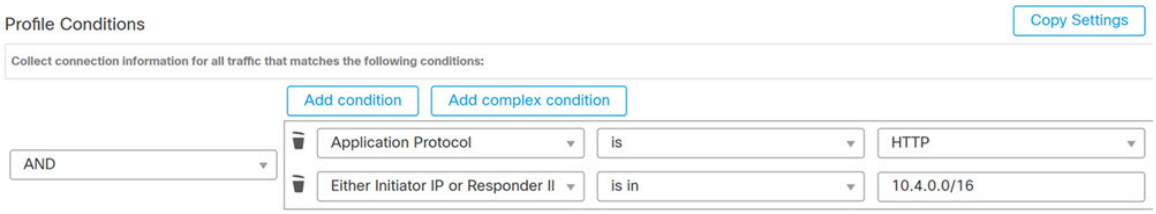
Simple traffic profile

If you wanted to constrain the profile and collect data only for a subnet, you can add a single condition, as shown in the following graphic.

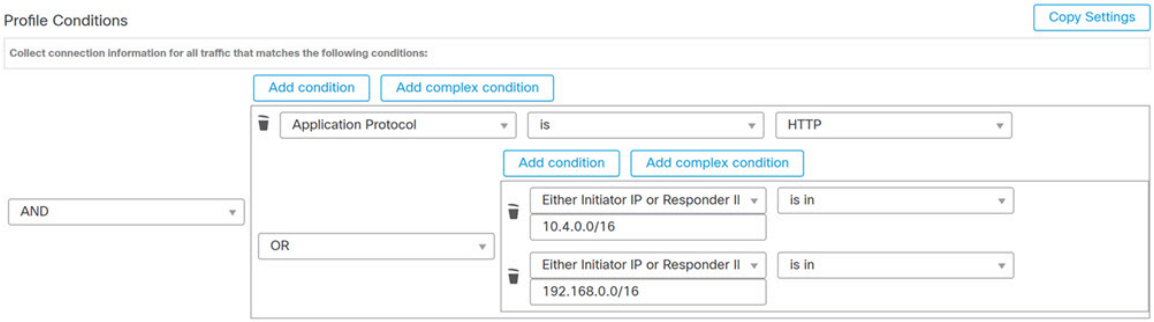


Complex traffic profile

This traffic profile contains two conditions linked by **AND**. This means that the traffic profile collects connection data only if both conditions are true. In this example, it collects HTTP connections for all hosts with IP addresses in a specific subnet.



In contrast, this traffic profile, which collects connection data for HTTP activity in either of two subnets, has three conditions, with the last constituting a complex condition.



Logically, the above traffic profile is evaluated as:

```
(A and (B or C))
```

Where...	Is the condition that states...
A	Application Protocol Name is HTTP
B	IP Address is in 10.4.0.0/16
C	IP Address is in 192.168.0.0/16

Requirements and prerequisites for traffic profiles

Category	Requirement
Model support	Any
Supported domains	Any
User roles	• Admin • Discovery Admin

Manage traffic profiles

Only if you create rules for active, complete traffic profiles, the system can trigger a correlation policy violation. Each traffic profile displays a slider indicating whether it is active and collecting data. You can view the status of the traffic profile's learning period in the progress bar.

In a multidomain deployment, you can edit traffic profiles created in the current domain. The system also displays selected traffic profiles from ancestor domains, which you cannot edit. To edit a traffic profile from a lower domain, switch to that domain.

Note The system does not display traffic profiles from ancestor domains if the profiles' conditions expose information about unrelated domains, including names, managed devices, and so on.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Traffic Profiles**.

Step 2 Manage your traffic profiles:

- Activate/Deactivate — To activate or deactivate a traffic profile, click the slider. Deactivating a traffic profile deletes its associated data. If you reactivate the profile, you must wait the length of its PTW before rules written against it will trigger.
- Create — To create a new traffic profile, click **New Profile** and use the instructions in Configure traffic profiles, on page 1032. You can also click **Copy** ( ) to edit a copy of an existing traffic profile.
- Delete — To delete a traffic profile, click **Delete** ( ), then confirm your choice.
- Edit — To modify an existing traffic profile, click **Edit** ( ) and use the instructions in Configure traffic profiles, on page 1032. If a traffic profile is active you can only change its name and description.
- Graph — To view the traffic profile as a graph, click **Graph** ( ). In a multidomain deployment, you cannot view the graph for a traffic profile that belongs to an ancestor domain if the graph exposes information about unrelated domains.

Configure traffic profiles

When you constrain a traffic profile to a higher-level domain, you aggregate and profiles the **same** type of traffic in **each** descendant leaf domains. The system builds a separate network map for each leaf domain. In a multidomain deployment, if you profile traffic across domains, you might see unexpected results.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Traffic Profiles**.

Step 2 Click **New Profile**.

Step 3 Enter a **Profile Name**, and optionally, a **Profile Description**.

Step 4 You can constrain the traffic profile using any of these options:

- Copy Settings — To copy settings from an existing traffic profile, click **Copy Settings**, choose the traffic profile you want to use, and click **Load**.
- Profile Conditions — To constrain the traffic profile using information from tracked connections, proceed as described in Add traffic profile conditions, on page 1033.
- Host Profile Qualification — To constrain the traffic profile using information from tracked hosts, proceed as described in Add host profile qualifications to a traffic profile, on page 1034.

- Profiling Time Window (PTW) — To change the **Profiling Time Window**, enter a time unit, then choose **hour(s)**, **day(s)**, or **week(s)**.
- Sampling Rate — Choose a **Sampling Rate**, in minutes.
- Inactive Period — Click **Add Inactive Period** and use the drop-down lists to specify when and how often you want the traffic profile to remain inactive. Inactive traffic profiles do not trigger correlation rules. Traffic profiles do not include data from inactive periods in profile statistics.

Step 5 Save the traffic profile:

- To save the profile and start collecting data immediately, click **Save & Activate**.
- To save the profile without activating it, click **Save**.

Add traffic profile conditions

Procedure

Step 1 In the traffic profile editor, under Profile Conditions, click **Add condition** or **Add complex condition** for each condition you want to add.

Conditions on the same level are evaluated together.

- To require that all conditions on the level that the operator controls are met, choose **AND**.
- To require that only one of the conditions on the level that the operator controls is met, choose **OR**.

Step 2 Specify a category, operator, and value for each condition as described in Syntax for traffic profile conditions, on page 1034 and Traffic profile conditions, on page 1029.

If you choose **is in** or **is not in** as the operator, you can select multiple values in a single condition as described in Use multiple values in a traffic profile condition, on page 1037.

When the category represents an IP address, choosing **is in** or **is not in** as the operator allows you to specify whether the IP address *is in* or *is not in* a range of IP addresses.

The following traffic profile collects information on a specific subnet. The category of the condition is **Initiator/Responder IP**, the operator is **is in**, and the value is `10.4.0.0/16`.



Related Topics

IP address conventions, on page 27

Add host profile qualifications to a traffic profile

Use host profile qualifications to filter connection data based on host characteristics such as operating system, applications, or other host attributes.

Procedure

Step 1 In the traffic profile editor, click **Add Host Profile Qualification**.

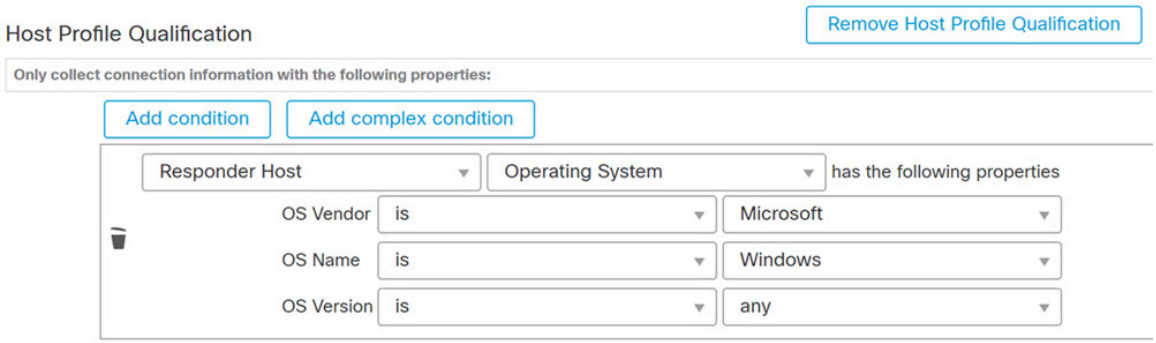
Step 2 Under Host Profile Qualification, click **Add condition** or **Add complex condition** for each condition you want to add. Conditions on the same level are evaluated together.

- To require that all conditions on the level that the operator controls are met, choose **AND**.
- To require that only one of the conditions on the level that the operator controls is met, choose **OR**.

Step 3 Specify a host type, category, operator, and value for each condition as described in Syntax for host profile qualifications in a traffic profile, on page 1035 and Traffic profile conditions, on page 1029.

If you choose **is in** or **is not in** as the operator, you can select multiple values in a single condition as described in Use multiple values in a traffic profile condition, on page 1037.

The following host profile qualification constrains a traffic profile such that it collects connection data only if the responding host in the detected connection is running a version of Microsoft Windows.



Syntax for traffic profile conditions

The following table describes how to build a traffic profile condition. The connection data available to build a traffic profile depends on several factors, including traffic characteristics and detection method.

Table 143: Syntax for traffic profile conditions

If you choose...	Choose an operator, then...
Application Protocol	Choose one or more application protocols.
Application Protocol Category	Choose one or more application protocol categories.

If you choose...	Choose an operator, then...
Client	Choose one or more clients.
Client Category	Choose one or more client categories.
Connection Type	Choose whether the profile uses connection data from traffic monitored by managed devices or from exported NetFlow records. If you do not specify a connection type, the traffic profile includes both.
Destination Country or Source Country	Choose one or more countries.
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants.
Initiator IP, Responder IP, or Initiator/Responder IP	Enter an IP address or range of IP addresses. The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.
NetFlow Device	Choose the NetFlow exporter whose data you want to use to create the traffic profile.
Responder Port/ICMP Code	Enter the port number or ICMP code.
Security Intelligence Category	Choose one or more a Security Intelligence categories. To use a Security Intelligence category for a traffic profile condition, that category must be set to **Monitor** instead of **Block** in your access control policy.
SSL Encrypted Session	Choose **Successfully Decrypted**.
Transport Protocol	Enter `TCP` or `UDP` as the transport protocol.
Web Application	Choose one or more web applications.
Web Application Category	Choose one or more web application categories.

Related Topics

Requirements for populating connection event fields, on page 776
IP address conventions, on page 27

Syntax for host profile qualifications in a traffic profile

Build a host profile qualification condition by first selecting the host to constrain your traffic profile. You can choose either **Responder Host** or **Initiator Host**. After you choose the host role, continue building your host profile qualification condition.

Although you can add hosts to the network map using NetFlow records, the available information about these hosts is limited. For example, there is no operating system data available for these hosts, unless you provide it using the host input feature. In addition, if your traffic profile uses connection data from exported NetFlow records, keep in mind that NetFlow records do not contain information about which host in the connection is the initiator and which is the responder. When the system processes NetFlow records, it uses an algorithm to determine this information based on the ports each host is using, and whether those ports are well-known.

To match against *implied* or generic clients, create a host profile qualification based on the application protocol used by the server responding to the client. When the client list on a host that acts as the initiator or source of a connection includes an application protocol name followed by **client**, that client may actually be an implied client. In other words, the system reports that client based on server response traffic that uses the application protocol for that client, not on detected client traffic.

For example, if the system reports **HTTPS client** as a client on a host, create a host profile qualification for **Responder Host** where **Application Protocol** is set to **HTTPS**, because HTTPS client is reported as a generic client based on the HTTPS server response traffic sent by the responder or destination host.

Table 144: Syntax for Host Profile Qualifications

If you choose...	Choose an operator, then...
Application Protocol > Application Protocol	Choose one or more application protocols.
Application Protocol > Application Port	Enter the application protocol port number.
Application Protocol > Protocol	Choose the protocol.
Application Protocol Category	Choose one or more application protocol categories.
Client > Client	Choose one or more clients.
Client > Client Version	Enter the client version.
Client Category	Choose one or more client categories.
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants.
Hardware	Enter a mobile device hardware model. For example, to match all Apple iPhones, enter `iPhone`.
Host Criticality	Choose a host criticality.
Host Type	Choose one or more host types. You can choose between a normal host or one of several types of network device.
IOC Tag	Choose one or more IOC tags.
Jailbroken	Choose **Yes** to indicate that the host in the event is a jailbroken mobile device or **No** to indicate that it is not.
MAC Address > MAC Address	Enter all or part of the MAC address of the host.
MAC Address > MAC Type	Choose whether the MAC type is **ARP/DHCP Detected**, that is, whether: • The system positively identified the MAC address as belonging to the host (**is ARP/DHCP Detected**) • The system is seeing many hosts with that MAC address because, for example, there is a router between the device and the host (**is not ARP/DHCP Detected**) • The MAC type is irrelevant (**is any**)

If you choose...	Choose an operator, then...
MAC Vendor	Enter all or part of the MAC vendor of hardware used by the host.
Mobile	Choose **Yes** to indicate that the host in the event is a mobile device or **No** to indicate that it is not.
NETBIOS Name	Enter the NetBIOS name of the host.
Network Protocol	Enter the network protocol number as listed in http://www.iana.org/assignments/ethernet-numbers.
Operating System > OS Vendor	Choose one or more operating system vendor names.
Operating System > OS Name	Choose one or more operating system names.
Operating System > OS Version	Choose one or more operating system versions.
Transport Protocol	Enter the name or number of the transport protocol as listed in http://www.iana.org/assignments/protocol-numbers.
VLAN ID	Enter the VLAN ID number of the host. The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal VLAN tags to constrain this configuration can have unexpected results.
Web Application	Choose one or more web applications.
Web Application Category	Choose one or more web application categories.
any available host attribute, including the default compliance allow list host attribute	Specify the appropriate value, which depends on the type of host attribute you choose: • If the host attribute type is Integer, enter an integer value in the range defined for the attribute. • If the host attribute type is Text, enter a text value. • If the host attribute type is List, choose a valid list string. • If the host attribute type is URL, enter a URL value.

Use multiple values in a traffic profile condition

When you build a condition and the syntax allows selection from a drop-down list, you can use multiple values.

To add a host profile qualification for traffic profiles where the host runs any UNIX version, use this procedure instead of creating several conditions with the OR operator.

Procedure

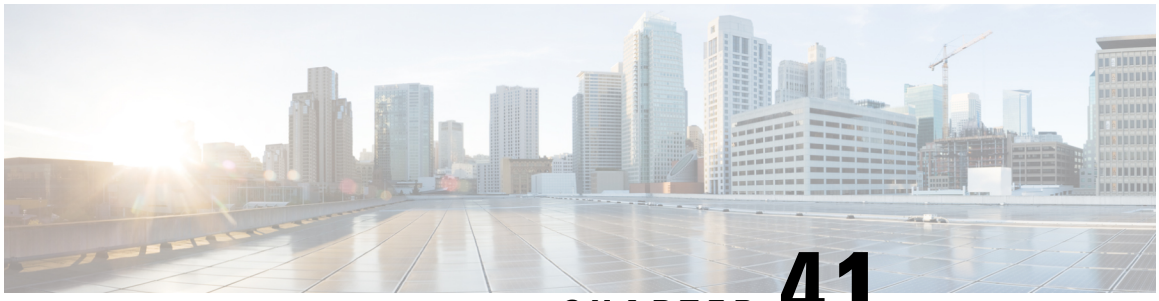
Step 1 While building a traffic profile or host profile qualification condition, choose **is in** or **is not in** as the operator. The drop-down list changes to a text field.

Step 2 Click anywhere in the text field or on the **Edit** link.

Step 3 Under **Available**, choose multiple values.

Step 4 Click the right arrow to move the selected entries to **Selected**.

Step 5 Click **OK**.



CHAPTER 41

Remediations

- Requirements and prerequisites for remediations, on page 1039
- Introduction to remediations, on page 1039
- Manage remediation modules, on page 1050
- Manage remediation instances, on page 1050
- Manage instances for a single remediation module, on page 1051

Requirements and prerequisites for remediations

This section lists the requirements and prerequisites before you implement security remediations.

Requirement	Supported values
Model support	Any
Supported domains	Any
User roles	• Admin • Discovery Admin

Introduction to remediations

A remediation is a program that

- launches in response to a correlation policy violation,
- generates a remediation status event when it runs, and
- includes details such as the remediation name, the correlation policy and rule that triggered it, and the exit status message.

Supported remediation modules

The system supports several remediation modules:

- **Cisco ISE Adaptive Network Control (ANC)**: Applies or clears ANC policies configured in Cisco ISE when a correlation policy violation occurs.
- **Cisco IOS Null Route**: Blocks traffic sent to a host or network involved in a correlation policy violation (requires Cisco IOS Version 12.0 or higher).
- **Nmap Scanning**: Scans hosts to determine the operating systems and servers running on them.
- **Set Attribute Value**: Sets a host attribute on a host involved in a correlation policy violation.



Note You can install custom modules to perform additional tasks. Refer to the *Firepower System Remediation API Guide*.

Implementing remediations

To implement a remediation:

- Create at least one instance for the remediation module you choose. Multiple instances can be created per module, each with unique configurations. For example, configure multiple instances of the Cisco IOS Null Route module to communicate with different routers.
- Add multiple remediations to each instance to specify the actions to be performed when a policy is violated.
- Associate remediations with rules in correlation policies so that the system launches the appropriate remediation in response to a violation.

Remediations and multitenancy

In a multidomain deployment, you can install custom remediation modules at any domain level. The system-provided modules belong to the Global domain.

Though you cannot add a remediation to an instance created in an ancestor domain, you can create a similarly configured instance in the current domain and add remediations to that instance. You can also use remediations created in ancestor domains as correlation responses.

Related Topics

Configuring external alerts with alert responses, on page 561
Nmap scanning
Add responses to rules and allow lists, on page 991

Cisco ISE EPS remediations

If you have Endpoint Protection Service (EPS) enabled and configured in your ISE deployment, you can configure your Firewall Management Center to launch remediations using ISE. When fully configured, ISE EPS remediations run the following Mitigation Actions on the source or destination host involved in a correlation policy violation:

- **`quarantine`**—Limits or denies an endpoint's access the network.
- **`unquarantine`**—Reverses an endpoint's quarantine status and allows full access to the network.

- **shutdown**—Deactivates an endpoint's network attached system (NAS) port to disconnect it from the network.

Exempt specific IP addresses from ISE EPS remediation if needed.

Note Your ISE version and configuration impact how you can use ISE. For example, you cannot use ISE-PIC to perform ISE EPS remediations. For more information, see the *User Control with ISE/ISE-PIC* chapter in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

For more information about ISE EPS actions, see the *Cisco Identity Services Engine User Guide*.

Configure ISE EPS remediations

You can respond to correlation policy violations by running ISE EPS remediations on the source or destination host.

Note ISE PIC cannot perform ISE EPS remediations.

Before you begin

- Configure EPS operations on your ISE server.
- See the chapter on configuring ISE/PIC in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Add a pxGrid mitigation instance as described in Add an ISE EPS instance, on page 1041.

Step 3 Add one or more ISE EPS remediations as described in Add ISE EPS remediations, on page 1042.

What to do next

Assign remediations as responses to correlation policy violations as described in Add responses to rules and allow lists, on page 991.

Add an ISE EPS instance

Add an ISE EPS instances so you can group individual remediations by logging type.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 From the **Add a New Instance** list, choose **pxGrid Mitigation(v1.0)** as the module type and click **Add**.

Step 3 Enter an **Instance Name** and **Description**.

Step 4 Set **Enable Logging** option to enable or disable system logging.

Step 5 Click **Create**.

What to do next

Create an ISE EPS remediation as described in Add attribute value remediations, on page 1049.

Related Topics

IP address conventions, on page 27

Add ISE EPS remediations

Create one or more ISE EPS remediations within an instance to run **Mitigation Actions** on the source or destination host involved in a correlation policy violation.

In a multidomain deployment, you cannot add a remediation to an instance created in an ancestor domain.

Before you begin

Create an ISE EPS instance as described in Add an ISE EPS instance, on page 1041.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Next to the instance where you want to add the remediation, click **View** ( ).

Step 3 In the **Configured Remediations** section, choose the **Mitigate Destination** or **Mitigate Source** and click **Add**.

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

Step 4 Enter a **Remediation Name** and **Description**.

Step 5 Choose a **Mitigation Action**: `quarantine`, `unquarantine`, or `shutdown`.

Step 6 (Optional) To exempt IP addresses or ranges from remediation, enter them into the **Allow List** box.

Step 7 Click **Create**, then click **Done**.

What to do next

Assign remediations as responses to correlation policy violations; see Add responses to rules and allow lists, on page 991.

Cisco IOS null route remediations

The Cisco IOS Null Route remediation module allows you to block an IP address or range of addresses using Cisco's "null route" command. This drops all traffic sent to a host or network by routing it to the router's NULL interface. This does not block traffic sent from the violating host or network.

Important considerations

Note Do not use a destination-based remediation as a response to a correlation rule that is based on a discovery or host input event. These events are associated with source hosts.

Caution When a Cisco IOS remediation is activated, there is no timeout period. To unblock the IP address or network, you must manually clear the routing change from the router.

Configure remediations for Cisco IOS routers

This task configures remediation responses on Cisco IOS routers that automatically block traffic when correlation policies are violated.

Cisco IOS remediations provide automated responses to security events detected by correlation rules. You can configure different types of blocking remediations based on your security requirements.

Note Do not use a destination-based remediation as a response to a correlation rule that is based on a discovery or host input event. These events are associated with source hosts.

Caution When a Cisco IOS remediation is activated, there is no timeout period. To unblock the IP address or network, you must manually clear the routing change from the router.

Before you begin

- Confirm that your Cisco router is running Cisco IOS 12.0 or higher.
- Confirm that you have level 15 administrative access to the router.

Procedure

Step 1 Enable Telnet on your Cisco router. Refer to the documentation provided with your router or IOS software for instructions.

Step 2 Add a Cisco IOS Null Route instance for each Cisco IOS router you want to use in Firewall Management Center. For more information, refer to Add a Cisco IOS instance, on page 1044.

Step 3 Create remediations for each instance based on the response you want the router to perform when correlation policies are violated:

- Add Cisco IOS block destination remediations, on page 1045
- Add Cisco IOS block destination network remediations, on page 1045
- Add Cisco IOS block source remediations, on page 1046
- Add Cisco IOS block source network remediations, on page 1047

What to do next

Assign remediations as responses to correlation policy violations; see Add responses to rules and allow lists, on page 991.

Add a Cisco IOS instance

Add a Cisco IOS instance to enable remediation actions on routers through the management system.

If you have multiple routers where you want to send remediations, create a separate instance for each router.

Before you begin

Configure Telnet access on the Cisco IOS router as described in the documentation provided with the router or IOS software.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 From the **Add a New Instance** list, choose **Cisco IOS Null Route** and click **Add**.

Step 3 Enter an **Instance Name** and **Description**.

Step 4 In the **Router IP** field, enter the IP address of the Cisco IOS router you want to use for the remediation.

Step 5 In the **Username** field, enter the Telnet user name for the router. This user must have level 15 administrative access on the router.

Step 6 In the **Connection Password** fields, enter the Telnet user's user password.

Step 7 In the **Enable Password** fields, enter the Telnet user's enable password. This is the password used to enter privileged mode on the router.

Step 8 In the **Allow List** field, enter IP addresses or ranges that you want to exempt from the remediation, one per line.

Note
The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.

Step 9 Click **Create**.

What to do next

Add specific remediations to be used by correlation policies as described in Add Cisco IOS block destination remediations, on page 1045, Add Cisco IOS block destination network remediations, on page 1045, Add Cisco IOS block source remediations, on page 1046, and Add Cisco IOS block source network remediations, on page 1047.

Related Topics

IP address conventions, on page 27

Add Cisco IOS block destination remediations

The Cisco IOS Block Destination remediation blocks traffic sent from the router to the destination host involved in a correlation policy violation. Do not use this remediation as a response to a correlation rule that is based on a discovery or host input event, as these events are associated with source hosts.

In a multidomain deployment, you cannot add a remediation to an instance created in an ancestor domain.

Before you begin

Add a Cisco IOS instance as described in Add a Cisco IOS instance, on page 1044.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Next to the instance where you want to add the remediation, click **View** ( ).

Step 3 In the **Configured Remediations** section, choose **Block Destination** and click **Add**.

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

Step 4 Enter a **Remediation Name** and **Description**.

Step 5 Click **Create**, then click **Done**.

What to do next

Assign remediations as responses to correlation policy violations; see Add responses to rules and allow lists, on page 991.

Add Cisco IOS block destination network remediations

The Cisco IOS block destination network remediation blocks traffic sent from the router to the network of the destination host involved in a correlation policy violation.

Do not use this remediation as a response to a correlation rule that is based on a discovery or host input event, as these events are associated with source hosts.

In a multidomain deployment, you cannot add a remediation to an instance created in an ancestor domain.

Before you begin

Add a Cisco IOS instance as described in Add a Cisco IOS instance, on page 1044.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Next to the instance where you want to add the remediation, click **View** ( ).

Step 3 In the **Configured Remediations** section, choose **Block Destination Network** and click **Add**.

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

Step 4 Enter a **Remediation Name** and **Description**.

Step 5 In the **Netmask** field, enter the subnet mask or use CIDR notation to describe the network that you want to block traffic to.

For example, to block traffic to an entire Class C network when a single host triggers a rule (this is not recommended), use `255.255.255.0` or `24` as the netmask.

As another example, to block traffic to 30 addresses that include the triggering IP address, specify `255.255.255.224` or `27` as the netmask. In this case, if the IP address `10.1.1.15` triggers the remediation, all IP addresses between `10.1.1.1` and `10.1.1.30` are blocked. To block only the triggering IP address, leave the field blank, enter `32`, or enter `255.255.255.255`.

Step 6 Click **Create**, then click **Done**.

What to do next

Assign remediations as responses to correlation policy violations; see Add responses to rules and allow lists, on page 991.

Related Topics

IP address conventions, on page 27

Add Cisco IOS block source remediations

The Cisco IOS Block Source remediation blocks traffic sent from the router to the source host involved in a correlation policy violation.

In a multidomain deployment, you cannot add a remediation to an instance created in an ancestor domain.

Before you begin

Add a Cisco IOS instance as described in Add a Cisco IOS instance, on page 1044.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Next to the instance where you want to add the remediation, click **View** ( ).

Step 3 In the **Configured Remediations** section, choose **Block Source** and click **Add**.

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

Step 4 Enter a **Remediation Name** and **Description**.

Step 5 Click **Create**, then click **Done**.

What to do next

Assign remediations as responses to correlation policy violations; see Add responses to rules and allow lists, on page 991.

Add Cisco IOS block source network remediations

The Cisco IOS block source network remediation blocks traffic from the router to the network of the source host involved in a correlation policy violation.

In a multidomain deployment, you cannot add a remediation to an instance created in an ancestor domain.

Before you begin

Add a Cisco IOS instance as described in Add a Cisco IOS instance, on page 1044.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Next to the instance where you want to add the remediation, click **View** ( ).

Step 3 In the **Configured Remediations** section, choose **Block Source Network** and click **Add**.

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

Step 4 Enter a **Remediation Name** and **Description**.

Step 5 In the **Netmask** field, enter the subnet mask or CIDR notation that describes the network that you want to block traffic to.

For example, to block traffic to an entire Class C network when a single host triggers a rule (this is not recommended), use `255.255.255.0` or `24` as the netmask.

As another example, to block traffic to 30 addresses that include the triggering IP address, specify `255.255.255.224` or `27` as the netmask. In this case, if the IP address `10.1.1.15` triggers the remediation, all IP addresses between `10.1.1.1` and `10.1.1.30` are blocked. To block only the triggering IP address, leave the field blank, enter `32`, or enter `255.255.255.255`.

Step 6 Click **Create**, then click **Done**.

What to do next

Assign remediations as responses to correlation policy violations; see Add responses to rules and allow lists, on page 991.

Related Topics

IP address conventions, on page 27

Nmap scan remediations

An Nmap scan remediation is a security response system that integrates with Nmap™, an open source active scanner for network exploration and security auditing, to trigger network scans when correlation policy violations occur.

Additional information

For more information about Nmap scanning, see Nmap scanning.

Set attribute value remediations

The attribute value remediation is a response mechanism that sets a host attribute value on the host where a triggering event occurred when responding to a correlation policy violation.

Text host attributes

For text host attributes, you can use the description from the event as the attribute value.

Configure set attribute remediations

Set attribute remediations allow you to modify specific attributes on managed devices to bring them into compliance with security policies or to address identified vulnerabilities.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Create a set attribute instance as described in Add a set attribute value instance, on page 1048.

Step 3 Add a set attribute remediation as described in Add attribute value remediations, on page 1049.

What to do next

Assign remediations as responses to correlation policy violations; see Add responses to rules and allow lists, on page 991.

Related Topics

Predefined Host Attributes, on page 888

User-Defined Host Attributes, on page 889

Add a set attribute value instance

Set an attribute value instance to create customized remediation actions.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 From the **Add a New Instance** list, choose **Set Attribute Value** and click **Add**.

Step 3 Enter an **Instance Name** and **Description**.

Step 4 Click **Create**.

What to do next

Create a set attribute remediation as described in Add attribute value remediations, on page 1049.

Add attribute value remediations

The set attribute value remediation sets a host attribute on a host involved in a correlation policy violation. Create a remediation for each attribute value you want to set. For text attributes, you can use the description from the triggering event as the attribute value.

In a multidomain deployment, you cannot add a remediation to an instance created in an ancestor domain.

Before you begin

Create a set attribute instance as described in Add a set attribute value instance, on page 1048.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Click **View** ( ) next to the instance where you want to add the remediation.

Step 3 In the **Configured Remediations** section, choose **Set Attribute Value** and click **Add**.

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

Step 4 Enter a **Remediation Name** and **Description**.

Step 5 To use this remediation in response to an event with source and destination data, choose an **Update Which Host(s) From Event** option.

Step 6 For text attributes, specify whether you want to **Use Description From Event For Attribute Value**:

- To use the description from the event as the attribute value, click **On** and enter the **Attribute Value** you want to set.
- To use the **Attribute Value** setting for the remediation as the attribute value, click **Off**.

Step 7 Click **Create**, then click **Done**.

What to do next

Assign remediations as responses to correlation policy violations; see Add responses to rules and allow lists, on page 991.

Manage remediation modules

Manage remediation modules to control automated security responses and network remediation capabilities.

In a multidomain deployment, the system displays remediation modules installed in the current domain, which you can delete. It also displays modules installed in ancestor domains, which you cannot delete. To manage remediation modules in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Policies** > **Actions** > **Modules**.

Step 2 Manage your remediation modules:

- Configure — To view the Module Detail page for a module and configure its instances and remediations, click **View** ( ). In a multidomain deployment, you cannot use the Module Detail page to add, delete, or edit instances in the current domain for a module installed in an ancestor domain. Instead, use the Instances page (**Policies** > **Actions** > **Instances**); see Manage remediation instances, on page 1050.
- Delete — To delete a custom module that is not in use, click **Delete** ( ). You cannot delete system-provided modules.
- Install — To install a custom module, click **Choose File**, browse to the module, and click **Install**. For more information, see the *Firepower System Remediation API Guide*.

Manage remediation instances

This task allows you to add, configure, or delete remediation instances that control how the system responds to security events across all remediation modules.

The Instances page lists all configured instances for all remediation modules.

In a multidomain deployment, the system displays remediation instances created in the current domain, which you can edit. It also displays instances created in ancestor domains, which you cannot edit. To edit remediation instances in a lower domain, switch to that domain.

Although you cannot add a remediation to an instance created in an ancestor domain, you can create a similarly configured instance in the current domain and add remediations to that instance. You can also use remediations created in ancestor domains as correlation responses.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Manage your remediation instances:

- Add—To add an instance, choose the remediation module for which you want to add an instance and click **Add**. For system-provided modules, see:
  - Add an ISE EPS instance, on page 1041
  - Add a Cisco IOS instance, on page 1044
  - *Cisco Secure Firewall Management Center Device Configuration Guide*
  - Add a set attribute value instance, on page 1048

  For help adding a custom module, see the documentation for that module, if available.
- Configure—To configure instance details and add remediations to the instance, click **View** ( ).
- Delete—To delete an instance that is not in use, click **Delete** ( ).

Manage instances for a single remediation module

Use the Module Detail page to manage all instances and remediations configured for a specific remediation module.

The Module Detail page displays all of the instances and remediations configured for a particular remediation module.

In a multidomain deployment, you can access the Module Detail page for remediation modules installed in the current domain and in ancestor domains. However, for modules installed in an ancestor domain, you cannot use the Module Detail page to add, delete, or edit instances in the current domain. Instead, use the Instances page (**Policies** > **Actions** > **Instances**); see Manage remediation instances, on page 1050.

Procedure

Step 1 Choose **Policies** > **Actions** > **Modules**.

Step 2 Next to the remediation module whose instances you want to manage, click **View** ( ).

Step 3 Manage your remediation instances:

- Add — To add an instance, click **Add**. For system-provided modules, see:
  - Add an ISE EPS instance, on page 1041
  - Add a Cisco IOS instance, on page 1044
  - *Cisco Secure Firewall Management Center Device Configuration Guide*
  - Add a set attribute value instance, on page 1048

  For help adding an instance for a custom module, see the documentation for that module, if available.
- Configure — To configure instance details and add remediations to the instance, click **View** ( ).

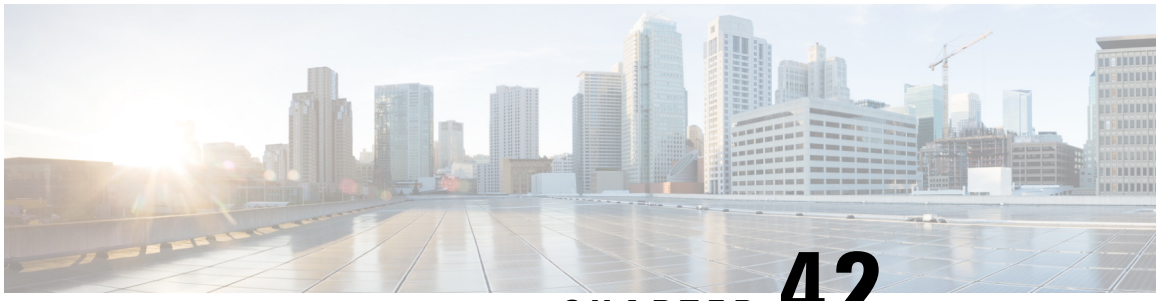
- Delete — To delete an instance that is not in use, click **Delete** ( ).



PART X

Reference

- Secure Firewall Management Center command line reference, on page 1055
- Communication Ports and Internet Access, on page 1063



CHAPTER 42

Secure Firewall Management Center command line reference

This reference explains the command line interface (CLI) for the Secure Firewall Management Center.

Note For Secure Firewall Threat Defense, see the *Cisco Secure Firewall Threat Defense Command Reference*.

- About the Secure Firewall Management Center CLI, on page 1055
- Secure Firewall Management Center CLI management commands , on page 1056
- Secure Firewall Management Center CLI show commands, on page 1057
- Secure Firewall Management Center CLI configuration commands, on page 1058
- Secure Firewall Management Center CLI system commands, on page 1058
- History for the Secure Firewall Management Center CLI, on page 1060

About the Secure Firewall Management Center CLI

The Secure Firewall Management Center CLI provides access to view and troubleshoot your Secure Firewall Management Center.

When you use SSH to log into the Firewall Management Center, you access the CLI. You can then access the Linux shell using the **expert** command, but this is strongly discouraged.

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

Caution Users with Linux shell access can obtain root privileges, which can present a security risk. For system security reasons, we strongly recommend:

- If you establish external authentication, make sure that you restrict the list of users with Linux shell access appropriately.
- Do not establish Linux shell users in addition to the pre-defined **admin** user.

You can use the commands described in this appendix to view and troubleshoot your Secure Firewall Management Center. You can also perform limited configuration operations.

Firewall Management Center CLI modes

The CLI encompasses four modes. The default mode, CLI Management, includes commands for navigating within the CLI itself. The remaining modes contain commands addressing three different areas of Firewall Management Center functionality; the commands within these modes begin with the mode name: `system`, `show`, or `configure`.

When you enter a mode, the CLI prompt changes to reflect the current mode. For example, to display version information about system components, you can enter the full command at the standard CLI prompt:

```
> show version
```

If you have previously entered **show** mode, you can enter the command without the **show** keyword at the show mode CLI prompt:

```
show> version
```

Secure Firewall Management Center CLI management commands

The CLI management commands provide the ability to interact with the CLI. These commands do not affect the operation of the device.

exit command

This command moves the CLI context up to the next highest CLI context level. Issuing this command from the default mode logs the user out of the current CLI session.

Syntax: `exit`

Example:

```
system> exit
>
```

expert command

This command invokes the Linux shell.

Syntax: `expert`

Example:

```
> expert
```

? (question mark) command

This command displays context-sensitive help for CLI commands and parameters. Use the question mark (`?`) command in these ways:

- To display help for the commands that are available within the current CLI context, enter a question mark (`?`) at the command prompt.
- To display a list of the available commands that start with a particular character set, enter the abbreviated command immediately followed by a question mark (`?`).
- To display help for a command's legal arguments, enter a question mark (`?`) in place of an argument at the command prompt.



Note The question mark (`?`) is not echoed back to the console.

Syntax:

```
?
abbreviated_command ?
command [arguments] ?
```

Example: `> ?`

Secure Firewall Management Center CLI show commands

Show commands provide information about the state of the appliance. These commands do not change the operational mode of the appliance and running them has minimal impact on system operation.

version command

Displays the product version and build, the UUID, and other information.

Syntax:

```
show version
```

Example:

```
> show version
-------------------[ fmc-austin ]--------------------
Model                     : Cisco Secure Firewall Management Center for VMware (66) Version
 7.6.0 (Build 1385)
UUID                      : a904b8b2-ca9a-11ee-a583-5e804c16b2fd
```

```
Rules update version      : 2024-05-13-001-vrt
LSP version               : lsp-rel-20240513-1955
VDB version               : 380
----------------------------------------------------
```

Secure Firewall Management Center CLI configuration commands

The configuration commands enable the user to configure and manage the system. These commands affect system operation.

password command

Allows the current CLI user to change their password.



Caution For system security reasons, we strongly recommend that you do not establish Linux shell users in addition to the pre-defined **admin** on any appliance.

After issuing the command, the CLI prompts the user for their current (or old) password, then prompts the user to enter the new password twice.

Syntax:

```
configure password
```

Example:

```
> configure password
Changing password for admin.
(current) UNIX password:
New UNIX password:
Retype new UNIX password:
passwd: password updated successfully
```

Secure Firewall Management Center CLI system commands

The system commands enable the user to manage system-wide files and access control settings.

Following sections describes the Secure Firewall Management Center CLI system commands:

generate-troubleshoot command

Generates troubleshooting data for analysis by Cisco.

Syntax:

```
system generate-troubleshoot option1 optionN
```

Where options are one or more of the following, space-separated:

- `SYS`: System Configuration, Policy, and Logs

- `PER`: Hardware Performance and Logs
- `NET`: Interface and Network Related Data
- `UPG`: Upgrade Data and Logs
- `SNT`: Snort Performance and Configuration
- `DES`: Detection Configuration, Policy, and Logs
- `VDB`: Discover, Awareness, VDB Data, and Logs
- `DBO`: All Database Data
- `LOG`: All Log Data
- `NMP`: Network Map Information
- `ALL`: Run all of the following options.

Example:

```
> system generate-troubleshoot VDB NMP
starting /usr/local/sf/bin/sf_troubleshoot.pl…
Please, be patient. This may take several minutes.
The troubleshoot options codes specified are VDB,NMP.
Getting filenames from [usr/local/sf/etc/db_updates/index]
Getting filenames from [usr/local/sf/etc/db_updates/base-6.2.3]
Troubleshooting information successfully created at
/var/common/results-06-14-2018—222027.tar.gz
```

lockdown command

Removes the `expert` command and access to the Linux shell on the device.



Caution This command is irreversible without a hotfix from Support. Use with care.

Syntax: `system lockdown`

Example:

```
> system lockdown
```

reboot command

Reboots the appliance.

Syntax: `system reboot`

Example:

```
> system reboot
```

restart command

Restarts the appliance application.

Syntax: `system restart`

Example:

```
> system restart
```

show max-logins command

Displays the maximum number of parallel logins for one user.

Syntax: `show max-logins`

Command default: The default value for the number of parallel/concurrent sessions for a single user is zero.

Example:

```
> show max-logins
```

configure max-logins command

Configures the maximum number of parallel logins for one user.

Syntax: `configure max-logins`

Example:

```
> configure max-logins [0-1024]
```

shutdown command

Shuts down the appliance.

Syntax: `system shutdown`

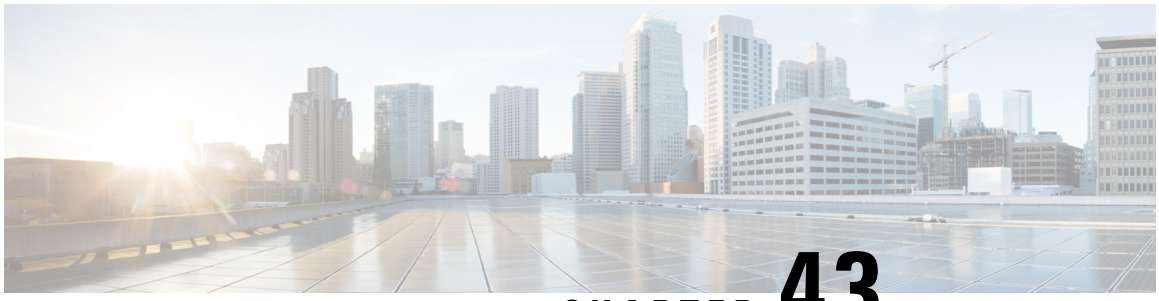
Example:

```
> system shutdown
```

History for the Secure Firewall Management Center CLI

This reference provides the chronological history of CLI features for the Secure Firewall Management Center, including version information, supported platforms, and implementation details.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Automatic CLI access for the Firewall Management Center	6.5	Any	When you use SSH to log into the Firewall Management Center, you automatically access the CLI. Although strongly discouraged, you can then use the CLI `expert` command to access the Linux shell. **Note** This feature deprecates the Version 6.3 ability to enable and disable CLI access for the Firewall Management Center. As a consequence of deprecating this option, the virtual Firewall Management Center no longer displays the **System** > **Configuration** > **Console Configuration** page, which still appears on physical Firewall Management Centers.
Ability to enable and disable CLI access for the Firewall Management Center	6.3	Any	New/Modified screens: A new check box available to administrators in Firewall Management Center web interface: **Enable CLI Access** on the **System** > **Configuration** > **Console Configuration** page. • Checked: Logging into the Firewall Management Center using SSH accesses the CLI. • Unchecked: Logging into Firewall Management Center using SSH provides access the Linux shell. This is the default state for fresh Version 6.3 installations as well as upgrades to Version 6.3 from a previous release. Supported platforms: Firewall Management Center
Firewall Management Center CLI	6.3	Any	Feature introduced. Initially supports the following commands: • `exit` • `expert` • `?` • `show version` • `configure password` • `system generate-troubleshoot` • `system lockdown` • `system reboot` • `system restart` • `system shutdown` Supported platforms: Firewall Management Center



CHAPTER 43

Communication Ports and Internet Access

- Guidelines for secure management communications, on page 1063
- Communication ports for Firewall Management Center, on page 1063
- Communication ports for managed devices, on page 1066
- Internet access requirements for Firewall Management Center, on page 1067
- Internet access requirements for managed devices, on page 1071
- Internet resources accessed by your browser, on page 1072

Guidelines for secure management communications

To safeguard the Firewall Management Center, you should install it on a protected internal network. Although the Firewall Management Center is configured to have only the necessary services and ports available, you must make sure that attacks cannot reach it (or any managed devices) from outside the firewall.

If the Firewall Management Center and its managed devices reside on the same network, you can connect the management interfaces on the devices to the same protected internal network as the Firewall Management Center. This allows you to securely control the devices from the Firewall Management Center. You can also configure multiple management interfaces to allow the Firewall Management Center to manage and isolate traffic from devices on other networks.

Regardless of how you deploy your appliances, inter-appliance communication is encrypted. However, you must still take steps to ensure that communications between appliances cannot be interrupted, blocked, or tampered with; for example, with a distributed denial of service (DDoS) or adversary-in-the-middle attack.

Communication ports for Firewall Management Center

The Firewall Management Center uses these ports to communicate. For deployments behind a network barrier—like an edge firewall—make sure you allow traffic on the required ports. Note that ports not required for essential or default operations remain closed until needed by a configuration or feature.

Inbound ports for Firewall Management Center

The Firewall Management Center accepts inbound traffic on these ports.

Table 145: Inbound ports for Firewall Management Center

Inbound port	Protocol/Feature	Details
22/tcp	SSH	Secure remote connections to the Firewall Management Center.
161/udp	SNMP	Allow access to MIBs via SNMP polling.
443/tcp	HTTPS	**Required.** Access the Firewall Management Center web interface.
443/tcp	HTTPS	Onboard an on-prem Firewall Management Center to Security Cloud Control with Secure Device Connector (on-prem).
443/tcp	HTTPS	Communicate with integrated and third-party products using the REST API.
443/tcp	HTTPS	Integrate with Secure Endpoint.
623/udp	SOL/LOM	Lights-Out Management (LOM) using a Serial Over LAN (SOL) connection.
1500/tcp 2000/tcp	Database access	Allow read-only access to the event database by a third-party client.
8302/tcp	eStreamer	Communicate with an eStreamer client.
8305/tcp	Appliance communications	**Required.** Securely communicate with managed devices. Also initiates connections on this port. Configurable. If you change this port, you must change it for *all* appliances in the deployment. We recommend you keep the default.
8307/tcp	Host input client	Communicate with a host input client.
8989/tcp	Cisco Support Diagnostics	Accepts authorized requests and transmits usage information and statistics. Also initiates connections on this port.

Outbound ports for Firewall Management Center

The Firewall Management Center initiates outbound traffic on these ports.

Table 146: Outbound ports for Firewall Management Center

Outbound port	Protocol/Feature	Details
7/udp 514/udp 6514/tcp	Syslog (audit logging)	Verify connectivity with the syslog server when configuring audit logging (7/udp). Send audit logs to a remote syslog server, when TLS is not configured (514/udp). Send audit logs to a remote syslog server, when TLS is configured (6514/tcp).
25/tcp	SMTP	Send email notices and alerts.

Outbound port	Protocol/Feature	Details
53/tcp 53/udp	DNS	**Required.** DNS
67/udp 68/udp	DHCP	DHCP
80/tcp	HTTP	Send and receive data from the internet. See Internet access requirements for Firewall Management Center, on page 1067.
80/tcp	HTTP	Download custom Security Intelligence feeds over HTTP.
80/tcp	HTTP	Download or query URL category and reputation data. This feature also uses 443/tcp.
80/tcp	HTTP	Display RSS feeds in the dashboard.
123/udp	NTP	Synchronize time.
162/udp	SNMP	Send SNMP alerts to a remote trap server.
389/tcp 636/tcp	LDAP	Communicate with an LDAP server for external authentication. Obtain metadata for detected LDAP users. Configurable.
443/tcp	HTTPS	Send and receive data from the internet. See Internet access requirements for Firewall Management Center, on page 1067.
443/tcp	HTTPS	Communicate with the Secure Malware Analytics Cloud (public or private).
443/tcp	HTTPS	Integrate with Secure Endpoint. Also accepts connections on this port.
443/tcp	HTTPS	Onboard an on-prem Firewall Management Center to Security Cloud Control with Cisco Security Cloud or Secure Device Connector (cloud).
1812/udp 1813/udp	RADIUS	Communicate with a RADIUS server for external authentication and accounting. Configurable.
5222/tcp	ISE	Communicate with an ISE identity source.
8305/tcp	Appliance communications	**Required.** Securely communicate with managed devices. Also accepts connections on this port. Configurable. If you change this port, you must change it for *all* appliances in the deployment. We recommend you keep the default.
8989/tcp	Cisco Support Diagnostics	Accepts authorized requests and transmits usage information and statistics. Also accepts connections on this port.
8989/tcp	Cisco Success Network	Transmit usage information and statistics.

Communication ports for managed devices

Managed devices use these ports to communicate. For deployments behind a network barrier—like an edge firewall—make sure you allow traffic on the required ports. Note that ports not required for essential or default operations remain closed until needed by a configuration or feature.

Inbound ports for managed devices

Managed devices accept inbound traffic on these ports.

Table 147: Inbound ports for managed devices

Inbound port	Protocol/Feature	Details
22/tcp	SSH	Secure remote connections to the device.
161/udp	SNMP	Allow access to MIBs via SNMP polling.
443/tcp	HTTPS	Communicate with integrated and third-party products using the REST API.
443/tcp	Remote access VPN (SSL/IPsec)	Allow secure VPN connections to your network from remote users.
500/udp 4500/udp	Remote access VPN (IKEv2)	Allow secure VPN connections to your network from remote users.
885/tcp	Captive portal	Communicate with a captive portal identity source.
8305/tcp	Appliance communications	**Required.** Securely communicate with the Firewall Management Center. Also initiates connections on this port. Configurable. If you change this port, you must change it for *all* appliances in the deployment. We recommend you keep the default.
8989/tcp	Cisco Support Diagnostics	Accepts authorized requests. Also initiates connections on this port.

Outbound ports for managed devices

Managed devices initiate outbound traffic on these ports. Managed devices also use ephemeral source ports for TCP and UDP traffic that they initiate. Make sure intervening network barriers allow response traffic for these connections.

Table 148: Outbound ports for managed devices

Outbound port	Protocol/Feature	Details
53/tcp 53/udp	DNS	DNS

Outbound port	Protocol/Feature	Details
67/udp 68/udp	DHCP	DHCP
123/udp	NTP	Synchronize time.
162/udp	SNMP	Send SNMP alerts to a remote trap server.
1812/udp 1813/udp	RADIUS	Communicate with a RADIUS server for external authentication and accounting. Configurable.
389/tcp 636/tcp	LDAP	Communicate with an LDAP server for external authentication. Configurable.
443/tcp	HTTPS	Send and receive data from the internet; see Internet access requirements for managed devices, on page 1071.
514/udp	Syslog (audit logging)	Send audit logs to a remote syslog server, when TLS is not configured.
8305/tcp	Appliance communications	Required. Securely communicate with the Firewall Management Center. Also accepts connections on this port. Configurable. If you change this port, you must change it for *all* appliances in the deployment. We recommend you keep the default.
8514/udp	Secure Network Analytics Manager	Send syslog messages to Secure Network Analytics using Security Analytics and Logging (On Premises).
8989/tcp	Cisco Support Diagnostics	Transmits usage information and statistics. Also accepts connections on this port.

Internet access requirements for Firewall Management Center

Internet access to the listed resources supports essential operations and configured features in your deployment.

The Firewall Management Center connects to the internet on ports 443/tcp (HTTPS) and 80/tcp (HTTP). You can configure a proxy server, except for NTP and whois. For some features, your location determines which resources the Firewall Management Center accesses. Some features also require managed devices to access the internet; refer to Internet access requirements for managed devices, on page 1071.

Threat intelligence

This table identifies the internet resources that the Firewall Management Center accesses for threat intelligence.

Table 149: Internet resources for threat intelligence

Feature	Reason	High availability	Resource
Intrusion rules for Snort 3	Download intrusion rules for Snort 3 (LSP).	Active peer downloads, syncs to standby.	updates-talos.sco.cisco.com updates-dyn-talos.sco.cisco.com
Intrusion rules for Snort 2	Download intrusion rules for Snort 2 (SRU).	Active peer downloads, syncs to standby.	talosintelligence.com support.sourcefire.com
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: • est.sco.cisco.com • regsvc.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

Feature	Reason	High availability	Resource
Malware Defense	Secure Malware Analytics 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
	Query for dynamic analysis results.	Both peers 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
Secure Endpoint	Receive malware events detected by Secure Endpoint from the cloud. Display malware events detected by the system in Secure Endpoint. Use centralized file Block and Allow lists created in Secure Endpoint to override dispositions from the cloud.	Both peers receive events. You must also configure the cloud connection on both peers (configuration is not synced).	Required Server Addresses for Proper Cisco Secure Endpoint & Malware Analytics Operations
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

Integrations

This table identifies the internet resources that the Firewall Management Center accesses for integrations.

Table 150: Internet resources for integrations

Feature	Reason	High availability	Resource
Cisco XDR	Send events to Cisco Security Cloud.	All units send events.	Cisco Secure Firewall Threat Defense and Cisco XDR Integration Guide

Support

This table identifies the internet resources that the Firewall Management Center accesses for support.

Table 151: Internet resources for support

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. Requires Version 7.2.4.	Each peer downloads its own certificates.	cisco.com/security/pki
Cisco Security Cloud/Security Services Exchange	Cisco Success Network Cisco Support Diagnostics	Both peers communicate.	api-sse.cisco.com api.eu.sse.itd.cisco.com api.apj.sse.itd.cisco.com api.au.sse.itd.cisco.com api.in.sse.itd.cisco.com
Licensing	Communicate with the Cisco Smart Software Manager.	Active peer communicates.	www.cisco.com 7.2.0–7.2.9: tools.cisco.com:443 7.2.10–7.2.x: smartreceiver.cisco.com
Upgrades	Download product (management center and device) upgrades.	Upgrade packages do not sync.	7.2.0–7.2.5: support.sourcefire.com 7.2.6–7.2.x: cdo-ftd-images.s3-us-west-2.amazonaws.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

Feature	Reason	High availability	Resource
RSS feeds	Display the Cisco Threat Research Blog on the dashboard.	Both peers communicate.	blog.talosintelligence.com blogs.cisco.com feeds.feedburner.com

Administration

This table identifies the internet resources that the Firewall Management Center accesses for administration.

Table 152: Internet resources for administration

Feature	Reason	High availability	Resource
Time synchronization (NTP)	Synchronize time in your deployment. Not supported with a proxy server.	Both peers communicate with the NTP server.	User configured
Whois	Request whois information for an external host. Not supported with a proxy server.	Any appliance requesting whois information must have internet access.	The whois client tries to guess the right server to query. If it cannot guess, it uses: • NIC handles: whois.networksolutions.com • IPv4 addresses and network names: whois.arin.net

Internet access requirements for managed devices

Internet access to the listed resources supports essential operations and configured features in your deployment.

Managed devices connect to the internet on ports 443/tcp (HTTPS) and 80/tcp (HTTP). You can configure a proxy server, except for NTP. For some features, your location determines which resources managed devices access. For deployments behind a network barrier—like an edge firewall—make sure you allow traffic on the required ports to the listed resources.

This table identifies the internet resources that managed devices access.

Table 153: Internet resources for managed devices

Feature	Reason	HA/Clustering	Resource
CA certificate bundles	Managed devices query for new CA certificates daily at a system-defined time. The local CA bundle contains certificates for accessing several Cisco services. Requires Version 7.2.4.	Each unit downloads its own certificates.	cisco.com/security/pki
Malware Defense	Submit files for dynamic analysis.	All units submit files.	fmc.api.threatgrid.com fmc.api.threatgrid.eu
Cisco Support Diagnostics	Accepts authorized requests and transmits usage information and statistics.	All units communicate.	api-sse.cisco.com:8989
Time synchronization	Synchronize time in your deployment. Not supported with a proxy server.	All units communicate with the NTP server.	User configured.

Internet resources accessed by your browser

When you use the Firewall Management Center, your browser may contact Google (google.com) or Amplitude (amplitude.com) web analytics servers to send non-personally-identifiable usage data to Cisco. For browser compatibility information, refer to the *Cisco Secure Firewall Management Center Compatibility Guide*.